

z/VM  
7.4

*Performance Toolkit Reference*



**Note:**

Before you use this information and the product it supports, read the information in [“Notices” on page 1087](#).

This edition applies to version 7, release 4 of IBM® z/VM® (product number 5741-A09) and to all subsequent releases and modifications until otherwise indicated in new editions.

Last updated: 2025-06-12

© **Copyright International Business Machines Corporation 2003, 2025.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

<b>Figures.....</b>	<b>xv</b>
<b>Tables.....</b>	<b>xxv</b>
<b>About this document.....</b>	<b>xxvii</b>
Intended Audience.....	xxvii
Where to Find More Information.....	xxvii
Links to Other Documents and Websites.....	xxvii
<b>How to provide feedback to IBM.....</b>	<b>xxix</b>
<b>Summary of Changes for z/VM: Performance Toolkit Reference.....</b>	<b>xxxi</b>
SC24-6303-74, z/VM 7.4 (June 2025) .....	xxxi
SC24-6303-74, z/VM 7.4 (September 2024).....	xxxi
SC24-6303-73, z/VM 7.3 (September 2023).....	xxxii
SC24-6303-73, z/VM 7.3 (August 2023).....	xxxii
SC24-6303-73, z/VM 7.3 (September 2022).....	xxxii
<b>Chapter 1. Performance Toolkit Subcommands.....</b>	<b>1</b>
Entering Performance Toolkit Subcommands from EXECs or CMS Programs.....	1
Return Codes.....	1
General Subcommands.....	3
Asterisk '*' .....	3
Equal sign '=' .....	4
Question mark '?' .....	5
BASIC.....	6
CANCEL.....	7
HELP.....	8
FCONTROL.....	10
FCONTROL ACTMSG.....	13
FCONTROL AUTOREFR.....	14
FCONTROL BENCHMRK.....	15
FCONTROL COLOR.....	18
FCONTROL DEFLOG.....	21
FCONTROL DEFSCRN.....	24
FCONTROL DEVGROUP.....	26
FCONTROL DUMPID.....	28
FCONTROL EMERGENC.....	29
FCONTROL FCXLEVEL.....	32
FCONTROL FORCEUSR.....	33
FCONTROL GDDMSPEC.....	38
FCONTROL HMA.....	40
FCONTROL LIMIT.....	42
FCONTROL MAINTID.....	46
FCONTROL MAXREC.....	47
FCONTROL MINPATHS.....	48
FCONTROL MODEL.....	50
FCONTROL MONCACHE.....	51
FCONTROL MONCOLL.....	53

FCONTROL MSGCLEAR.....	63
FCONTROL MSGWAIT.....	64
FCONTROL NUMBERS.....	65
FCONTROL PFKEY.....	66
FCONTROL PROCESS.....	70
FCONTROL RELOAD.....	75
FCONTROL RMTLINES.....	76
FCONTROL SCROLL.....	77
FCONTROL SEARCH.....	79
FCONTROL SECUSER.....	80
FCONTROL SETEVENT.....	81
FCONTROL SETTINGS.....	83
FCONTROL TIMESTAMP.....	88
FCONTROL UCLASS.....	89
FCONTROL UPDTCMS.....	90
FCONTROL UPDTSCRN.....	91
FCONTROL USERBOTL.....	92
FCONTROL USERHDR.....	93
FCONTROL USERVAR.....	94
FCONTROL USRLIMIT.....	97
Basic Mode Subcommands.....	100
CLEAR.....	101
CMS.....	102
CP.....	103
DELETE.....	104
FCONAPPC.....	105
FCONRMT.....	106
MODELEVEL.....	108
HISTDATA.....	109
MONITOR.....	110
MONSCAN.....	111
QUIT.....	113
REDISP.....	114
REPLY.....	115
RETURN.....	116
TRNDSCAN.....	117
'ENTER'.....	118
Redisplay Mode Subcommands.....	119
BACKWARD.....	120
BOTTOM.....	121
FORWARD.....	122
LEFT.....	123
RETURN.....	124
RIGHT.....	125
TOP.....	126
/ (LOCATE).....	127
'ENTER'.....	128
Performance Monitor Mode Subcommands.....	129
AGELLOG.....	130
AUXLOG.....	131
AVAILLOG.....	132
AVERAGE.....	133
AVLA2GLG.....	134
AVLB2GLG.....	135
BACKWARD.....	136
BASMODE.....	137
BENCHMRK.....	138
BFSFILES.....	139

BFSLOCKS.....	140
BFSPIPES.....	141
BOTTOM.....	142
CACHDBSE.....	143
CACHDLOG.....	144
CACHELOG.....	145
CACHEXT.....	146
CHANNEL.....	147
CMS.....	148
CORREL.....	149
CP.....	152
CPLACT.....	153
CPLCONF.....	154
CPLMENU.....	155
CPOWNLOG.....	156
CPSALL.....	157
CPSERV.....	158
CPSFAILS.....	159
CPSTOCP.....	160
CPSTOVM.....	161
CPU.....	162
CPUMENU.....	163
CTLUNIT.....	164
CURRENT.....	165
DASDLOG.....	166
DEMNDLOG.....	167
DEVCONF.....	168
DEVICE.....	169
DEVLOG.....	170
DEVMENU.....	171
DSPACESH.....	172
DSVBKACT.....	173
DSVCLOG.....	174
DSVSLOG.....	175
EVNIC.....	176
EVSWITCH.....	177
EXCEPT.....	178
FCHANNEL.....	179
FORWARD.....	180
GLONACT.....	181
GRAPHDET / PLOTDET.....	182
GRAPHICS.....	184
GRAPHSUM / PLOTSUM.....	185
GRAPHVAR / PLOTVAR.....	189
GVNIC.....	192
GVSWITCH.....	193
HIPSOCK.....	194
HISTDATA.....	195
HPALIAS.....	196
HPFLOG.....	197
HPSHARE.....	198
IDLEUSER.....	199
INTERIM.....	200
IOCHANGE.....	201
IOPQUEUE.....	202
IOPROCLG.....	203
ISFEACT.....	204
ISFECONF.....	205

ISFELOG .....	206
ISFLACT.....	207
ISFLALOG.....	208
ISFLCONF.....	209
ISFLLOG.....	210
LCHANNEL.....	211
LEFT.....	212
LGRDATA.....	213
LGRELOG.....	214
LINUX.....	215
LINUXMEN.....	216
LOCATE.....	217
LOCKACT.....	218
LOCKLOG.....	219
LPAR.....	220
LPARCONF.....	221
LPARLOG.....	222
LPARLOGM.....	223
LPARMENU.....	224
LSHARACT.....	225
LXCPU.....	226
LXCPULOG.....	227
LXFILSYS.....	228
LXMEM.....	229
LXMEMLOG.....	230
LXNETLOG.....	231
LXNETWRK.....	232
MDCACHE.....	233
MDCSTOR.....	234
MENU.....	235
MONDATA.....	236
MONSET.....	237
MTCORELG.....	238
MTCTYPLG.....	239
MTUSERS.....	240
MTUSRLOG.....	241
NEXTSAMP.....	242
NSS.....	243
PAGELOG.....	244
PCIACT.....	245
PCICONF.....	246
PCILOG.....	247
PCIMENU.....	248
PHYSLOG.....	249
PINLOG.....	250
PRCLOG.....	251
PRINT.....	252
PRINTRMT.....	255
PRIVOP.....	257
PROCCONF.....	258
PROCLOG.....	259
PROCMENU.....	260
PROCSUM.....	261
PUCFGLOG.....	262
PUORGLOG.....	263
QDIO.....	264
QDIOLOG.....	265
QEBSM.....	266

QEBSMLOG.....	267
REDHIST.....	268
REDISP.....	269
RESET.....	270
RESPALL.....	271
RESPUP.....	272
RETURN.....	273
RIGHT.....	274
RSKDET.....	275
RSKDISK.....	276
RSKLOG.....	277
RSKMENU.....	278
RSKSTOR.....	279
SCHDATA.....	280
SCHEDLOG.....	281
SCSI.....	282
SCSILOG.....	283
SEEKDET.....	284
SEEKDIST.....	285
SEEKDLOG.....	286
SEEKLLOG.....	287
SEEKLOC.....	288
SFS.....	289
SFSIOLOG.....	290
SFSLOG.....	291
SFSMENU.....	292
SFSREQ.....	293
SKIPSAMP.....	294
SORT.....	295
SPOOL.....	296
SSICONF.....	297
SSIMENU.....	298
SSISCHLG.....	299
SSISMILG.....	300
STEALLOG.....	301
STORAGE.....	302
STORLOG.....	303
STORMENU.....	304
SUBPLOG.....	305
SXSAVAIL.....	306
SXSDEFER.....	307
SXSPAGE.....	308
SXSUTIL.....	309
SYSCONF.....	310
SYSLOAD.....	311
SYSLOG.....	312
SYSMENU.....	313
SYSSET.....	314
SYSSUMLG.....	315
SYSTEM.....	316
SYSTRANS.....	317
TCPACTLG.....	318
TCPBPLOG.....	319
TCPCONF.....	320
TCPDATLG.....	321
TCPDOSLG.....	322
TCPICMP.....	323
TCPIOLOG.....	324

TCPLINKS.....	325
TCPMENU.....	326
TCPSESS.....	327
TCPUSERS.....	328
TOP.....	329
TOPOLOG.....	330
UCOMM.....	331
UCOMMLOG.....	332
UCONF.....	333
UDEFMENU.....	334
UPAGE.....	335
UPAGELOG.....	336
UPGACT.....	337
UPGACTLG.....	338
UPGMENU.....	339
UPGUTL.....	340
UPGUTLLG.....	341
UQDIO.....	342
UQDIOLOG.....	343
URESP.....	344
URESPLOG.....	345
USER.....	346
USERLOG.....	347
USRMPLOG.....	348
USRPRCTM.....	349
USRTMLOG.....	350
USTAT / USTATG.....	351
USTATLOG.....	352
USTLOG.....	353
USTMPLOG.....	354
UTRANDET.....	355
UTRANLOG.....	356
UTRANS.....	357
VDISKS.....	358
VMRM.....	359
VNIC.....	360
VOLLOG.....	361
VOLUME.....	362
VSWITCH.....	363
/ (LOCATE).....	364
'ENTER'.....	365

## **Chapter 2. Performance Toolkit Reports..... 367**

Report Layout Overview.....	367
Interval.....	367
Function Key Definitions.....	369
Data Consistency.....	369
Data Field Values.....	369
Data Completeness.....	370
FCX100, CPU Load and Transactions – CPU.....	371
FCX101/FCX181 Remote Performance Log – REDISP, REDHIST.....	377
Redisplay of General CPU Performance Data.....	378
Redisplay of User Transaction Data.....	379
Redisplay of Storage Utilization Data.....	381
Redisplay of User Extremes.....	383
FCX102, System Counters – SYSTEM.....	385
FCX103, Storage Utilization – STORAGE.....	392



FCX104, Privileged Operations Screen – PRIVOP.....	400
FCX105, Performance History Data Screens - GRAPHDET/PLOTDET, GRAPHSUM/PLOTSUM, GRAPHVAR/PLOTVAR .....	407
FCX106, CP IUCV Services Screen – CPSESV.....	408
FCX107, Channel Load Screen – CHANNEL.....	410
FCX108, General I/O Device Screen – DEVICE or DEVICE DASD.....	412
FCX109, CP-Owned Device Screen – DEVICE CPOWNED.....	416
FCX110, I/O Device Details Screen – DEVICE.....	420
FCX111, CU-Cached Disks Screen – CACHDBSE.....	428
FCX112, User Resource Usage – USER.....	431
FCX113, User Page Data Screen – UPAGE.....	435
FCX114, User Wait States Screen – USTAT/USTATG.....	438
FCX115, User Resource Details – USER <i>userid</i> .....	443
FCX116, Shared File System Server Screen – SFS.....	455
FCX117, Shared File System Server Details Screen – SFS.....	457
FCX119, SPOOL File Summary Screen – SPOOL or SPOOL NOSYSTEM.....	461
FCX120, SPOOL File User Summary Screen – SPOOL USER.....	464
FCX121, SPOOL File User Details Screen – SPOOL.....	466
FCX122, Remote CP Command Execution Screen – CP.....	469
FCX123, Remote CMS Command Execution Screen – CMS.....	470
FCX124, Performance Data Selection Menu Screen – MONITOR, MENU.....	471
FCX126, LPAR Load Screen – LPAR.....	472
FCX127, User Response Time Screen – URESP.....	475
FCX128, Graphics Selection Menu Screen – GRAPHICS.....	477
FCX129, User Resources per Transaction Screen – UTRANS.....	478
FCX130, System Counters Screen – SYSTRANS.....	480
FCX131, I/O Device Configuration Screen – DEVCONF.....	482
FCX132, User Communication Screen – UCOMM.....	484
FCX133, Shared Segments Screen – NSS.....	487
FCX134, Shared Data Spaces Screen – DSPACESH.....	489
FCX135, User Wait State Log Screen – USTLOG.....	491
FCX136, Response Time Log (All) Screen – RESPALL.....	494
FCX137, Response Time Log (UP) Screen – RESPUP.....	496
FCX138, Minidisk Cache Log Screen – MDCACHE.....	498
FCX139/140/141/142, CP IUCV Services Log Screen – CPSALL, CPSTOCP, CPSTOVM, CPSFAIL.....	500
FCX143, Paging Log Screen – PAGELOG.....	502
FCX144, Processor Log Screen – PROCLOG.....	504
FCX145, Scheduler Queue Log Screen – SCHEDLOG.....	507
FCX146, Auxiliary Storage Log Screen – AUXLOG.....	509
FCX147, Virtual Disks in Storage Screen – VDISKS.....	511
FCX148, User-Defined Performance Screen – FCONTROL DEFSCRN.....	513
FCX149, Monitor Settings – MONSET.....	515
FCX150, Shared File System Log Screen – SFSLOG.....	518
FCX151, Shared File System I/O Log Screen – SFSIOLOG.....	520
FCX152, Shared File System Requests Log Screen – SFSREQ.....	522
FCX153, Exception Log Screen – EXCEPT.....	524
FCX154, System Settings Screen – SYSSET.....	526
FCX155, Monitor Data Statistics Screen – MONDATA.....	530
FCX156, DASD Seeks Locations Screen – SEEKLOC.....	532
FCX157, DASD Seeks Distances Screen – SEEKDIST.....	534
FCX158, DASD Seeks Details Screen – SEEKDET.....	536
FCX160, History Data Selection Menu Screen – HISTDATA.....	539
FCX161, LPAR Channel Load Screen – LCHANNEL.....	540
FCX162, User Resource Usage Log – USERLOG <i>userid</i> .....	542
FCX163, User Page Data Log Screen – UPAGELOG.....	543
FCX164, User Wait States Log Screen – USTATLOG.....	544
FCX165, User Response Time Log Screen – URESPLOG.....	545
FCX166, User Resources per Transaction Log Screen – UTRANLOG.....	546

FCX167, User Communication Log Screen – UCOMMLOG.....	547
FCX168, General I/O Device Data Log Screen – DEVLOG.....	548
FCX169, Control Unit Cache Log Screen – CACHDLOG.....	552
FCX170, CPOWNER Device Log Screen – CPOWNLOG.....	553
FCX171, DASD Seeks Locations Log Screen – SEEKLLOG.....	556
FCX172, DASD Seeks Distances Log Screen – SEEKDLOG.....	557
FCX173, Benchmark Log Selection Menu Screen – BENCHMRK.....	558
FCX174, User Transaction Details Screen – UTRANDET.....	559
FCX175, Scheduler Data Log Screen – SCHDATA.....	564
FCX176, Cached Control Units Screen – CTLUNIT.....	566
FCX177, Cache Extended Functions Screen – CACHEXT.....	569
FCX178, Minidisk Cache Storage Log Screen – MDCSTOR.....	573
FCX179, System Facilities Log Screen – SYSLOG.....	575
FCX180, System Configuration – SYSCONF.....	577
FCX182, General User Communication Log Screen – UCOMMLOG.....	585
FCX183, DASD Performance Log Screen – DASDLOG.....	586
FCX184, Cache Extended Functions Log Screen – CACHELOG.....	587
FCX185, I/O Configuration Changes Screen – IOCHANGE.....	590
FCX186, Remote Basic Mode Screen – BASMODE.....	593
FCX187, User-Defined Performance Data Logs Screen – FC DEFLOG.....	594
FCX188, Multitasking Users Screen – MTUSERS.....	596
FCX189, Multitasking User Log Screen – MTUSRLOG.....	598
FCX190, Byte File System File Activity Log Screen – BFSFILES.....	599
FCX191, Byte File System PIPEs and Linking Activity Log Screen – BFSPIPES.....	601
FCX192, Byte File System Locking Activity Log Screen – BFSLOCKS.....	604
FCX193, System Selection Menu Screen – SYSMENU.....	606
FCX194, Variable Correlation Screen – CORREL.....	607
FCX195, Extended History Data Log Screen – REDHIST.....	609
FCX198, System Load Overview Screen – FCONRMT, SYSLOAD.....	629
FCX199, Remote History Data Screen.....	630
FCX201, FCONXREP Summary History Files Screen - FCONXREP.....	631
FCX202, LPAR Load Log – LPARLOG.....	632
FCX203, TCP/IP Performance Data Menu Screen – TCPMENU.....	636
FCX204, General TCP/IP Activity Log Screen – TCPACTLG.....	638
FCX205, General TCP/IP Data Transfer Log Screen – TCPDATLG.....	640
FCX206, TCP/IP ICMP Messages Log Screen – TCPICMP.....	642
FCX207, TCP/IP TCP and UDP Sessions Log Screen – TCPSESS.....	644
FCX208, TCP/IP Links Activity Log Screen – TCPLINKS.....	646
FCX210, TCP/IP Buffer Pools Log Screen – TCPBPLOG.....	648
FCX212, TCP/IP Server Configuration Screen – TCPCONF.....	650
FCX213, User-Defined Displays Selection Menu Screen – UDEFMENU.....	654
FCX214, TCP/IP TCP and UDP Sessions Log Screen – TCPUSERS.....	655
FCX215, FICON Channel Load Screen – FCHANNEL.....	657
FCX216, Reusable Server Kernel Performance Data Menu Screen – RSKMENU.....	659
FCX217, Reusable Server Kernel Performance Details Screen – RSKDET.....	660
FCX218, Reusable Server Kernel Performance Log Screen – RSKLOG.....	665
FCX219, Reusable Server Kernel Disk Performance Log Screen – RSKDISK.....	667
FCX220, Reusable Server Kernel Storage Log Screen – RSKSTOR.....	669
FCX221, SFS and BFS Logs Selection Menu Screen – SFSMENU.....	671
FCX222, TCP/IP I/O Activity Log Screen – TCPIOLOG.....	672
FCX223, Linux Systems Selection Screen – LINUX.....	674
FCX224, Linux Details Selection Screen – LINUX.....	675
FCX225, System Summary Log Screen – SYSSUMLG.....	677
FCX226, User Configuration – UCONF.....	679
FCX227, Linux Network Activity Screen – LXNETWRK.....	682
FCX228, Linux Filesystem Usage Screen – LXFILSYS.....	684
FCX229, Linux Memory Utilization and Activity Screen – LXMEM.....	686
FCX230, Linux CPU Utilization Screen – LXCPU.....	689

FCX231, Hipersocket Channel Activity Screen – HIPSOCK.....	691
FCX232, I/O Processor Log Screen – IOPROCLG.....	693
FCX233, TCP/IP Denial of Service Log Screen – TCPDOSLG.....	695
FCX234, Processor/Core Configuration Log Screen – PROCCONF.....	697
FCX235, LPAR Configuration Log Screen – LPARCONF.....	699
FCX236, Processor Load and Configuration Logs Menu – PROCMENU.....	701
FCX237, I/O Priority Queueing Screen – IOPQUEUE.....	702
FCX238, Idle Users Screen – IDLEUSER.....	704
FCX239, Processor Summary Log Screen – PROCSUM.....	705
FCX240, Virtual Switch Activity Screen – VSWITCH.....	707
FCX241, VM Resource Manager Screen – VMRM.....	709
FCX242, Linux Data Selection Menu Screen – LINUXMEN.....	711
FCX243, General Linux CPU Utilization Screen – LXCPU.....	712
FCX244, General Linux Memory Utilization Screen – LXMEN.....	714
FCX245, General Linux Networking Activity Screen – LXNETWRK.....	716
FCX246, Linux CPU Load Log Screen – LXCPULOG.....	718
FCX247, Linux Memory Utilization Log Screen – LXMENLOG.....	719
FCX248, Linux Network Activity Log Screen – LXNETLOG.....	720
FCX249, SCSI Device – SCSI.....	721
FCX250, SCSI Resource Utilization Log – SCSILOG.....	722
FCX251, QDIO Activity Screen – QDIO.....	723
FCX252, QDIO Activity Log Screen – QDIOLOG.....	725
FCX253, Storage Utilization Log Screen – STORLOG.....	726
FCX254, Available List Log Screen – AVAILLOG.....	728
FCX255, QEBSM Activity Screen – QEBSM.....	730
FCX256, QEBSM Activity Log Screen – QEBSMLOG.....	732
FCX257, User QDIO Activity Screen – UQDIO.....	733
FCX258, User QDIO Activity Log Screen – UQDIOLOG.....	735
FCX259, Demand Scan Log Screen – DEMNDLOG.....	736
FCX260, Storage Management Logs Menu Screen – STORMENU.....	739
FCX261, SXS Available Page Queues Management Log Screen – SXSAVAIL.....	740
FCX262, SXS Page Management Log Screen – SXSPAGE.....	742
FCX263, SXS Deferred Tasks Management Log Screen – SXSDEFER.....	744
FCX264, SXS Storage Utilization Log Screen – SXSUTIL.....	746
FCX265, Spin Lock Log Screen – LOCKLOG.....	748
FCX266, General Virtual Switch Description Screen – GVSWITCH.....	750
FCX267, Extended Virtual Switch Activity Screen – EVSWITCH.....	753
FCX268, General Virtual Network Device Description Screen – GVNIC.....	755
FCX269, Virtual Network Device Activity Screen – VNIC.....	757
FCX270, Extended Virtual Network Device Activity Screen – EVNIC.....	759
FCX271, SSI Data Menu Screen – SSIMENU.....	760
FCX272, ISFC End Point Configuration Screen – ISFECONF.....	761
FCX273, ISFC End Point Activity Screen – ISFEACT.....	764
FCX274, ISFC Logical Link Activity Screen – ISFLACT.....	766
FCX275, ISFC Logical Link Configuration Screen – ISFLCONF.....	768
FCX276, SSI Configuration Screen – SSICONF.....	770
FCX277, SSI State Change Synchronization Activity Log Screen – SSISCHLG.....	772
FCX278, SSI State/Mode Information Log Screen – SSISMILG.....	773
FCX279, ISFC End Point Activity Log Screen – ISFELOG.....	775
FCX280, ISFC Logical Link Activity Log Screen – ISFLLOG.....	776
FCX281, ISFC Logical Link Activity By-Time Log Screen – ISFLALOG.....	777
FCX282, HPF I/O Device Screen – DEVICE HPF.....	778
FCX283, HPF I/O Device Performance Log Screen – HPFLOG.....	782
FCX284, I/O Device Data Selection Menu Screen – DEVMENU.....	785
FCX285, LGR Event Log Screen – LGRELOG.....	786
FCX286, LGR Data Screen – LGRDATA.....	789
FCX287, System Topology Machine Organization Screen – TOPOLOG.....	792
FCX288, Multiprocessor User Activity Log Screen – USRMPLOG.....	794

FCX289, User Paging Menu Screen – UPGMENU.....	797
FCX290, User Page Activity Screen – UPGACT.....	798
FCX291, User Page Activity Log Screen – UPGACTLG.....	800
FCX292, User Page Utilization Data Screen – UPGUTL.....	801
FCX293, User Page Utilization Data Log Screen – UPGUTLLG.....	803
FCX294, Available List Data Below 2G Screen – AVLB2GLG .....	804
FCX295, Available List Data Above 2G Screen – AVLA2GLG .....	806
FCX296, Steal Statistics Screen – STEALLOG.....	807
FCX297, Age List Log Screen – AGELLOG.....	809
FCX298, Logical Core Organization Log Screen – PUORGLOG.....	811
FCX299, Processor Unit (Core and Threads) Configuration Log Screen – PUCFGLOG.....	813
FCX300, Dispatch Vector Configuration Change Log Screen – DSVCLG.....	816
FCX301, Dispatch Vector Activity Screen – DSVBKACT.....	818
FCX302, Real Core Utilization Log Screen – PHYSLOG.....	820
FCX303, DSVBK Steals per Processor Log Screen – DSVSLOG.....	822
FCX304, Processor Log Screen – PRCLOG.....	825
FCX305, Logical Partition Activity Menu Screen – LPARMENU.....	828
FCX306, Logical Partition Share Screen – LSHARACT.....	829
FCX307, Logical Partition Logs Menu Screen – LPARLOGM.....	832
FCX308, CPU Pool Configuration Screen – CPLCONF.....	833
FCX309, CPU Pool Activity Screen – CPLACT.....	835
FCX310, PCI Function Menu Screen – PCIMENU.....	837
FCX311, PCI Function Configuration Screen – PCICONF.....	838
FCX312, PCI Function Activity Screen, Format 0 – PCIACT.....	840
FCX313, PCI Function Activity Log Screen, Format 0 – PCILOG.....	842
FCX314, Pinned Storage Log Screen – PINLOG.....	843
FCX315, Multiprocessor User Wait States Log Screen – USTMPLOG.....	845
FCX316, Subpool Storage Log Screen – SUBPLOG.....	848
FCX317, Global Networking Object Activity Screen – GLONACT.....	849
FCX318, PCI Function Activity Screen, Format 1 – PCIACT ROCE.....	851
FCX319, PCI Function Activity Log Screen, Format 1 – PCILOG.....	852
FCX320, PCI Function Activity Screen, Format 2 – PCIACT ZEDC.....	853
FCX321, PCI Function Activity Log Screen, Format 2 – PCILOG.....	854
FCX322, PCI Function Activity Screen, Format 3 – PCIACT ISM.....	855
FCX323, PCI Function Activity Log Screen, Format 3 – PCILOG.....	856
FCX324, CPU Pool Menu Screen – CPLMENU.....	857
FCX325, CPU Activity Selection Menu Screen – CPUMENU.....	858
FCX326, Spin Lock Collision Activity Screen – LOCKACT.....	859
FCX327, HyperPAV Alias Activity Screen – HPALIAS.....	861
FCX328, HyperPAV Share Activity Screen – HPSHARE .....	863
FCX329, DASD Volume Performance Screen – VOLUME.....	865
FCX330, DASD Volume Performance Log Screen – VOLLOG.....	868
FCX331, SMT Metrics Per Core Type Log Screen – MTCTYPLG.....	871
FCX332, SMT Metrics Per Core Log Screen – MTCORELG.....	873
FCX333, User Processor Time Screen – USRPRCTM.....	875
FCX334, User Processor Time Log Screen – USRTMLOG.....	877

## **Chapter 3. Performance Toolkit Messages..... 879**

## **Appendix A. Graphical Displays of Performance History Data.....955**

Graphics Selection Menu.....	958
Detailed Performance History Display.....	961
Example Using GDDM Graphics.....	962
Example Using Simple Character Plot.....	962
Summary Performance History Display.....	963
Performance Trend Display.....	964
Daily Performance Profile.....	965

Variable Correlation Display.....	966
Variable Correlation Display with Simple Plot.....	966
Variable Correlation Display Using GDDM Graphics.....	968
More about graphics.....	968
Definition of 'User Variables'.....	968
Generation of Cumulative Graphics.....	968
Example of Cumulative Graphics with User Variables.....	969
<b>Appendix B. Numerical Analysis of Performance History Data.....</b>	<b>971</b>
Numerical Analysis Using FCONXREP Procedure.....	971
Interactive Use of FCONXREP Procedure.....	971
Using FCONXREP with a Detailed History Log.....	973
Batch Use of FCONXREP Procedure.....	974
Sample Output.....	975
<b>Appendix C. Printed Reports.....</b>	<b>979</b>
Table of contents for printed reports.....	979
Layout of print file header:.....	979
<b>Appendix D. HISTLOG and HISTSUM File Records.....</b>	<b>981</b>
<b>Appendix E. Extended Trend File Records.....</b>	<b>983</b>
Format of Physical Records.....	984
<b>Appendix F. Extended Summary File Records.....</b>	<b>985</b>
Summary Record Headers.....	985
<b>Appendix G. Trend File Variables Selection.....</b>	<b>987</b>
Performance Variables of Trend Records X'FC01'.....	989
Performance Variables of Trend Records X'FC02'.....	1014
Performance Variables of Trend Records X'FC03'.....	1020
Performance Variables of Trend Records X'FC04'.....	1022
Performance Variables of Trend Records X'FC05'.....	1025
Performance Variables of Trend Records X'FC06'.....	1032
Performance Variables of Trend Records X'FC07'.....	1033
Performance Variables of Trend Records X'FC08'.....	1034
Performance Variables of Trend Records X'FC3A'.....	1037
Performance Variables of Trend Records X'FC3C'.....	1040
Performance Variables of Trend Records X'FC41', X'FC42' & X'FC43'.....	1042
Performance Variables of Trend Records X'FC44', X'FC45' & X'FC46'.....	1048
Performance Variables of Trend Records X'FC51'.....	1050
Performance Variables of Trend Records X'FC61'.....	1051
Performance Variables of Trend Records X'FC65'.....	1053
Performance Variables of Trend Records X'FC68'.....	1056
Performance Variables of Trend Records X'FC6D'.....	1058
Performance Variables of Trend Records X'FC6F'.....	1061
Performance Variables of Trend Records X'FC71'.....	1062
Performance Variables of Trend Records X'FCA2'.....	1065
Performance Variables of Trend Records X'FCA4'.....	1071
Performance Variables of Trend Records X'FCA6'.....	1072
Performance Variables of Trend Records X'FCA8'.....	1080
Performance Variables of Trend Records X'FCA9'.....	1084
<b>Notices.....</b>	<b>1087</b>
IBM Online Privacy Statement.....	1088
Trademarks.....	1088

<b>Bibliography.....</b>	<b>1091</b>
Where to Get z/VM Information.....	1091
z/VM Base Library.....	1091
z/VM Facilities and Features.....	1092
Prerequisite Products.....	1094
Related Products.....	1094
 <b>Index.....</b>	 <b>1095</b>

---

# Figures

1. General Layout of Performance Monitor Reports.....	367
2. Layout of CPU Load and Transactions report (FCX100).....	371
3. Layout of Remote Performance Log Report (FCX101).....	377
4. Layout of performance redisplay screen: CPU data.....	378
5. Layout of performance redisplay screen: Transaction data.....	380
6. Layout of performance redisplay screen: Storage Utilization data.....	382
7. Layout of performance redisplay screen: Max. User data.....	384
8. Layout of System Counters Report (FCX102).....	385
9. Layout of Storage Utilization Report (FCX103).....	392
10. Layout of Privileged Operations Screen (FCX104).....	400
11. Layout of CP IUCV Services Screen (FCX106).....	408
12. Layout of Channel Load Screen (FCX107).....	410
13. Layout of General I/O Device Screen (FCX108).....	412
14. Layout of CP-Owned Device Screen (FCX109).....	416
15. Layout of I/O Device Details Screen (FCX110).....	420
16. Layout of CU-Cached Disks Screen (FCX111).....	428
17. Layout of User Resource Usage report (FCX112).....	431
18. Layout of User Page Data Screen (FCX113).....	435
19. Layout of User Wait States Screen (FCX114).....	438
20. Layout of Graphic User Wait State Display.....	441
21. Layout of User Resource Details report (FCX115).....	443
22. Layout of Shared File System Server Screen (FCX116).....	455
23. Layout of Shared File System Server Details Screen (FCX117).....	457

24. Layout of SPOOL File Summary Screen (FCX119).....	461
25. Layout of SPOOL File User Summary Screen (FCX120).....	464
26. Layout of SPOOL File User Details Screen (FCX121).....	466
27. Layout of Remote CP Command Execution Screen (FCX122).....	469
28. Layout of Remote CMS Command Execution Screen (FCX123).....	470
29. Layout of Performance Data Selection Menu Screen (FCX124).....	471
30. Layout of LPAR Load Screen (FCX126).....	472
31. Layout of User Response Time Screen (FCX127).....	475
32. Layout of User Resources per Transaction Screen (FCX129).....	478
33. Layout of System Transaction Statistics Screen (FCX130).....	480
34. Layout of I/O Device Configuration Screen (FCX131).....	482
35. Layout of User Communication Screen (FCX132).....	484
36. Layout of Shared Segments Screen (FCX133).....	487
37. Layout of Shared Data Spaces Screen (FCX134).....	489
38. Layout of User Wait State Log Screen (FCX135).....	491
39. Layout of Response Time Log (All) Screen (FCX136).....	494
40. Layout of Response Time Log (UP) Screen (FCX137).....	496
41. Layout of Minidisk Cache Log Screen (FCX138).....	498
42. Layout of CP IUCV Services Log Screen (FCX139/140/141/142).....	500
43. Layout of Paging Log Screen (FCX143).....	502
44. Layout of Processor Log Screen (FCX144).....	504
45. Layout of Scheduler Queue Log Screen (FCX145).....	507
46. Layout of Auxiliary Storage Log Screen (FCX146).....	509
47. Layout of Virtual Disks in Storage Screen (FCX147).....	511
48. Example of a User-Defined Performance Screen (FCX148).....	513



49. Layout of Monitor Settings report (FCX149).....	515
50. Layout of Shared File System Log Screen (FCX150).....	518
51. Layout of Shared File System I/O Log Screen (FCX151).....	520
52. Layout of Shared File System Requests Log Screen (FCX152).....	522
53. Layout of Exception Log Screen (FCX153).....	524
54. Layout of System Settings Screen (FCX154).....	526
55. Layout of Monitor Data Statistics Screen (FCX155).....	530
56. Layout of DASD Seeks Locations Screen (FCX156).....	532
57. Layout of DASD Seeks Distances Screen (FCX157).....	534
58. Layout of DASD Seeks Details Screen (FCX158).....	536
59. Layout of History Data Selection Menu Screen (FCX160).....	539
60. Layout of LPAR Channel Load Screen (FCX161).....	540
61. Layout of User Resource Usage Log report (FCX162).....	542
62. Layout of User Page Data Log Screen (FCX163).....	543
63. Layout of User Wait States Log Screen (FCX164).....	544
64. Layout of User Response Time Log Screen (FCX165).....	545
65. Layout of User Resources per Transaction Log Screen (FCX166).....	546
66. Layout of User Communication Log Screen (FCX167).....	547
67. Layout of General I/O Device Data Log Screen (FCX168) - Example 1.....	548
68. Layout of General I/O Device Data Log Screen (FCX168) - Example 2.....	548
69. Layout of Control Unit Cache Log Screen (FCX169).....	552
70. Layout of CPOWNED Device Log Screen (FCX170).....	553
71. Layout of DASD Seeks Locations Log Screen (FCX171).....	556
72. Layout of DASD Seeks Distances Log Screen (FCX172).....	557
73. Layout of Benchmark Log Selection Menu Screen (FCX173).....	558

74. Layout of User Transaction Details Screen (FCX174).....	559
75. Layout of Scheduler Data Log Screen (FCX175).....	564
76. Layout of Cached Control Units Screen (FCX176).....	566
77. Layout of Cache Extended Functions Screen (FCX177).....	569
78. Layout of Minidisk Cache Storage Log Screen (FCX178).....	573
79. Layout of System Facilities Log Screen (FCX179).....	575
80. Layout of System Configuration report (FCX180).....	578
81. Layout of General User Communication Log Screen (FCX182).....	585
82. Layout of DASD Performance Log Screen (FCX183).....	586
83. Layout of Cache Extended Functions Log Screen (FCX184).....	587
84. Layout of I/O Configuration Changes Screen (FCX185).....	590
85. Layout of Remote Basic Mode Screen (FCX186).....	593
86. A Sample User-Defined Performance Data Logs Screen (FCX187).....	594
87. Layout of Multitasking Users Screen (FCX188).....	596
88. Layout of Multitasking User Log Screen (FCX189).....	598
89. Layout of Byte File System File Activity Log Screen (FCX190).....	599
90. Layout of Byte File System PIPEs and Linking Activity Log Screen (FCX191).....	601
91. Layout of Byte File System Locking Activity Log Screen (FCX192).....	604
92. Layout of System Selection Menu Screen (FCX193).....	606
93. Layout of Variable Correlation Screen (FCX194).....	607
94. Layout of Extended History Data Log Screen (FCX195).....	609
95. Layout of System Load Overview Screen (FCX198).....	629
96. Layout of Remote History Data Screen (FCX199).....	630
97. Layout of FCONXREP Summary History Files Screen (FCX201).....	631
98. Layout of LPAR Load Log report (FCX202).....	632

99. Layout of TCP/IP Performance Data Menu Screen (FCX203).....	636
100. Layout of General TCP/IP Activity Log Screen (FCX204).....	638
101. Layout of General TCP/IP Data Transfer Log Screen (FCX205).....	640
102. Layout of TCP/IP ICMP Messages Log Screen (FCX206).....	642
103. Layout of TCP/IP TCP and UDP Sessions Log Screen (FCX207).....	644
104. Layout of TCP/IP Links Activity Log Screen (FCX208).....	646
105. Layout of TCP/IP Buffer Pools Log Screen (FCX210).....	648
106. Layout of TCP/IP Server Configuration Screen (FCX212).....	650
107. Layout of User-Defined Displays Selection Menu Screen (FCX213).....	654
108. Layout of TCP/IP Users Activity Log Screen (FCX214).....	655
109. Layout of FICON Channel Load Screen (FCX215).....	657
110. Layout of Reusable Server Kernel Performance Data Menu Screen (FCX216).....	659
111. Layout of Reusable Server Kernel Performance Details Screen (FCX217).....	660
112. Layout of Reusable Server Kernel Performance Log Screen (FCX218).....	665
113. Layout of Reusable Server Kernel Disk Performance Log Screen (FCX219).....	667
114. Layout of Reusable Server Kernel Storage Log Screen (FCX220).....	669
115. Layout of SFS and BFS Logs Selection Menu Screen (FCX221).....	671
116. Layout of TCP/IP I/O Activity Log Screen (FCX222).....	672
117. Layout of Linux Systems Selection Screen (FCX223).....	674
118. Layout of Linux Details Selection Screen (FCX224).....	675
119. Layout of System Summary Log Screen (FCX225).....	677
120. Layout of User Configuration report (FCX226).....	679
121. Layout of Linux Network Activity Screen (FCX227).....	682
122. Layout of Linux Filesystem Usage Screen (FCX228).....	684
123. Layout of Linux Memory Utilization and Activity Screen (FCX229).....	686

124. Layout of Linux CPU Utilization Screen (FCX230).....	689
125. Layout of Hipersocket Channel Activity Screen (FCX231).....	691
126. Layout of I/O Processor Log Screen (FCX232).....	693
127. Layout of TCP/IP Denial of Service Log Screen (FCX233).....	695
128. Layout of Processor/Core Configuration Log Screen (FCX234).....	697
129. Layout of LPAR Configuration Log Screen (FCX235).....	699
130. Layout of Processor Load and Configuration Logs Menu (FCX236).....	701
131. Layout of I/O Priority Queueing Screen (FCX237).....	702
132. Layout of Idle Users Screen (FCX238).....	704
133. Layout of Processor Summary Log Screen (FCX239).....	705
134. Layout of Virtual Switch Activity Screen (FCX240).....	707
135. Layout of VM Resource Manager Screen (FCX241).....	709
136. Layout of Linux Data Selection Menu Screen (FCX242).....	711
137. Layout of General Linux CPU Utilization Screen (FCX243).....	712
138. Layout of General Linux Memory Utilization Screen (FCX244).....	714
139. Layout of General Linux Networking Activity Screen (FCX245).....	716
140. Layout of Linux CPU Load Log Screen (FCX246).....	718
141. Layout of Linux Memory Utilization Log Screen (FCX247).....	719
142. Layout of Linux Network Activity Log Screen (FCX248).....	720
143. Layout of SCSI Device report (FCX249).....	721
144. Layout of SCSI Resource Utilization Log report (FCX250).....	722
145. Layout of QDIO Activity Screen (FCX251).....	723
146. Layout of QDIO Activity Log Screen (FCX252).....	725
147. Layout of Storage Utilization Log Screen (FCX253).....	726
148. Layout of Available List Log Screen (FCX254).....	728

149. Layout of QEBSM Activity Screen (FCX255).....	730
150. Layout of QEBSM Activity Log Screen (FCX256).....	732
151. Layout of User QDIO Activity Screen (FCX257).....	733
152. Layout of User QDIO Activity Log Screen (FCX258).....	735
153. Layout of Demand Scan Log Screen (FCX259).....	736
154. Layout of Storage Management Logs Menu Screen (FCX260).....	739
155. Layout of SXS Available Page Queues Management Log Screen (FCX261).....	740
156. Layout of SXS Page Management Log Screen (FCX262).....	742
157. Layout of SXS Deferred Tasks Management Log Screen (FCX263).....	744
158. Layout of SXS Storage Utilization Log Screen (FCX264).....	746
159. Layout of Spin Lock Log Screen (FCX265).....	748
160. Layout of General Virtual Switch Description Screen (FCX266).....	750
161. Layout of Extended Virtual Switch Activity Screen (FCX267).....	753
162. Layout of General Virtual Network Device Description Screen (FCX268).....	755
163. Layout of Virtual Network Device Activity Screen (FCX269).....	757
164. Layout of Extended Virtual Network Device Activity Screen (FCX270).....	759
165. Layout of SSI Data Menu Screen (FCX271).....	760
166. Layout of ISFC End Point Configuration Screen (FCX272).....	761
167. Layout of ISFC End Point Activity Screen (FCX273).....	764
168. Layout of ISFC Logical Link Activity Screen (FCX274).....	766
169. Layout of ISFC Logical Link Configuration Screen (FCX275).....	768
170. Layout of SSI Configuration Screen (FCX276).....	770
171. Layout of SSI State Change Synchronization Activity Log Screen (FCX277).....	772
172. Layout of SSI State/Mode Information Log Screen (FCX278).....	773
173. Layout of ISFC End Point Activity Log Screen (FCX279).....	775

174. Layout of ISFC Logical Link Activity Log Screen (FCX280).....	776
175. Layout of ISFC Logical Link Activity By-Time Log Screen (FCX281).....	777
176. Layout of HPF I/O Device Screen (FCX282).....	778
177. Layout of HPF I/O Device Performance Log Screen (FCX283).....	782
178. Layout of I/O Device Data Selection Menu Screen (FCX284).....	785
179. Layout of LGR Event Log Screen (FCX285).....	786
180. Layout of LGR Data Screen (FCX286).....	789
181. Layout of System Topology Machine Organization Screen (FCX287) .....	792
182. Layout of Multiprocessor User Activity Log Screen (FCX288).....	794
183. Layout of User Paging Menu Screen (FCX289).....	797
184. Layout of User Page Activity Screen (FCX290).....	798
185. Layout of User Page Activity Log Screen (FCX291).....	800
186. Layout of User Page Utilization Data Screen (FCX292).....	801
187. Layout of User Page Utilization Data Log Screen (FCX293).....	803
188. Layout of Available List Data Below 2G Screen (FCX294).....	804
189. Layout of Available List Data Above 2G Screen (FCX295).....	806
190. Layout of Steal Statistics Screen (FCX296).....	807
191. Layout of Age List Log Screen (FCX297).....	809
192. Layout of Logical Core Organization Log Screen (FCX298) .....	811
193. Processor Unit (Core and Threads) Configuration Log Screen (FCX299) .....	813
194. Dispatch Vector Configuration Change Log Screen (FCX300).....	816
195. Dispatch Vector Activity Screen (FCX301) .....	818
196. Real Core Utilization Log Screen (FCX302) .....	820
197. DSVBK Steals per Processor Log Screen (FCX303) .....	822
198. Layout of Processor Log Screen (FCX304).....	825

199. Logical Partition Activity Menu Screen (FCX305) .....	828
200. Logical Partition Share Screen (FCX306) .....	829
201. Logical Partition Logs Menu Screen (FCX307) .....	832
202. Layout of CPU Pool Configuration Screen (FCX308).....	833
203. Layout of CPU Pool Activity Screen (FCX309).....	835
204. Layout of PCI Function Menu Screen (FCX310).....	837
205. Layout of PCI Function Configuration Screen (FCX311).....	838
206. Layout of PCI Function Activity Screen, Format 0 (FCX312).....	840
207. Layout of PCI Function Activity Log Screen, Format 0 (FCX313).....	842
208. Layout of Pinned Storage Log Screen (FCX314).....	843
209. Layout of Multiprocessor User Wait States Log Screen (FCX315).....	845
210. Layout of Subpool Storage Log Screen (FCX316).....	848
211. Layout of Global Networking Object Activity Screen (FCX317).....	849
212. Layout of PCI Function Activity Screen, Format 1 (FCX318).....	851
213. Layout of PCI Function Activity Log Screen, Format 1 (FCX319).....	852
214. Layout of PCI Function Activity Screen, Format 2 (FCX320).....	853
215. Layout of PCI Function Activity Log Screen, Format 2 (FCX321).....	854
216. Layout of PCI Function Activity Screen, Format 3 (FCX322).....	855
217. Layout of PCI Function Activity Log Screen, Format 3 (FCX323).....	856
218. Layout of CPU Pool Menu Screen (FCX324).....	857
219. Layout of CPU Activity Selection Menu Screen (FCX325).....	858
220. Layout of Spin Lock Collision Activity Screen (FCX326).....	859
221. Layout of HyperPAV Alias Activity Screen (FCX327).....	861
222. Layout of HyperPAV Share Activity Screen (FCX328).....	863
223. Layout of DASD Volume Performance Screen (FCX329).....	865

224. Layout of DASD Volume Performance Log Screen (FCX330).....	868
225. Layout of SMT Metrics Per Core Type Log Screen (FCX331).....	871
226. Layout of SMT Metrics Per Core Log Screen (FCX332).....	873
227. Layout of User Processor Time Screen (FCX333).....	875
228. Layout of User Processor Time Log Screen (FCX334).....	877
229. Graphics selection menu.....	958
230. Detailed performance history display with GDDM graphics.....	962
231. Detailed performance history display using simple plot.....	963
232. Summary performance history display using simple plot.....	964
233. Layout of performance trend display (GDDM).....	965
234. Layout of daily performance profile display (GDDM).....	966
235. Layout of screen for variable correlation plots.....	967
236. Layout of GDDM history data display in cumulative form.....	970
237. Layout of FCONXREP Summary History Files Screen (FCX201).....	972
238. Layout of FCONXREP screen for detailed history files.....	973
239. Example of 'FCONXREP' output (file FCONXREP LISTING).....	976
240. Layout of print file table of contents.....	979
241. Layout of print file header.....	980
242. Trend file layout.....	984



---

# Tables

1. Available performance graphics..... 956



## About this document

---

This document, with the *z/VM: Performance Toolkit Guide*, contains all information that is required for using Performance Toolkit for z/VM. Performance Toolkit for z/VM was previously known as the VM/ESA Full Screen Operator Console and Graphical Real Time Monitor, FCON/ESA, or FCON.

New users should read the introductory information in the *z/VM: Performance Toolkit Guide* first, then return to this document for reference information about Performance Toolkit subcommands, screens, and messages.

**Note:** The topic collections in *z/VM: Performance Toolkit Reference* and *z/VM: Performance Toolkit Guide* do not contain information about z/VM Performance Data Pump. z/VM Performance Data Pump is licensed with Performance Toolkit for z/VM but does not support, depend upon, or interact with Performance Toolkit for z/VM in any way.

z/VM Performance Data Pump (Data Pump) converts machine-readable z/VM monitor and SFS data into a generic text-based data stream. Modern tools can use the data stream to display real-time performance dashboards, aggregate real-time data for long-term usage analysis, or integrate with existing enterprise observability solutions.

For more information, see *z/VM Performance Data Pump* in *z/VM: Performance* and *Data Pump Messages* in *z/VM: Other Components Messages and Codes*.

## Intended Audience

---

This document is intended for use by system programmers and operators for performance analysis and general system operation. People using the document are expected to have a good general knowledge of z/VM.

## Where to Find More Information

---

For more information about z/VM functions, see the documents listed in the “Bibliography” on page 1091.

Other manuals of the z/VM system reference library may be needed, especially for performance monitoring. Please refer to the bibliography of your operating system for locating the appropriate documentation.

## Links to Other Documents and Websites

The PDF version of this document contains links to other documents and websites. A link from this document to another document works only when both documents are in the same directory or database, and a link to a website works only if you have access to the Internet. A document link is to a specific edition. If a new edition of a linked document has been published since the publication of this document, the linked document might not be the latest edition.



## How to provide feedback to IBM

---

We welcome any feedback that you have, including comments on the clarity, accuracy, or completeness of the information. See [How to send feedback to IBM](#) for additional information.



# Summary of Changes for z/VM: Performance Toolkit Reference

---

This information includes terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations for the current edition are indicated by a vertical line (|) to the left of the change.

## SC24-6303-74, z/VM 7.4 (June 2025)

---

This edition includes changes to support product changes that are provided or announced after the general availability of z/VM 7.3 and 7.4.

### [7.4 VM66826, 7.3 VM66825] z/VM Performance Data Pump enhancements

With the PTFs for APARs VM66826 (7.4) and VM66825 (7.3), z/VM 7.4 and 7.3 provide support for a new z/VM Performance Data Pump dashboard. The dashboard provides a graphical representation of IBM z17 family power consumption metrics from the z/VM monitor stream. The display includes information for the following components:

- CPC
- LPAR
- CPU
- I/O
- Memory

The information can be used to calculate or approximate guest level apportionment.

Information about setting up Data Pump is available. See [z/VM Performance Data Pump](#) in *z/VM: Performance*.

Information is available about other services that process the Data Pump output stream, including links to Grafana sample dashboards. See [z/VM Performance Data Pump](#) (<https://www.vm.ibm.com/related/perfkit/datapump/>).

## SC24-6303-74, z/VM 7.4 (September 2024)

---

This edition supports the general availability of z/VM 7.4. Note that the publication number suffix (-74) indicates the z/VM release to which this edition applies.

### Miscellaneous updates for z/VM 7.4

The following topics are updated:

- [“FCX195, Extended History Data Log Screen – REDHIST”](#) on page 609
- [“FCX300, Dispatch Vector Configuration Change Log Screen – DSVCLLOG”](#) on page 816
- [“FCX303, DSVBK Steals per Processor Log Screen – DSVSLOG”](#) on page 822
- [“FCX304, Processor Log Screen – PRCLOG”](#) on page 825
- [“Performance Variables of Trend Records X'FC01'”](#) on page 989
- [“Performance Variables of Trend Records X'FC02'”](#) on page 1014

## SC24-6303-73, z/VM 7.3 (September 2023)

---

This edition includes changes to support product changes that are provided or announced after the general availability of z/VM 7.3.

### [VM66678, VM66709] Warning Track Interruption Facility

With the PTFs for APARs VM66678 (CP) and VM66709 (Performance Toolkit), z/VM 7.3 exploits a feature of Processor Resource/Systems Manager (PR/SM) called the *warning-track-interruption facility*. z/VM's exploitation of this facility helps improve guest response time and overall performance of workloads that are run on vertical-low or vertical-medium logical processors.

## SC24-6303-73, z/VM 7.3 (August 2023)

---

This edition includes changes to support product changes that are provided or announced after the general availability of z/VM 7.3.

### [VM66687] z/VM Performance Data Pump

With the PTF for APAR VM66687, z/VM 7.3 supports z/VM Performance Data Pump. z/VM Performance Data Pump (Data Pump) converts machine-readable z/VM monitor and SFS data into a generic text-based data stream. Modern tools can use the data stream to display real-time performance dashboards, aggregate real-time data for long-term usage analysis, or integrate with existing enterprise observability solutions.

z/VM Performance Data Pump is licensed with Performance Toolkit for z/VM but does not support, depend upon, or interact with Performance Toolkit for z/VM in any way.

For more information, see [z/VM Performance Data Pump](#) in *z/VM: Performance* and [Data Pump Messages](#) in *z/VM: Other Components Messages and Codes*.

### Miscellaneous updates for August 2023

The following section is updated:

- [“FCX109, CP-Owned Device Screen – DEVICE CPOWNED” on page 416](#)

## SC24-6303-73, z/VM 7.3 (September 2022)

---

This edition supports the general availability of z/VM 7.3. Note that the publication number suffix (-73) indicates the z/VM release to which this edition applies.

### Miscellaneous updates for z/VM 7.3

The following sections are updated:

- [“FCX236, Processor Load and Configuration Logs Menu – PROCMENU” on page 701](#)
- [“FCX249, SCSI Device – SCSI” on page 721](#)



---

# Chapter 1. Performance Toolkit Subcommands

There are subcommands which can be entered and executed in basic control mode, re-display mode and performance monitor mode. Some subcommands are meaningful in multiple modes and some are only meaningful in one specific mode.

The subcommands are entered on the command line of Performance Toolkit like any CP or CMS command, but they are not re-displayed in the output area, nor are they included in the log file of Performance Toolkit.

Be aware that subcommands executed over a remote connection to the PERFSVM machine will be executed on the PERFSVM machine itself and *not* on the remote user's machine.

---

## Entering Performance Toolkit Subcommands from EXECs or CMS Programs

Most Performance Toolkit subcommands can be executed from within CMS programs or EXEC procedures. The interface built into Performance Toolkit for executing commands will trap the CMS command 'FCONCMD' and pass its arguments on for execution by the Performance Toolkit command handling subroutines.

Two methods are supported:

- *Delayed execution* of the Performance Toolkit subcommand: Execution of the command string FCONCMD FCSTACK subcmd arg1 arg2 arg3 ...

will cause the subcommand and its arguments to be stacked in the command queue, and to be executed sequentially (first-in / first-out) like any other command entered from the command line. Execution of stacked commands will start only when the current CMS command or EXEC has ended. Most commands can be submitted and correctly executed in this way.

Exception: commands which generate GDDM graphics may execute, but the program will require manual intervention (QUIT-function) to end the graphics program.

- *Immediate execution* of the Performance Toolkit subcommand: Execution of the command FCONCMD subcmd arg1 arg2 arg3 ...

will cause immediate execution of the subcommand specified. In many cases this is what is desired (e.g. changing the screen update mode from 'UPDTCMS IMMED' to 'UPDTCMS DELAYED' in a procedure which will use full screen displays). In other cases it will not produce the desired results because:

1. The corresponding subcommand set a flag which would cause execution of the command (e.g. selection of a specific performance display) when the program is looking for more work, or
2. Subsequent subcommands entered with the FCONCMD interface may reset such a flag before the previous command has actually been executed.

Use of the 'FCONCMD FCSTACK' format for delayed execution is the preferred method.

## Return Codes

The return code from the FCONCMD module does **not** indicate problems with the actual Performance Toolkit subcommand execution. The possible return codes and their meaning are:

**-3**

Can indicate either of the following:

- FCONCMD module not found, or
- FCONCMD module called natively, i.e. not running under control of Performance Toolkit, or a mismatch of the FCONCMD and PERFKIT module levels.

Checking the return code for FCONCMD is a simple way of verifying whether an EXEC procedure does, or does not, run under Performance Toolkit.

**0**

This is the usual return code you will see when the FCONCMD module is executed under control of Performance Toolkit, regardless of the success of the subcommand execution.

**>0**

Any positive return code is the number of the PROCESS entry just created by means of an 'FC PROCESS ..' command, i.e. such a return code should never be the result of executing any other subcommand.

You may need the PROCESS entry number when coding automated procedures that have to dynamically define and delete PROCESS entries. This number is required for deleting a specific entry again by means of the 'DELETE' argument of the 'FC PROCESS ..' subcommand.

## General Subcommands

---

### Asterisk '\*'

#### Purpose

The asterisk indicates that the subcommand entered is to be treated as a comment.

#### Format

➤ \* — *data* ➤

#### Usage

The 'command' will be passed to CP and included in the basic mode display when entered in either basic mode or redisplay mode; it will simply be ignored when entered in any other mode.

## Equal sign '='

### Purpose

The 'equal' sign indicates that the last subcommand entered should be re-executed.

### Format

» = «

### Parameters

No parameters are allowed or tested for.

## Question mark '?'

### Purpose

The question mark retrieves the most recently entered subcommands from an internal command buffer where all subcommands entered (except '?' and '=') are stacked. Entering '?' for the first time retrieves the last subcommand entered from the command buffer.

### Format

» ? «

### Parameters

No parameters are allowed or tested for.

# BASIC

## Purpose

The BASIC subcommand forces an immediate return to Performance Toolkit 'basic' mode, regardless of the current operating mode.

## Format

➤ BASIC ➤

## Parameters

No parameters are allowed or tested for.

# CANCEL

## Purpose

The CANCEL subcommand forces an immediate exit from Performance Toolkit, regardless of the current operating mode.

## Format

➤ CANCEL ➤

## Parameters

No parameters are allowed or tested for.


# HELP

## Purpose

The HELP subcommand activates CMS HELP for the Performance Toolkit. Specific HELP information can be selected by additional arguments.

For calling the general CMS HELP facility, you will have to enter the subcommand 'CMS HELP .... ..'.

## Format



## Parameters

### PERFKIT

Selects an initial HELP menu for Performance Toolkit. The same information is shown if the command is entered without any additional arguments, or if HELP mode is entered by pressing PF1.

### ssssss

Is any Performance Toolkit subcommand for which HELP information is displayed. In the case of the subcommand 'FCCONTROL' you can enter an additional argument to select the corresponding HELP text.

**Example:** To display HELP information for the 'FCCONTROL PROCESS' subcommand, enter the subcommand:

```
HELP FC PROCESS
```

When different subcommand descriptions are available for different modes, HELP text corresponding to the current Performance Toolkit mode will be shown if available, otherwise the search order

1. Basic mode HELP text
2. Redisplay mode HELP text
3. Performance monitor mode HELP text

will be used for locating the appropriate HELP information.

### Fields

Specifies that HELP information for performance monitor fields is to be shown. The information shown depends on the current status of the screen:

- If a performance monitor screen is being displayed for which HELP information is available, the field descriptions for the selected screen will be shown.
- A HELP menu with a selection of all available field descriptions will be shown in all other cases (screen not in performance monitor mode, or no specific HELP information available for the selected performance display).

### Msg {FCXxxxxnnx}

Specifies that HELP information for a message is to be shown, where *FCXxxxxnnx* is the message prefix. If the message prefix *FCXxxxxnnx* is specified, the corresponding HELP text will be shown. When 'Msg' is selected without a message prefix, the data shown depends on the current status of the screen:



- If a Performance Toolkit message is shown in the message area while the subcommand 'HELP MSG' is entered, HELP information for that message will be selected.
- If the message area is empty, a general instruction will be shown which explains how to select HELP information for a specific message.

## **FCONTROL**

### **Purpose**

The FCONTROL subcommand allows tailoring of many Performance Toolkit functions for specific uses. All remaining general subcommands are different arguments of this command.

## Format

» FControl	ACTmsg	»
	AUTORefr	
	BENCHmrk	
	COlo(u)r	
	DEFLog	
	DEFScr	
	DEVGROU	
	DUMPid	
	EMergenc	
	FCXLevel	
	FORCeusr	
	GDDMspe	
	HMA	
	LImit	
	MAINTid	
	MAxrec	
	MINPaths	
	MOdel	
	MONCAche	
	MONcoll	
	MSGClear	
	MSGWait	
	NUmbers	
	PFkey	
	PRocess	
	RELoad	
	RMTlines	
	SCroll	
	SEarch	
	SECuser	
	SETEvent	
	SETtings	
	TIMStamp	
	UClass	
	UPdtcms	
	UPDTScr	
	USerbotl	
	USERHdr	
	USERVar	
	USRlimit	

## Usage


1. Because most of these FCONTROL subcommands have many options, the full command format will be shown only in the detailed description of the subcommand.
2. Usually you will not see a reply to acknowledge receipt of FCONTROL subcommands (except in the case of QUERY options). The absence of an error message indicates that the command was correct and that it has been executed.

# FCONTROL ACTMSG

## Purpose

The FCONTROL ACTMSG subcommand specifies the number of action messages (line type CPAMSG) that are to be left pending at the top of the screen while the virtual machine running Performance Toolkit is logged on to a display terminal. Action messages will be scrolled only while their number exceeds the value defined with the ACTMSG subcommand. ACTMSG is initialized to a value of ten.

## Format



## Parameters

**n1**  
Specifies the number of action messages that are not to be scrolled. Minimum is 0 and the maximum is 99 (will keep *all* action messages on any screen).

**Query**  
Inserts and displays the current value of ACTMSG on the command line. QUERY is the default.

## Usage

1. The ACTMSG setting will be ignored while the virtual machine running Performance Toolkit is in disconnected state *if the logical display would otherwise become locked*.  
This change in the scrolling logic is implemented in order to prevent a lock-up of the (disconnected) logical display which could lead to an abend due to storage constraints when none of the new output lines can be unstacked and inserted into the basic mode display.

# FCONTROL AUTOREFR

## Purpose

The FCONTROL AUTOREFR subcommand specifies whether the display should be updated automatically in an APPC/VM remote performance retrieval session. The initial default operation mode is to *not* refresh the screen automatically.

## Format



## Parameters

### ON

Indicates that autorefresh mode is to be activated for this client machine. The refresh interval will automatically be set to the server's data collect interval, i.e. it may vary, depending on the server you connect to.

### OFF

Indicates that autorefresh mode is to be deactivated.

### Query

Inserts and displays the current AUTOREFR setting on the command line. QUERY is the default.

## Usage


1. The command is to be entered in the *client* virtual machine that wants to activate auto refresh mode.
2. The string (A) will be appended to the system identifier in the top right hand corner of remote performance data displays for APPC/VM sessions while autorefresh mode is active.

# FCONTROL BENCHMRK

## Purpose

The FCONTROL BENCHMRK subcommand defines any I/O devices, users, ISFC end points, ISFC logical links, or PCI functions that are to be "benchmarked", that is, for which individual "per sample" logs are to be created. Creating such logs will help in the detailed analysis of the behavior of a specific object.

## Format



## Parameters

### Query

Displays the list of currently-benchmarked objects and their properties. This is the default if no arguments are entered with the subcommand.

### DEvice *nnnn*

Defines the device number (*nnnn*) of an I/O device that is to be benchmarked.

### FILE

Lets the benchmarking data be written to volatile in-storage logs and to disk files during the specified period.

### *hh:mm*

Are the start and end times respectively of the period during which disk logs are to be created. The time must be entered in exactly this format, that is, two digits each for hours and minutes, separated by a colon.

### USer *userid*

Defines the user identification (*userid*) of a virtual machine that is to be benchmarked.

### EPoint *epoint*

Defines the ISFC end point number that is to be benchmarked, in hexadecimal. Leading zeroes can be omitted.

### LLink *partner*

Defines the name of a ISFC logical link that is to be benchmarked.

### PCifunc *pcifid*

Defines the PCI function ID that is to be benchmarked, in hexadecimal. Leading zeroes can be omitted.

### OFF

Disables benchmarking for the specified object, or for all objects already benchmarked if FC BENCHMRK OFF is used.

## Usage

1. The following files can be created:

I/O device files:

***devno* CACHDLG**

For the CACHDLOG

***devno* CACHELG**

For the CACHELOG

***devno* CPOWNLG**

For the CPOWNLOG

***devno* DEVLOG**

For the DEVLOG

***devno* HPFLOG**

For the HPFLOG

***devno* QDIOLOG**

For the QDIOLOG

***devno* QEBSMLG**

For the QEBSMLOG

***devno* SCSILOG**

For the SCSILOG

***devno* SEEKDLG**

For the SEEKDLOG

***devno* SEEKLLG**

For the SEEKLLOG

***devno* VOLLOG**

For the VOLLOG

User files:

***userid* LXCPULG**

For the LXCPULOG

***userid* LXMEMLG**

For the LXMEMLOG

***userid* LXNETLG**

For the LXNETLOG

***userid* MTUSRLG**

For the MTUSRLOG

***userid* UCOMMLG**

For the UCOMMLOG

***userid* UPACTLG**

For the UPGACTLG

***userid* UPAGELG**

For the UPAGELOG

***userid* UPUTLLG**

For the UPGUTLLG

***userid* UQDIOLG**

For the UQDIOLOG

***userid* URESPLG**

For the URESPLOG

***userid* USERLOG**

For the USERLOG



***userid* USRMPLG**

For the USRMPLG

***userid* USRTMLG**

For the USRTMLOG

***userid* USTATLG**

For the USTATLOG

***userid* USTMPLG**

For the USTMPLOG

***userid* UTRANLG**

For the UTRANLOG

ISFC end point files:

***epoint* ISFELOG**

For the ISFELOG

ISFC logical link files:

***partner* ISFLLOG**

For the ISFLLOG

PCI function files:

***pcifid* PCIFLOG**

For the PCILOG (Format 0)

***pcifid* PCIF1LG**

For the PCILOG (Format 1)

***pcifid* PCIF2LG**

For the PCILOG (Format 2)

***pcifid* PCIF3LG**

For the PCILOG (Format 3)

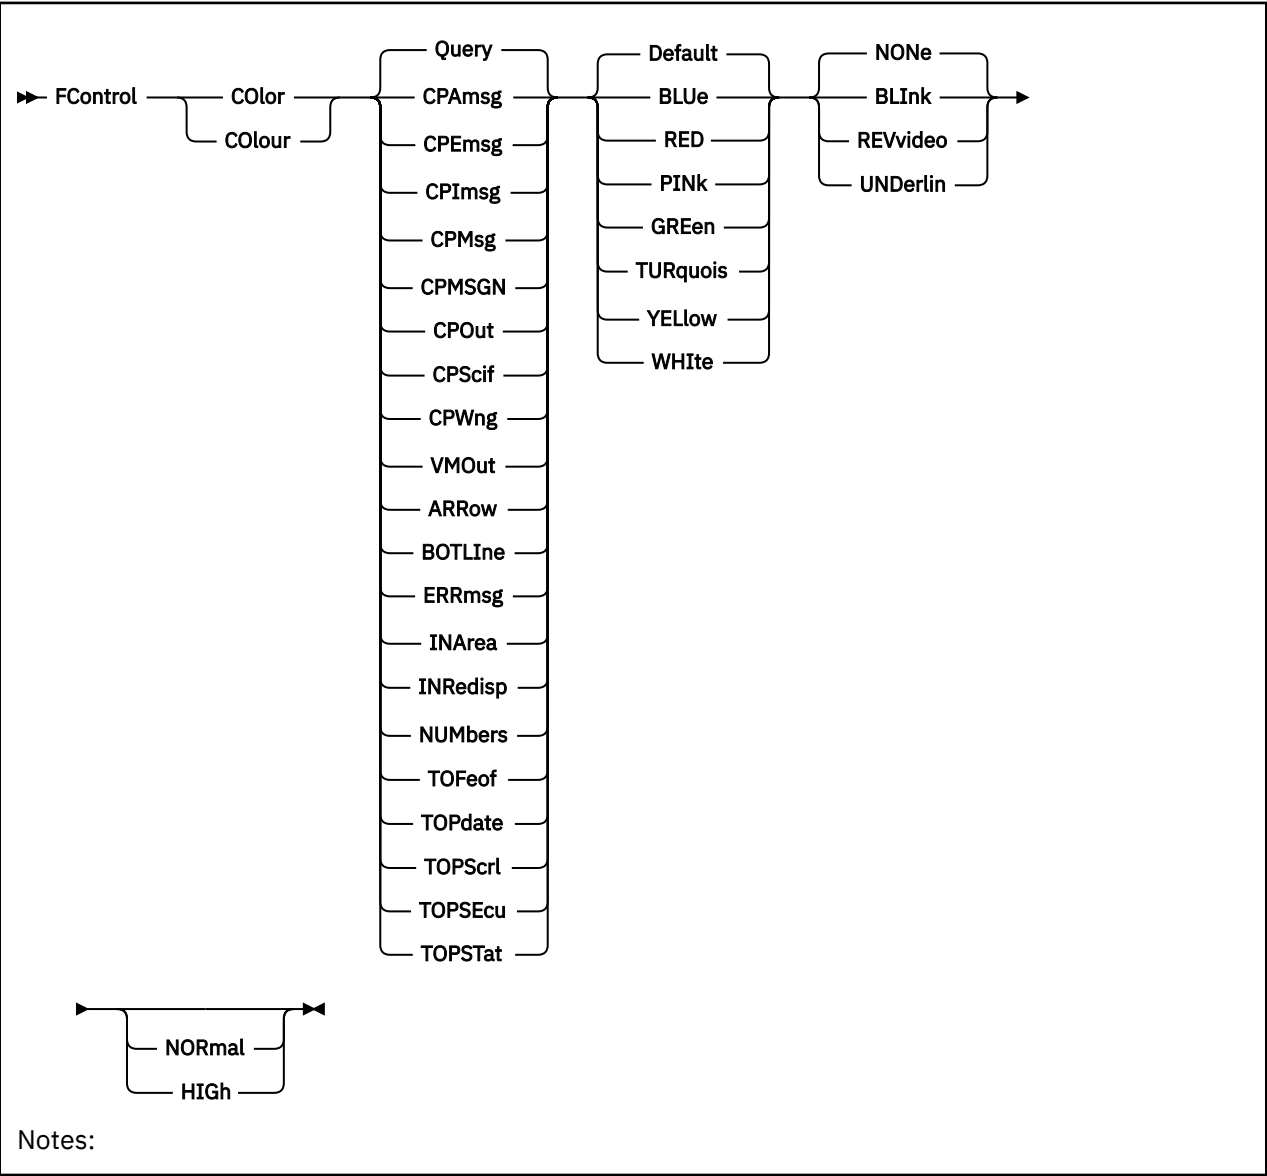
2. Up to three log generations will automatically be kept, with a "1" or "2" suffix for previous days' logs as for the console and the detailed performance history log.

# FCONTROL COLOR

## Purpose

This subcommand allows specification of extended color and highlighting to be used for all the different fields displayed by Performance Toolkit. The subcommand can be entered on any screen, but the extended color and highlighting will be used only on terminals with the appropriate support (for example: 3279 for color and highlighting, 3290 for extended highlighting only).

## Format



## Parameters

### Query

Displays the current color settings for all fields. This is also the default if no field name is entered.

### fieldname

If not 'Query'

CP output line definitions:

**CPAmsg**

CP message line that is treated as a Performance Toolkit action message (i.e. contains preceding carriage return characters)

See the "Action Messages" section in the "General System Operation" chapter of the [z/VM: Performance Toolkit Guide](#) for more information.

**CPEmsg**

CP error message (EMSG)

**CPImsg**

CP information message (IMSG)

**CPMsg**

CP message (MSG)

**CPMSGN**

CP noheader message

**CPOut**

CP command output line

**CPScif**

Output line from another machine which we received via the Single Console Image Facility (SCIF)

**CPWng**

CP warning

**VMOut**

Output line generated by the virtual machine (e.g. CMS)

Performance Toolkit field definitions:

**ARRow**

The command pointer

**BOTLine**

Bottom line of the display

**ERRmsg**

Performance Toolkit error messages

**INArea**

Command input area

**INRedisp**

CP and CMS commands redisplayed in the general display area

**NUMbers**

The line numbers of Performance Toolkit on the left side of the display area (inserted only if numbers have been set on, see ["FCONTROL NUMBERS"](#) on page 65).

**TOFeof**

Top-of-file and end-of-file indicators shown at the beginning and end of the console log of Performance Toolkit in re-display mode

**TOPdate**

The date field on the left of the top line

**TOPScrl**

Active scrolling information displayed in the center part of the top line

**TOPSEcu**

Secondary user information (or Performance Toolkit header) displayed in the center part of the top line

**TOPStat**

Screen mode field ('RE-DISPLAY' or 'PERF. MONITOR') displayed on the right of the top line when not in basic mode

The next two parameters are the extended highlighting you want to select and the color. It does not matter which is entered first.

Color Specification:

**DEFAULT**

The default colors of Performance Toolkit are to be used. Color will be blue for the following fields: CPOUT, VMOUT, CPMSG, CPMSG, CPSCIF, TOPDATE, and NUMBERS. All other fields will be displayed in white. On 3270 and 3278 type displays the 'white' fields will be highlighted while the 'blue' fields will be displayed with normal intensity.

**BLUe, RED, PINK, GREen, TURquois, YELlow and WHItE**

(When specified) are the colors that should be used instead of the default ones.

Highlighting:

**NONE**

No extended highlighting (the default)

**BLInk**

Line should be blinking

**REVvideo**

Reverse video

**UNDERlin**

Complete line is to be underlined

The following two arguments apply only to screens without extended highlighting support. They can also be entered from screens with extended highlighting support, but the changed values will not become active.

**NORMAL**

Normal intensity

**HIGH**

High intensity

**Usage**

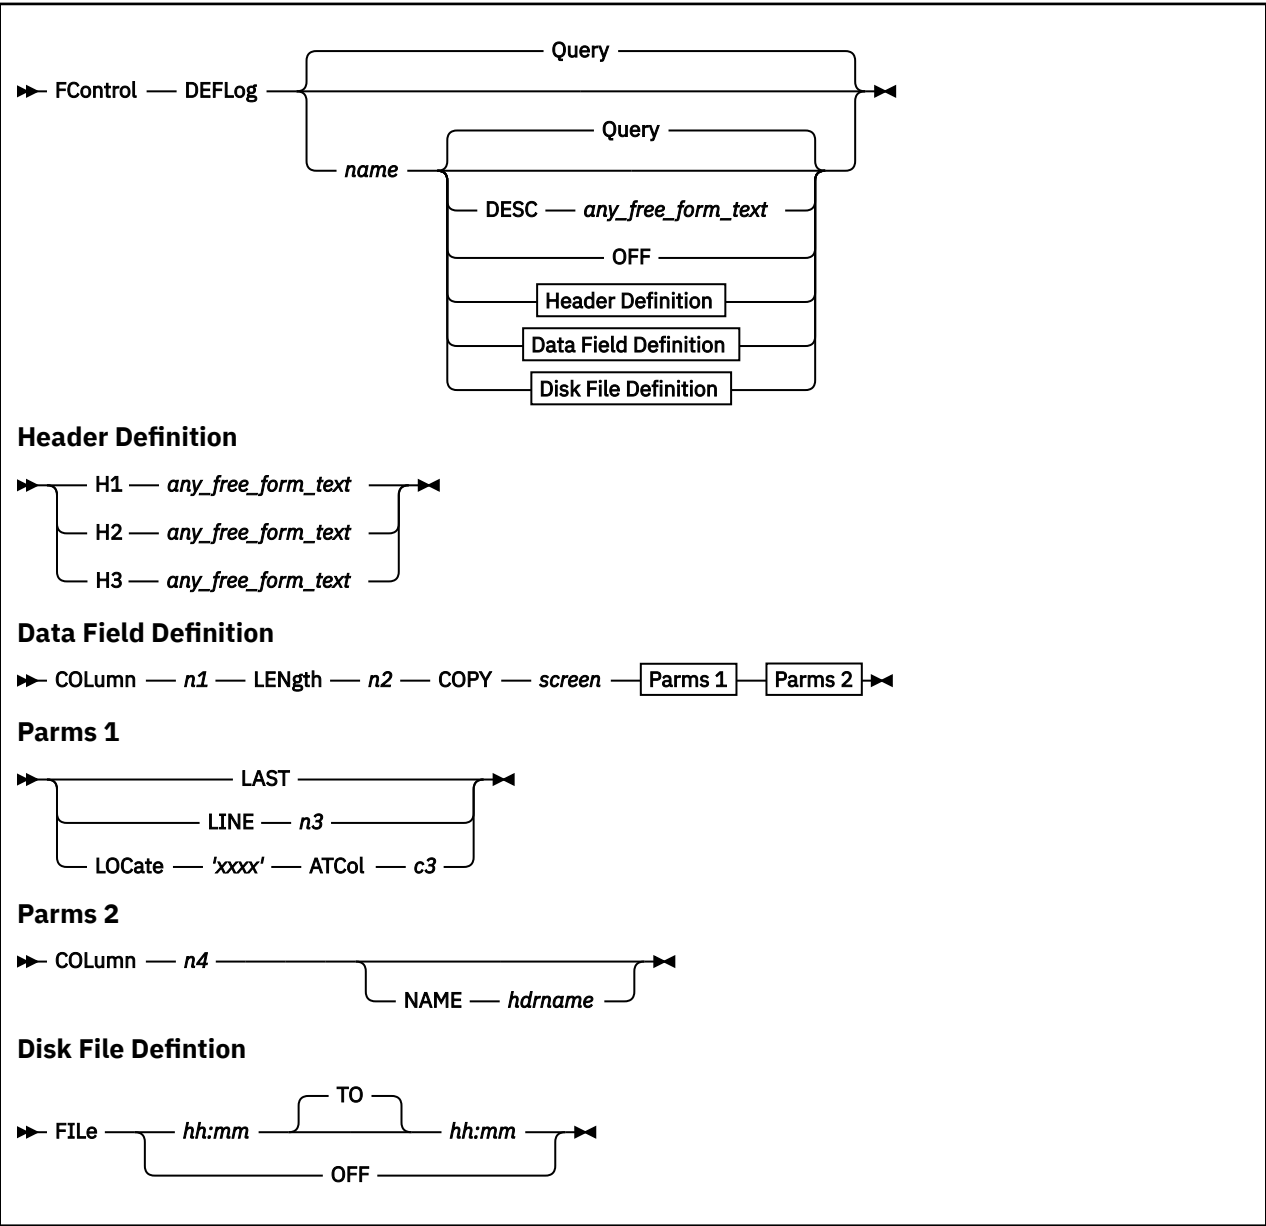
1. The first parameter is the name of the field whose attributes you want to specify. Please note that default values for both highlighting and color will be set once a valid field name has been entered. These default values will become active unless they are replaced by valid color and/or highlighting definitions in the same command.

# FCONTROL DEFLOG

## Purpose

The FCONTROL DEFLOG subcommand allows defining your own 'by time' performance data logs, to be built from parts of standard performance displays and user-defined text. It causes a User Defined Performance Data Logs Screen (FCX187) to be displayed.

## Format



## Parameters

### Query

Displays the names of all user defined logs. This is also the default action if no screen name is entered.

### name

Defines the name of the new display which is also the command you will have to enter for selecting it

**Query**

Lets the definitions of the selected user defined log be displayed. This is also the default action if no further arguments are entered.

**DESC**

The log description, where *any\_freeform\_text* will be used to describe the log in the UDEFMENU.

**OFF**

Deletes the definitions for the selected user defined log.

**H1/H2/H3**

Indicate that header lines number 1, 2 or 3 respectively are to be defined, where *any\_freeform\_text* stands for the text that is to be inserted.

**COLUMN n1**

Specifies the leftmost column of the field in the new performance log *name* where the information is to be placed which is defined by the remaining arguments

**LENGTH n2**

Defines the length of the field in the new performance log

**COPY screen**

Defines the name of the standard performance display from which lines are to be copied

'Standard' refers here to one of the general displays which can be selected by entering a single command, without additional arguments, i.e. you cannot specify any of the 'detailed' displays with data for specific users or I/O devices.

*Special case:* Use the argument COWNED to specify the CP owned device display as input.

Source Line Specification:

**LAST**

Indicates that the last line of the source display is to be copied from. Copying from the last line is usually a requirement when using any of the 'by time' log displays as input, where the most recent data fields are always found in the last line.

**LINE n3**

Defines a fixed line in the 'source' display where the source field is located

**LOCate 'xxxx' ATCol n3**

Can be used to locate the source line on displays with variable output format.

**'xxxx'**

Is a string which is to be searched for. It must be enclosed in single quotes and can have a maximum length of eight characters.

**ATCol c3**

Defines the column *c3* where the string 'xxxx' is to be looked for. These arguments are mandatory in order to reduce the overhead required for locating a line.

*Note that the column number c3 refers to the string's position in the internal output buffer. The column number seen on the display must be incremented for each display attribute found to the left of the string: Attributes take up two bytes in the buffer, but only one position on the actual display.*

Source Column and Optional Header Information:

**COLUMN n4**

Defines the leftmost column on the selected line in the 'source' display from where on data are to be copied, for length *n2*

**NAME hdrname**

Defines an optional character string which, when defined, will be inserted in header line 3 above the columns containing the data fields defined by the command.

The string:

- Must not contain blanks (a blank will be assumed to indicate the end of the string)
- Will be right-adjusted if it is shorter than the length *n2* defined for the field

- Will be truncated to the field length *n2* if it is longer.

Output File on Disk:

### **FILE**

Lets the benchmarking data be written not only to volatile in-storage logs but also to disk files during the specified period.

#### **hh:mm**

Are the start and end times respectively of the period during which disk logs are to be created. The time must be entered in exactly this format, i.e. two digits each for hours and minutes, separated by a colon.

### **OFF**

Deactivates logging to disk file.

The file name of the disk file will be the name of the user defined log, and the file type is always UDEFLOG

Up to three log generations will automatically be kept, with a '1' or '2' suffix for previous days' logs as for the console and the detailed performance history log.

## **Usage**

1. The first ten columns are reserved for the time stamp with the sample interval end time; they cannot be overwritten by performance data fields.
2. The number of field definitions for any single user-defined performance log is limited to 20.
3. The maximum width of the new log is 132 characters. Field definitions extending beyond this width cannot be accepted.
4. The source data may contain display attributes for highlighting specific fields. These are 2-byte strings which end up taking only one position on the output screen.

While the DEFLOG logic should compensate for attribute strings on columns preceding the string you want to copy, you may copy:

- Only the second part of such a string if the 'source' display column you specified is just one position too low (this second part may appear as one of the characters '-' or 'Y', or as another strange character)
  - Part of a trailing attribute string if the field length you specified was a bit too long
  - Less of the actual data than expected if the field length specified spans two or more fields, with intervening attribute strings (the two bytes of each attribute must be accounted for in your 'length' definition even though each attribute will take up only the space of one character in the final display).
5. Set your field definitions so that only the actual data fields are copied, or, if you intend to copy the attribute(s) too, remember to take into account the additional data length to be copied (one additional byte per attribute).
  6. In the resulting screen, keep in mind the following notes:
    - a. *Current/Average Data:* The fields will be copied just as they have been generated for the source screen, i.e. on performance displays where different values are displayed for the 'current' and 'average' settings, averages will be copied while the average setting is active. This is probably not what you want to see on your log, so do not set averages while collecting this kind of log data.
    - b. *Data Availability:* Valid data can, obviously, be copied from the source screen only when the prerequisites for generating it are met.

## **Results**

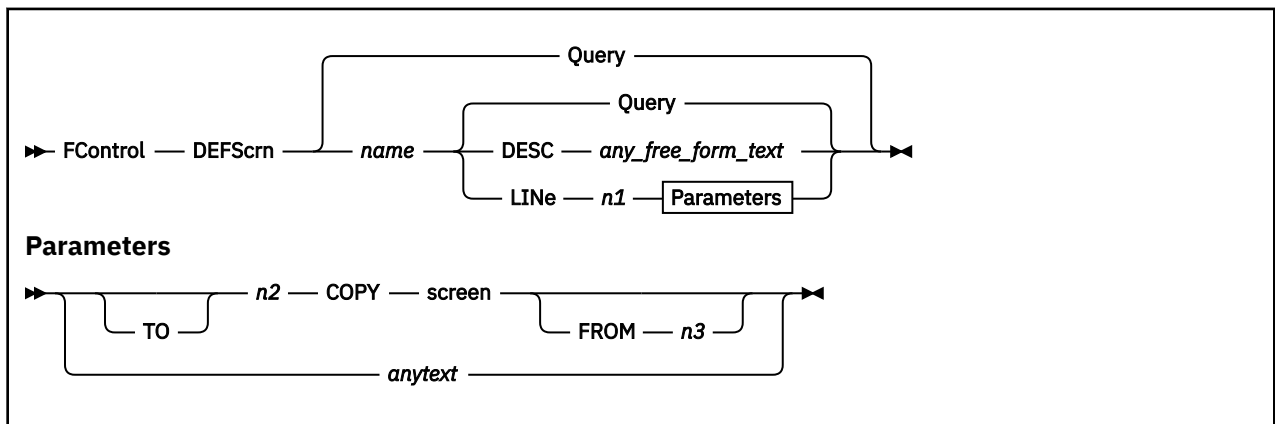
See [“FCX187, User-Defined Performance Data Logs Screen – FC DEFLOG” on page 594.](#)

# FCONTROL DEFSCRN

## Purpose

The FCONTROL DEFSCRN subcommand allows you to define your own performance displays, to be built from parts of the standard performance displays and user-defined text. It causes a User Defined Performance Screen (FCX148) to be displayed.

## Format



## Parameters

### Query

Displays the names of all user defined screen. This is the default action if no screen name is entered.

### name

Defines the name of the new display which is also the command you will have to enter for selecting it

### Query

Lets the definitions of the selected user defined screen be displayed. This is the default action if no further arguments are entered.

### DESC

The log description, where *any\_freeform\_text* will be used to describe the log in the UDEFMENU.

### LINE n1

Specifies the line number in the new display *name* where the information is to be placed which is defined by the remaining arguments

### {TO} n2

Specifies the last line number of a range of lines in the new display which are to be filled with data from the same standard performance display.

When not specified, only one line is retrieved from the standard display and copied to line number *n1*.

### COPY screen

Defines the name of the standard performance display from which lines are to be copied.

'Standard' refers here to one of the general displays which can be selected by entering a single command, without additional arguments, i.e. you cannot specify any of the 'detailed' displays with data for specific users or I/O devices.

*Special case:* Use the argument CPOWNER to specify the CP owned device display as input.

### FROM n3

Defines the first line in the 'source' display where lines are to be copied into the new display. Copying will start at the first line of the source display if the FROM argument is omitted.



**'any text'**

Defines a string of characters, enclosed by single quotes, which is to be placed on line *n1* of the new screen.

**Usage**

1. The number of definition commands for any single user-defined display is limited to eight, and the maximum number of lines which can be defined for the new performance screen is limited to 70.
2. Valid data can be copied from the source screen only when the prerequisites for generating it are met. Make sure that you include only specifications for copying from performance screens which are likely to be available.

**Results**

See [“FCX148, User-Defined Performance Screen – FCONTROL DEFSCRN” on page 513](#).

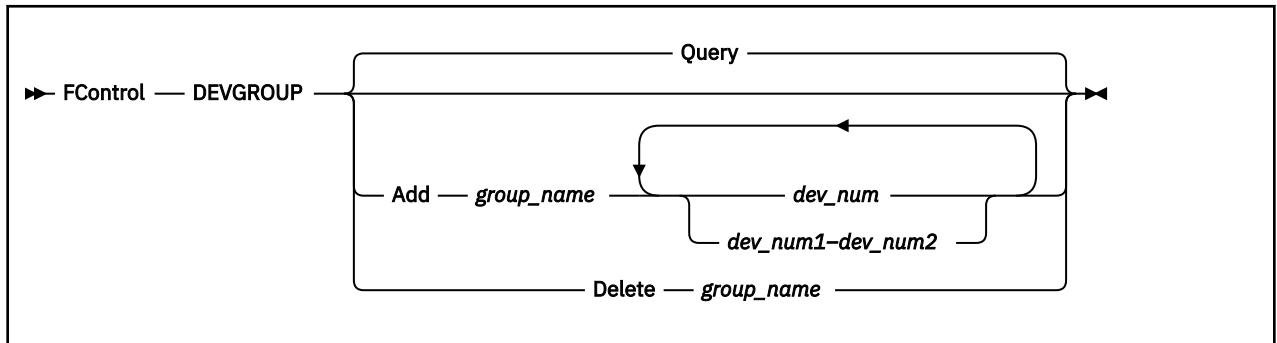
# FCONTROL DEVGROUP

## Purpose

The FCONTROL DEVGROUP subcommand allows you to define device groups or volume groups for which separate averages are to be calculated and displayed on the following screens:

- “FCX108, General I/O Device Screen – DEVICE or DEVICE DASD” on page 412
- “FCX282, HPF I/O Device Screen – DEVICE HPF” on page 778
- “FCX329, DASD Volume Performance Screen – VOLUME” on page 865.

## Format



## Parameters

### Query

Lists the device group or volume group definitions that are currently active. This is the default.

### Add

Adds a set of devices, specified by the device numbers, to the device group or volume group specified by *group\_name*.

### Delete

Deletes the device group or volume group specified by *group\_name* from the list of active device groups or volume groups and removes it from reports.

### *group\_name*

Specifies the 1- to 8-character name of the device group or volume group with which the FCONTROL DEVGROUP command operates.

### *dev\_num*

Specifies the number of the device to be added to the device group or volume group specified by *group\_name*. For a volume, this is the device number of the base RDEV.

### *dev\_num1-dev\_num2*

Indicates that you can specify a list or range of device numbers. If you specify a range of device numbers, all valid devices in the range, including the first and last devices you specify, are added to the device group or volume group specified by *group\_name*.

## Usage Notes

1. FCONTROL DEVGROUP specifications become effective immediately after they have been entered.
2. If you are defining a volume group, include only non-PAV DASD RDEVs, PAV base RDEVs, or HyperPAV base RDEVs. Any other types of RDEVs are excluded from the calculation.
3. If you issue multiple FCONTROL DEVGROUP ADD commands for the same device group or volume group name, all devices specified are added to that same group.

4. If a range that is being added to a device group or volume group intersects or overlaps with a range that was previously defined for the group, the resulting range will include all specified range values.
5. A specific device can belong to several device groups or volume groups.
6. If you request a device report for a specific device type (DEVICE DASD or DEVICE HPF, for example), you will still see device group rows for every device group you defined, whether or not those rows are meaningful for the specific report you requested.
7. If you define a device group that consists of devices of multiple types — for example, if you mix CTCs and DASDs in the same device group — the device group statistics that Performance Toolkit calculates might not be useful or meaningful.

# FCONTROL DUMPID

## Purpose

The FCONTROL DUMPID subcommand defines the destination address where dumps are to be sent in case of an abend with automatic dump. Dumps are sent to the virtual reader of the machine where Performance Toolkit has been running at the time of the abend if no other address has been specified, or it has been specified as an asterisk (\*). The CP VMDUMP command is used for creating these dumps.

## Format



## Parameters

### *userid*

Defines the user identification of the virtual machine where automatic dumps are to be sent after an abend of Performance Toolkit. An asterisk may be specified to send the dumps to the virtual reader of the machine where Performance Toolkit has been running.

### **SYSTEM**

Defines the target destination for automatic dumps as it has been specified on the SYSTEM\_USERIDS configuration file statement. (This is the user ID designated to receive any CP abend dumps.) For further details, see the description of this parameter in the help for the CP VMDUMP command.

### **Query**

Inserts and displays the current target destination for abend dumps. This is the default if no *userid* is specified.

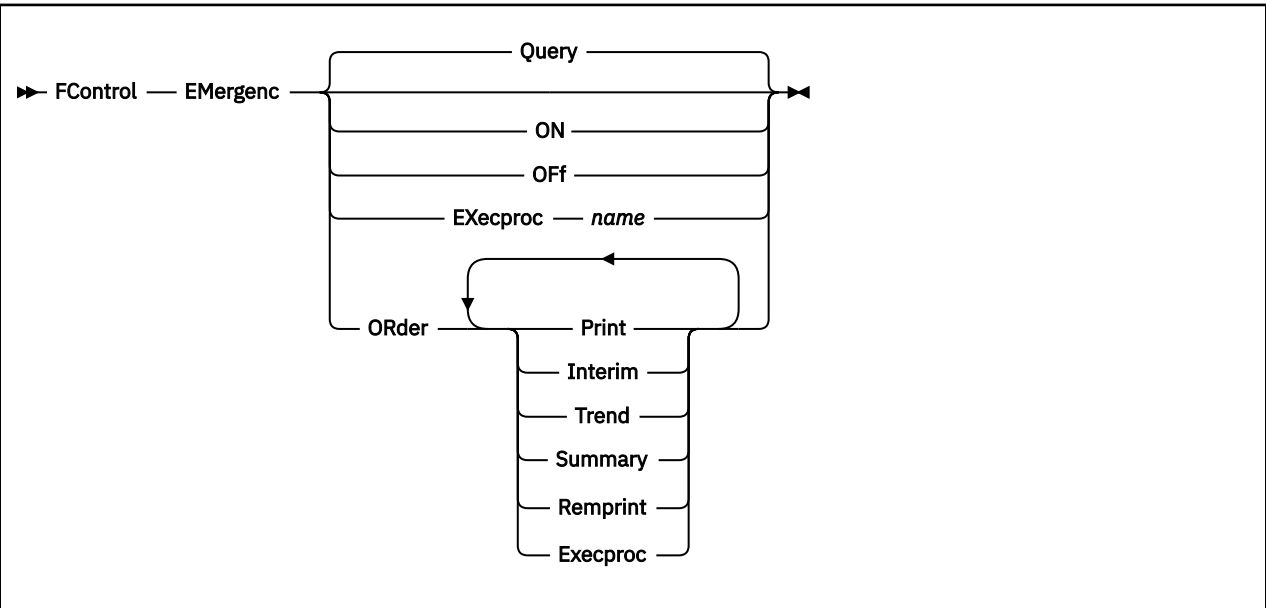
# FCONTROL EMERGENC

## Purpose

The FCONTROL EMERGENC subcommand allows you to control Performance Toolkit's emergency safeguard feature (ESF). With this feature enabled, Performance Toolkit will automatically save any unsaved data collected during normal monitoring in the event of an emergency, such as a z/VM shutdown or a sudden shortage in virtual memory. Data that can be saved includes printed reports, as well as trend and summary files.

In addition to saving data, Performance Toolkit will also call any user-defined EXEC procedure you've specified when an emergency event occurs.

## Format



## Parameters

### Query

Returns a list of all the emergency safeguard feature (ESF) specifications as they are currently set. This is the default if no further argument is entered.

ESF status can be one of the following:

#### Disabled

No special actions are taken in emergency situations. This is the default setting.

#### Enabled

Emergency events detection activated, ready to take specified actions.

#### Suspended

Emergency events detection activated, but they will be ignored because permanent data collection is not active. Use FC MONCOLL ON command to activate it.

### ON

Activates the emergency safeguard feature.

### OFF

Deactivates the emergency safeguard feature.

**Execproc name**

Defines *name* EXEC as the CMS EXEC procedure that will be called upon detection of an emergency event. Note that you must also specify the order in which all actions, including this EXEC procedure, will be specified, using an additional FC EMERGENC ORDER command.

For more information on this parameter, see Usage Note “4” on page 30.

**Order**

Allows specification of the actions that will be done, as well as their order, when an emergency event is detected. Multiple selections are possible. Note that each subsequent FC EMERGENC ORDER command invocation will clear any previously defined actions.

The valid actions are as follows:

**Print**

Indicates that the print reports defined with the FC MONCOLL RESET *hh:mm:ssP* command(s) will be created.

**Interim**

Indicates that the print reports defined with the FC MONCOLL RESET *hh:mm:ssI* command(s) will be created.

**Trend**

Indicates that the trend records defined with the FC MONCOLL RESET *hh:mm:ssT* command(s) will be processed and saved.

**Summary**

Indicates that the summary records defined with the FC MONCOLL RESET *hh:mm:ssS* command(s) will be processed and saved.

**Remprint**

Indicates that the print reports defined with the FC SETEVENT ... PRINTRMT command(s) will be created.

**Execproc**

Indicates that the CMS EXEC procedure previously defined by the FC EMERGENC EXECPROC command will be called.

**Usage**

1. If an emergency event occurs and any of the PRINT, INTERIM, TREND or SUMMARY actions have been specified with the FC EMERGENC ORDER command, Performance Toolkit will attempt to save any collected data by scanning the reset chain (as defined by the FC MONCOLL RESET command) to find any appropriate actions that have been scheduled before the next reset. If any such actions are found, Performance Toolkit will perform the corresponding report printing, trend/summary file creation, or both, just as if they had been defined using the FC MONCOLL RESET command. Performance Toolkit will avoid any unnecessary duplicate actions.
2. For the REMPRINT action, the operation is similar to that described above, except in this case Performance Toolkit will scan the event chain defined with the FC SETEVENT command, looking for only those event elements that reference the PRINTRMT command. If any scheduled PRINTRMT commands are found, Performance Toolkit will process them, avoiding any duplicate actions.
3. When performing the emergency safeguard actions, there are special timing considerations that must be kept in mind regarding TREND and SUMMARY reports. Because the overall time available for performing these safeguard actions may be limited by the system (as defined by the CP SET SHUTDOWNTIME command), care should be taken when defining the order in which these actions are taken, so that the most valuable data is saved first.

Note that Performance Toolkit keeps trend data in records with a fixed size of 4K bytes, while summary data records have variable lengths. Because of this format difference, trend files can be written to disk much faster.

4. Special notes for the EXECPROC option:

- a. Remember that you must specify a user-defined EXEC procedure in two separate commands – an FC EMERGENC EXECPROC *name* command to define the actual EXEC procedure, and an FC EMERGENC ORDER command to specify the order in which all actions, including this EXEC procedure, will be specified.
- b. A user-defined EXEC procedure will be notified of the emergency event reason by way of a special parameter that will be passed to it. The possible parameter values and their meanings are as follows:

**\*SHUTDOWN**

A signal shutdown event has been detected.

**\*NOMEM**

Performance Toolkit ran out of virtual storage.

**\*NOCONS**

A permanent Performance Toolkit console I/O error (it was detached or redefined with a different device number).

**\*UNKNOWN**

All others (this value should not occur).

A sample parameter handling code is shown below:

```
/* Sample Emergency Safeguard Feature EXEC procedure */
Arg emgparm .
Select
  When emgparm = '*SHUTDOWN' Then reason = 'SHUTDOWN'
  When emgparm = '*NOMEM'   Then reason = 'Performance Toolkit memory shortage'
  When emgparm = '*NOCONS'  Then reason = 'Permanent console I/O error'
  When emgparm = '*UNKNOWN' Then reason = 'unknown reason'
  Otherwise Nop
End
'tell OP1 Emergency exit is being taken due to' reason
```

- c. When writing a user-defined EXEC procedure, remember that certain restrictions may apply, due to the nature of the emergency of event. In the case of a system shutdown, for instance, the same timing consideration as described in Usage Note “3” on [page 30](#) will apply. If the procedure is called due to a lack of virtual storage, any further storage-intensive functions may lead to a CMS system abend. (Although note that a special reserved free storage buffer will be provided to help mitigate this possibility.)

## FCONTROL FCXLEVEL

### Purpose

The FCONTROL FCXLEVEL subcommand allows displaying the Performance Toolkit program version, release, modification level and service level.

Message FCXxxx500I will be issued with this information. In addition, the service level information for each part in the module used for this invocation of Performance Toolkit will be written to the log. See [“MODLEVEL” on page 108](#) for the details on how this information is obtained.

### Format

►► FControl — FCXLevel ◄◄

### Parameters

No parameters are allowed or tested for.

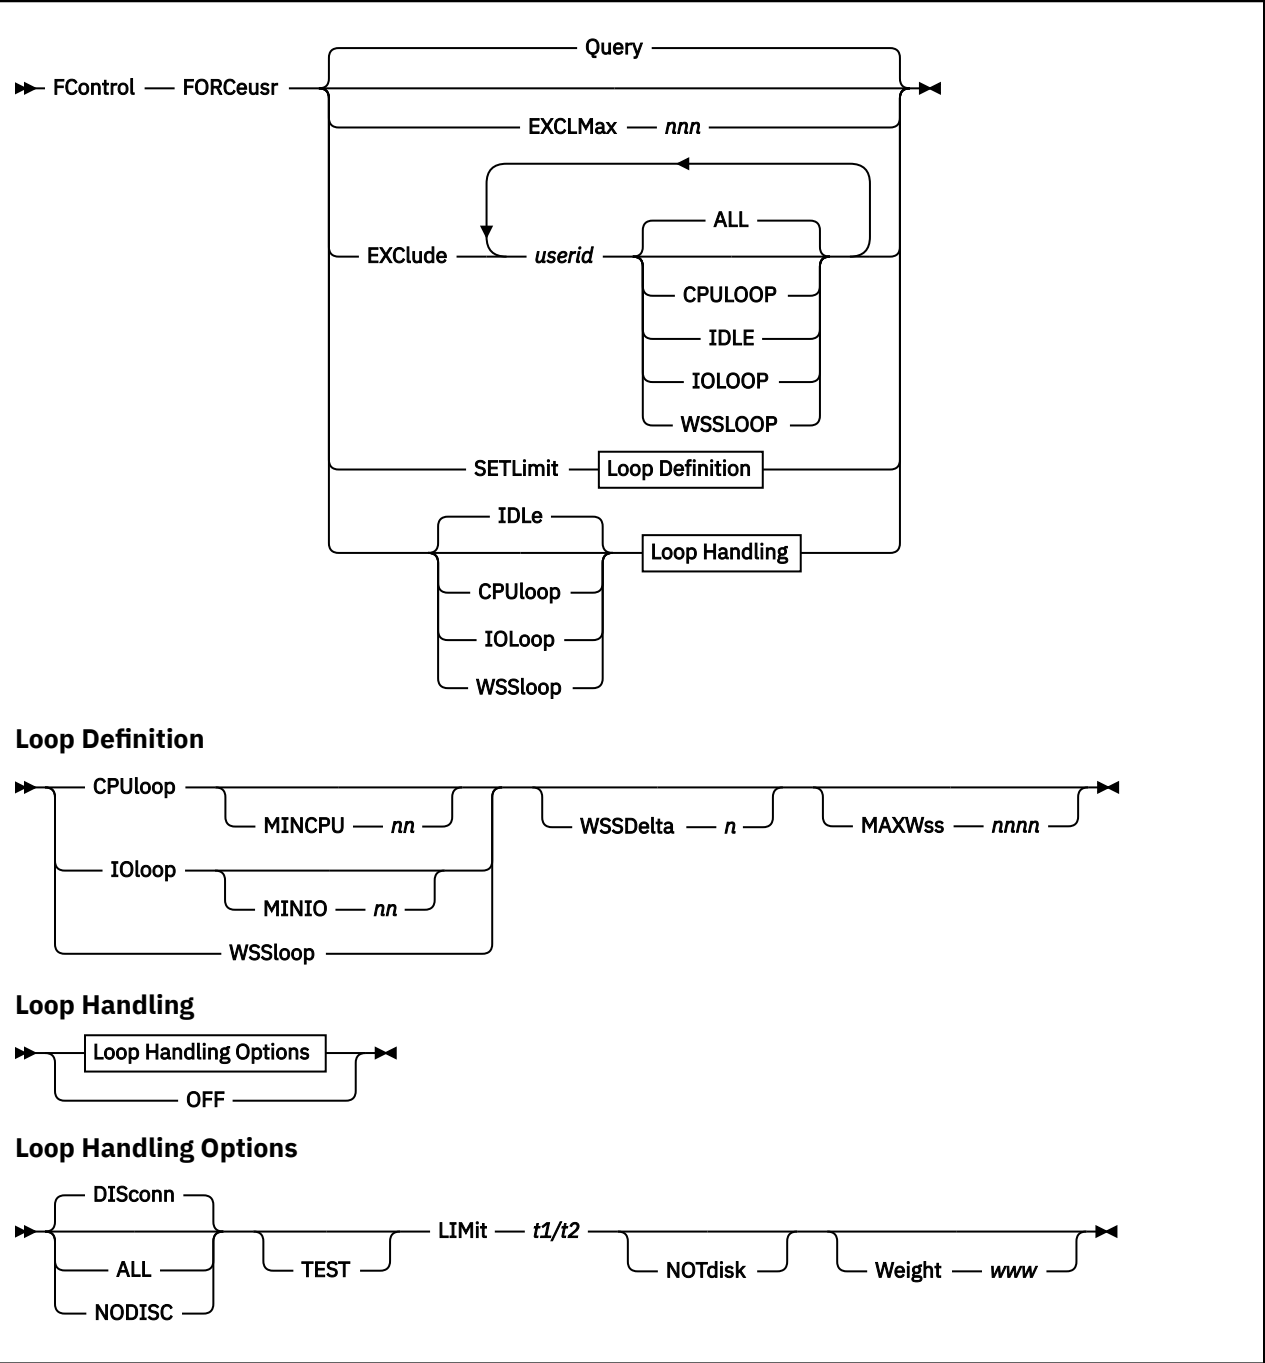


# FCONTROL FORCEUSR

## Purpose

The FCONTROL FORCEUSR subcommand is intended for controlling the user monitoring function. It allows monitoring and, optionally, forcing of virtual machines which have been looping, or have been idle for a long time. Forcing of such machines will free the resources (e.g. CPU, main storage) they occupied, and can improve the performance of the remaining users of the system. Make sure, however, that all users are informed of such a policy before activating this facility, and that they understand the implications (e.g. data stored onto a temporary minidisk will be lost when a machine is forced). The command is valid, and looping/idle user detection active, only when permanent performance data collection has been activated.

## Format



## Parameters

### **EXCLMax *nnn***

Defines the maximum number of users which must fit into the 'exclude' table. The number *nnn* can be any value between 10 and 2000. A default table size of 500 will be used if the 'EXCLUDE *userid*' argument is used without a prior explicit table size specification.

### **EXclude *userid***

Specifies which virtual machines are to be excluded from being monitored / forced, and, optionally, for which cases they are to be excluded. 'Wildcard' characters '\*' and '%' can be used for creating generic entries which could apply to several machines; they have the same meaning as for the CMS 'LISTFILE' command. Several userids can be specified with a single command, and several EXCLUDE commands can be entered. All userids from all EXCLUDE commands will be inserted into a common exclude table, up to the maximum which fits into the table (defined by EXCLMAX argument, or default of 500).

Users can only be added to the exclude list, but not removed, you will have to re-initialize PERFKIT if you want to clear the table.

### **ALL**

Indicates that the user is to be excluded from FORCEUSR processing for *all* the monitored cases (i.e. CPULoop, IDLE and IOLOOP). ALL is the default setting which will be assumed if no specific cases are specified.

### **CPULoop**

Indicates that FORCEUSR CPULoop processing is not to be applied for the user even if the virtual machine appears to be in a CPU loop

### **IDLE**

Indicates that FORCEUSR IDLE processing is not to be applied for the user

### **IOLOOP**

Indicates that FORCEUSR IOLOOP processing is not to be applied for the user

### **WSSLOOP**

Indicates that FORCEUSR WSSLOOP processing is not to be applied for the user

Multiple selections are possible, i.e. you can exclude a virtual machine from both IOLOOP and IDLE state monitoring, but leave CPU loop detection active. The above arguments must be entered in full length.

## **Query**

Returns a list of all the FORCEUSR specifications as they are currently set. This is the default if no further argument is entered.

## **SETLimit**

Allows specification of the criteria used to determine whether a virtual machine should be assume to be in a loop of some sort.

### **CPULoop**

Indicates that the criteria for detecting a CPU loop are to be set

### **MINCPU *nn***

Specifies the minimum CPU load (in percent) which must be reached before a virtual machine can be a CPULoop candidate.

The value is initially set to 10(%).

### **IOLoop**

Indicates that the criteria for detecting an I/O loop are to be set

### **MINIO *nn***

Specifies the minimum virtual I/O rate which must be reached before a virtual machine can be an IOLOOP candidate.

The value is initially set to 500 (I/Os per second).

**WSSloop**

Indicates that the criteria for detecting a WSS loop are to be set

**WSSDelta**

Specifies the maximum difference in a virtual machine's working set size which is to be tolerated for the selected kind of loop. The initial WSSDELTA values are

**1**

For the CPULoop case

**4**

For the IOLOOP case

**15**

For the WSSLOOP case

Note that all looping virtual machines should, by definition, be permanently executing the same instructions all over again, and their working sets should thus always be the same. If the loop is large enough, however, the working set size can change by several pages, and this is especially true in the case of any loop which involves I/O activity.

Machines will be assumed to be not looping if their working set size changes by more than the specified max. delta between consecutive monitor sample intervals.

**MAXWSS**

Specifies the maximum working set size allowed for a virtual machine for the selected kind of loop.

The initial MAXWSS values are set to 10000 for all three loop cases.

Although a loop could theoretically touch any number of pages, i.e. there is no guarantee at all that a machine with a large working set is not looping, experience has shown that setting an upper limit for the working set can increase the probability of catching only real loopers, and reduce the risk of treating a non-looping machine as a looper.

Machines will be assumed to be not looping if their working set size exceeds the maximum number of pages at the end of a sample interval.

**CPULoop**

Indicates that the following arguments refer to CPU loop detection. A machine is considered to be in a CPU loop when it has used at least the total CPU percentage specified with the MINCPU argument, has not had any IUCV activity nor executed a single non-spoiled I/O and not more than two I/Os to spoiled UR devices in an interval *and* when the number of pages in its working set has not changed by more than the WSSDELTA value between consecutive measurements, *and* when it has been found in either the dispatch or eligible list at the end of each interval and was never found dormant or test-idle by user state sampling.

Note that this may apply also to some CPU bound machines which are *not* looping, and be sure to put the corresponding userids on the exclude list.

**IDLe**

Indicates that the following arguments refer to idle machine monitoring. A machine is considered to be idle when it has used less than 10ms of total CPU, no emulation CPU, has not executed a single non-spoiled I/O and not more than two I/Os to spoiled UR devices in an interval.

**IOLoop**

Indicates that the following arguments refer to I/O loop detection. A machine is considered to be in an I/O loop when it has a virtual I/O rate of at least the value specified with the MINIO argument *and* when the number of pages in its working set has not changed by more than the WSSDELTA value between consecutive measurements, *and* when it has been found in either the dispatch or eligible list at the end of each interval and was never found dormant or test-idle by user state sampling.

This code should allow to detect, for instance, users which have ended up in an I/O loop to their virtual console, but the 'I/O loop' condition may also be fulfilled by large machines which are *not* looping. Be sure to put the corresponding userids on the exclude list.

**WSSloop**

Indicates that the following arguments refer to WSS loop detection. A machine is considered to be in a 'constant WSS loop' if it has used at least the percentage of total CPU specified with the MINCPU argument, has executed any number of I/Os, *and* when the number of pages in its working set has not changed by more than the WSSDELTA value between consecutive measurements, *and* when it has been found in either the dispatch or eligible list at the end of each interval and was never found dormant or test-idle by user state sampling.

**ALL**

Specifies that machines are to be forced regardless of logged-on or disconnected state.

**DISconn**

Specifies that machines are to be forced only if they are disconnected. Machines which are logged on to a terminal will not be forced.

**NODISC**

Specifies that machines are to be forced only if they are logged on to a terminal. Disconnected machines will not be forced.

**TEST**

Starts user monitoring in test mode: looping machines, or machines which have been idle, will receive warning messages that they are about to be forced when the set warning threshold has been exceeded, but instead of actually forcing the machine when the FORCE threshold has been reached a message will be generated to \*, telling you which machines would have been forced if the facility had been really activated.

This test mode is intended to help you in testing the completeness of your exclude list (i.e. the list of users excluded from FORCE processing by corresponding 'FC FORCEUSR EXCLUDE ...' commands). Run the facility in test mode first, and add all machines which should not have been forced to the exclude list, before activating actual FORCE processing.

**LImit t1/t2**

Defines the periods, in minutes, for which a virtual machine may loop, or remain idle, before some action is taken:

**t1**

Is the period after which a warning message will be sent to the machine to inform the user of the current status, and, if t2 is not zero and TEST mode has not been set, of the time that remains to take corrective action before the machine will actually be forced. One of the following messages will be sent:

Machine has been idle for over *t1* min., will be forced in *nn* min.

Machine in CPU loop for over *t1* min., will be forced in *nn* min.

Machine in I/O loop for over *t1* min., will be forced in *nn* min.

Message number 338 will also be generated to inform you of the machine's status (not generated for IDLE users unless FORCE interval has been set to 0).

**t2**

Is the period after which the machine is to be forced off the system. Message number 339 will be generated to inform you that a machine has been forced, and of the reason for doing so.

Acceptable values for the warning threshold *t1* range from 1 to 254 minutes and for the FORCE threshold *t2* from 0 to 255 minutes, and *t2* must at least be equal to the warning threshold *t1* if *t2* is not 0. Generation of the warning message to the user and of message number 338 will be suppressed if the *t2* interval is set equal to the *t1* interval.

A value of '0' for *t2* Indicates that only warning messages are to be generated, but no machines are to be forced.

**NOTdisk**

Indicates that users are only to be forced if they do not own temporary minidisks. Users with T-disks will only receive warnings, but will not actually be forced.

The same logic is applied to users with VDISKS (virtual disks in storage).

**Weight *www***

Sets the 'weight' *www* of the exception message when the selected case is encountered, where *www* can be any number in the range 0 to 255. A default value of 10 will be assumed when the WEIGHT argument has not been specified.

The WEIGHT specification is relevant only when a machine's exception data are sent to a central monitor machine (activated with the REMSEND argument of the FC MONCOLL command). All applicable weights will then be added, and the result divided by 10 to obtain the exception severity code which will be used to set the color of the load bar in the central monitoring machine.

The 'weight' number for a 'loop' case will be added once per looping user, i.e. it may be added several times if multiple users are looping.

**OFF**

Disables monitoring of users and FORCE processing for the selected function (CUPLOOP, IDLE, IOLOOP or WSSLOOP).

**Usage**

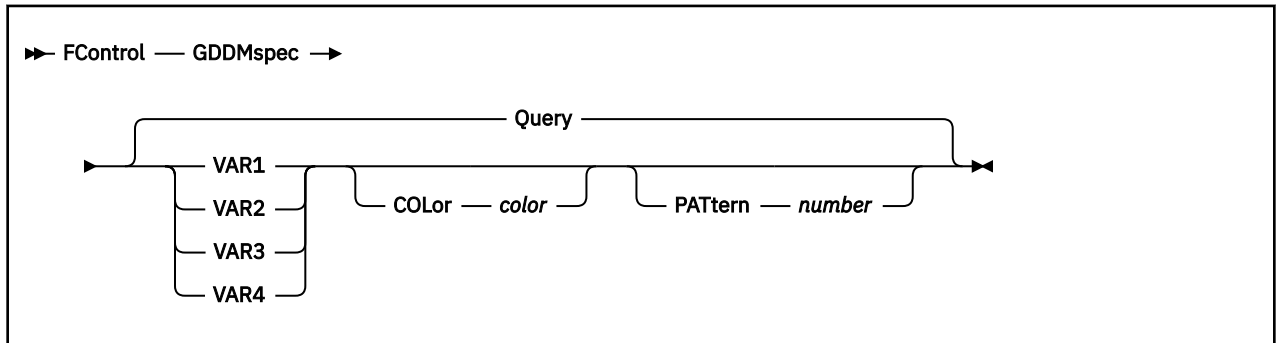
1. See the "User Monitoring" section in the [z/VM: Performance Toolkit Guide](#) for more information on user monitoring, and for hints on the use of the FC FORCEUSR command.

# FCONTROL GDDMSPEC

## Purpose

The FCONTROL GDDMSPEC subcommand allows setting colors and shading patterns which are to be used when creating GDDM graphics.

## Format



## Parameters

### Query

Lists colors and patterns defined for variables 1 to 4. This is the default if no further argument is entered.

### VAR1-VAR4

Specifies whether the remaining arguments are intended for the first, second, third or fourth Y-variable

### COLor

Indicates that the color for the selected variable is to be changed. The next argument *color* will be assumed to be the color specification; it must be one of the following:

#### BLUe

For blue

#### RED

for red

#### PINK

For pink

#### GREen

For green

#### TURquois

For turquoise

#### YELLow

For yellow

#### WHItE

For white

#### DARKblue

For dark blue (blue)

#### ORAnge

For orange (red)

#### PURple

For purple (red)

**DRKgreen**

For dark green (green)

**DRKTurqu**

For dark turquoise (turquoise)

**MUStard**

For mustard (yellow)

**GRAY**

For gray (white)

**BROwn**

For brown (red)

The colors used for simple plots (using native screen colors and plot characters shown with reverse video) are adapted also when the colors for GDDM graphics are changed, but the whole range of colors is not available then. The colors used for simple plots are shown in parenthesis where the original selection cannot be used.

**PATtern**

Indicates that the default shading pattern for the selected variable is to be changed.

***number***


Must be a value between 0 and 16, corresponding to the numbers used when setting a pattern using the GSPAT macro. See the appropriate GDDM Programming Reference manual for examples of the GDDM-defined shading patterns.

# FCONTROL HMA

## Purpose

The FCONTROL HMA subcommand configures the High Memory Area (HMA) storage residing above the 2G line, that is, divides it into the internal page pool and print buffer. It can also display the current HMA settings and usage statistics.

## Format



## Parameters

### Query

Displays the current HMA configuration and its usage statistics. This is the default action if no arguments are entered with the subcommand.

### PAGE

Specifies that the required percentage of the HMA storage between 2G and 4G is to be allocated to the internal page pool. The rest of the HMA storage will be used as the print work buffer.

### Query

Inserts the current HMA page pool percent setting into the command line. This is the default if no additional arguments are entered.

### nn

Defines the HMA page pool percentage to be set. *nn* can be any decimal number between 0 and 100. The default is 0 if this value is not set.

## Usage

1. The FC HMA PAGE *nn* subcommand must be included in the Performance Toolkit initialization control file (FCONX \$PROFILE for non-batch invocation or FCONX SETTINGS for batch run) to be effective. The HMA page pool percentage cannot be changed after the program has been initialized.
2. To minimize the risk of conventional memory below 2G shortage condition and for performance reasons, it is recommended that the HMA page pool size be defined as a nonzero value. The settings shipped with the Performance Toolkit are virtually optimal: the HMA is defined as 2G.2G, and FC HMA PAGE 10 sets the page pool as 10% \* 2G = 200M. This is sufficient to cover almost every configuration.

## Example

1. Here is an FC HMA QUERY example:

```

HMA storage 2048M.2048M usage statistics:

<----- Page pool ----->
<----- at 00000000_80000000 ----->
      Total  Alloc  Used   Free
Pages  52428    18    18     0
Memory  205M    72K    72K     0
%       -----  .034  .034  .000

<----- Print buffer ----->

```



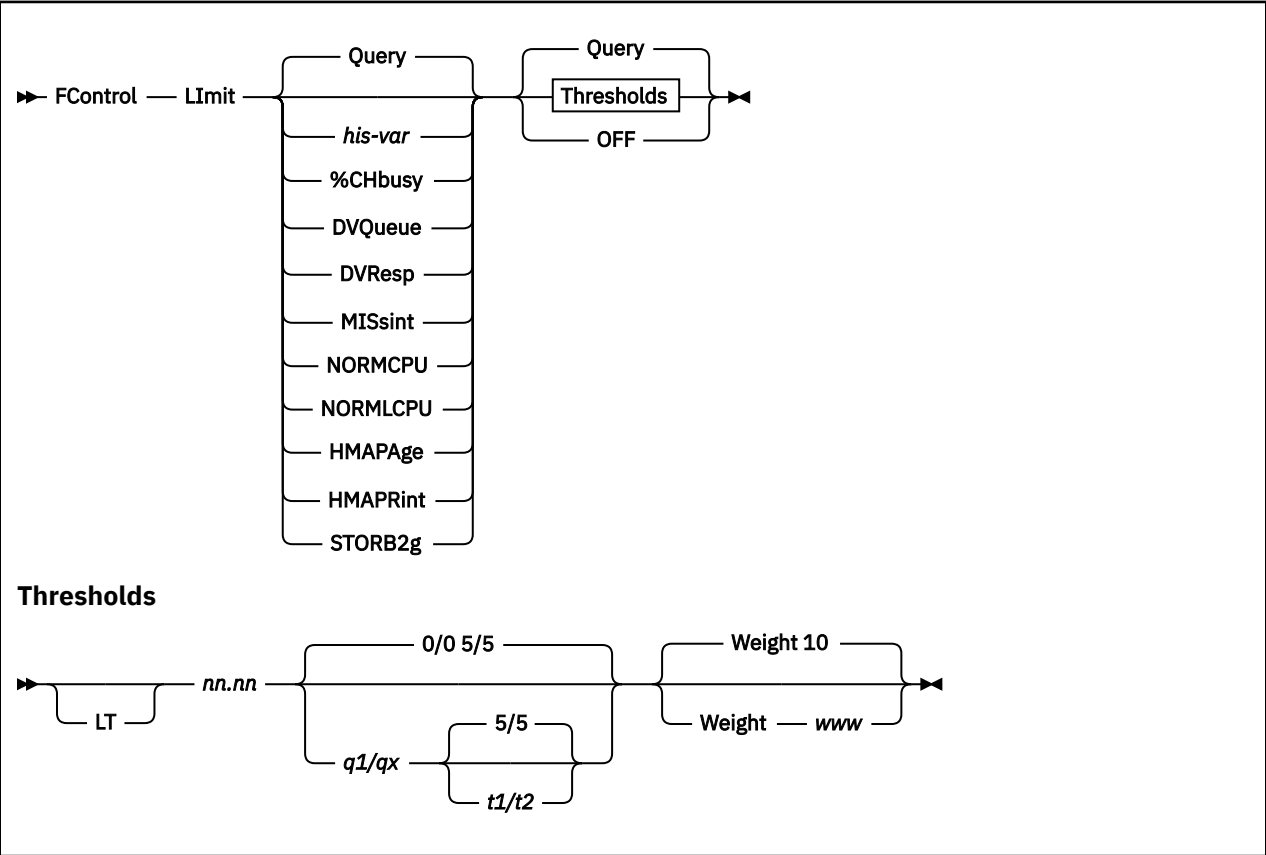
```
<----- at 00000000_8CCCC000 ----->  
      Total  Alloc  Used  Free  
Blocks 13422k  8994  8994    0  
Memory 1843M 1265K 1265K    0  
%      -----  .067  .067  .000
```

# FCONTROL LIMIT

## Purpose

The FCONTROL LIMIT subcommand can define thresholds to Performance Toolkit for many performance indicators. When a limit has been exceeded for a certain amount of time, messages are issued to inform you of the problem, so that you can take the necessary actions. Thresholds can be defined only if permanent performance data collection has been activated previously. See [“FCONTROL MONCOLL”](#) on page 53 (and specifically the CPMON parameter).

## Format



## Parameters

### *his-var*

Is a HISTLOG/HISTSUM performance variable for which a threshold is to be set or displayed. Thresholds can be set for any of the HISTLOG and HISTSUM performance variables. For more information about these performance variables, see [Appendix D, “HISTLOG and HISTSUM File Records,”](#) on page 981.

### **CPU**

The total CPU load

### **%VEC**

The total vector facility load

### **IO/S**

The total I/O rate

### **PG/S**

The total page rate (reads and writes)

**ACT**

The number of users active in an interval

**C1ES**

85% threshold of class 1 elapsed time

**TR-Q**

The number of interactive users in queue

**NT-Q**

The number of non-interactive users in queue

**%PQ**

The percentage of in-queue users found in page wait

**%IQ**

The percentage of in-queue users found in I/O wait

**%EL**

The percentage of in-queue users found in resource wait

**%ST**

The percentage of pageable DPA in use

**%PGSL**

The percentage of system PAGE area on DASD in use

**%SPSL**

The percentage of system SPOOL area in use

**Note:** If you want to use a variable that begins with the pound sign (#), which would otherwise be interpreted as the logical line-end character, you must precede it with a double quotation mark ("), which is the logical escape character. For example, for the variable #CPU, you must specify "#CPU.

Note that these characters *cannot* be used when accessing Performance Toolkit remotely. For more information, see the "Logical Line End and Escape Characters" section of the "General System Operation" chapter in the *z/VM: Performance Toolkit Guide*.

Message number 315 is displayed when the set limit is exceeded (or, if the LT argument has been specified, if the measured value is lower than the threshold).

**%CHbusy**

Indicates that a threshold for the channel-busy percentage is to be set or displayed. Message number 318 is displayed when the set limit is exceeded.

**DVQueue**

Indicates that the I/O request queue length threshold for disks is to be set or displayed. Message number 319 is displayed when the set limit is exceeded.

**DVResp**

Indicates that the device response time limit, in milliseconds, is to be set or displayed. Message number 320 indicates an exceeded limit.

Note that the DVQUEUE and the DVRESP thresholds will be monitored for DASD only, and that message 320 is displayed only if a minimum I/O activity of five I/O operations per second was found for the disk.

**MISsint**

Indicates that queue counts and message intervals for the missing interrupt detection logic are to be set or displayed. Message number 354 is displayed when a missing interrupt is detected. Note that the threshold value *nn.nn* is ignored for this case, although you *must* enter it, because the remaining arguments are positional.

**NORMCPU**

Indicates that a threshold for the normalized CPU load is to be set or displayed, that is, for the relative CPU load based on the total processing power of the system, instead of a single processor. This is a value between 0 and 100%, regardless of the number of processors.

## **NORMLCPU**

Indicates that a threshold for the normalized logical CPU load is to be set or displayed. NORMLCPU is calculated similarly to NORMCPU, except that it is based on the sum of active time and voluntary wait time instead of total elapsed time. The distinction is important for systems running in a constrained LPAR environment or second level under another VM.

## **HMAPage**

Indicates that a threshold for the HMA page pool allocation is to be set or displayed. Message number 323 is displayed when the set limit is exceeded.

## **HMAPrint**

Indicates that a threshold for the HMA print buffer allocation is to be set or displayed. Message number 324 is displayed when the set limit is exceeded.

## **STORB2g**

Indicates that a threshold for the conventional storage <2G allocation is to be set or displayed. Message number 325 is displayed when the set limit is exceeded.

## **Query**

Lists the current threshold definitions. This is the default if no further argument is entered.

## **Thresholds:**

### **LT**

Specifies that exception messages are to be generated if the measure value is *lower* than the defined limit *nn.nn*.

The LT argument is meaningful for HISTLOG and HISTSUM performance variables only; it will be ignored when specified with any other variable. For more information about these performance variables, see [Appendix D, "HISTLOG and HISTSUM File Records," on page 981](#).

### **nn.nn**

Is the maximum value above (or, if the LT argument has been specified, *below*) which notification is desired. The initial setting is OFF, that is, the respective values are not monitored or displayed.

### **q1/qx**

Are the minimum numbers of class 1 users and of class 0/2/3 users that must be in-queue to make the corresponding reading valid. The reason for this co-requisite is that limits set for normal production could be exceeded even on an otherwise very lightly-loaded system when a single user does some non-characteristic work (extremely I/O-bound save jobs during off-hours, for example). Specify *q1/qx* values high enough so that messages are issued only when the measured value is significant.

### **t1/t2**

*t1* is the minimum period (in minutes) during which the threshold must have been exceeded before a message is issued.

*t2* is the time interval (also in minutes) after which the message is to be repeated until the monitored value decreases below the set limit. The minimum that can be set is 1 minute.

The maximum value that can be set for *t1* or *t2* is 60 minutes; they both default to 5 minutes if omitted.

Consider also the length of the CP monitor SAMPLE interval when setting these time limits - where the *t1* and *t2* values are shorter than the data collection interval, threshold checking may be performed on the values of just a single sample.

## **Weight *www***

Sets the weight (*www*) of the exception message when the set threshold has been exceeded for the selected variable, where *www* can be any number in the range 0 to 255. A default value of 10 is assumed when the WEIGHT argument has not been specified.

The WEIGHT specification is relevant only when a machine's exception data are sent to a central monitor machine (activated via the REMSEND argument of the FC MONCOLL command, see "Remote Performance Monitoring Facility" in the [z/VM: Performance Toolkit Guide](#)). All applicable weights are

then added and the result is divided by 10 to obtain the exception severity code, which will be used to set the color of the load bar in the central monitoring machine.

The weight number for exception messages referring to channel or device data may be added several times if multiple channels or I/O devices exceed the set thresholds.

**OFF**

Indicates that the corresponding value is not to be monitored any longer.

**Query**

Displays the current limit for the selected value in the command input area. This is the default.

**Usage Note:**


See also the "Setup of Threshold Monitoring" section in the [\*z/VM: Performance Toolkit Guide\*](#) for hints on how to set thresholds.

# FCONTROL MAINTID

## Purpose

The FCONTROL MAINTID subcommand defines the destination address where high priority alert messages are to be sent under certain conditions. Some examples of such conditions are: insufficient storage for internal work areas, divide exceptions, etc.

## Format



## Parameters

### *userid*

Defines the user identification of the virtual machine where high priority alert messages are to be sent.

### *nodeid*

Defines the node identification for the virtual machine where high priority alert messages are to be sent. It is assumed that *userid* is a machine on the local system if no *nodeid* is specified.

### Query


Inserts and displays the current target destination for alert messages. This is the default if no *userid* is specified.

# FCONTROL MAXREC

## Purpose

The FCONTROL MAXREC subcommand specifies the number of records to be written to the '*mmdyy* CONLOG' file before issuing an intermediate 'CLOSE' command (the file will actually be closed only once Performance Toolkit would go back to its normal wait state). MAXREC is initialized to a value of ten. Please be aware that the CONLOG file may be incomplete after a VM system restart. MAXREC should be increased only if the LOG is of no great importance, but making it too small will increase general overhead.

## Format



## Parameters

### *n1*

Specifies the min. number of records to be written before closing the 'CONLOG' file. Minimum is 1 and the highest value accepted is 200.

### Query

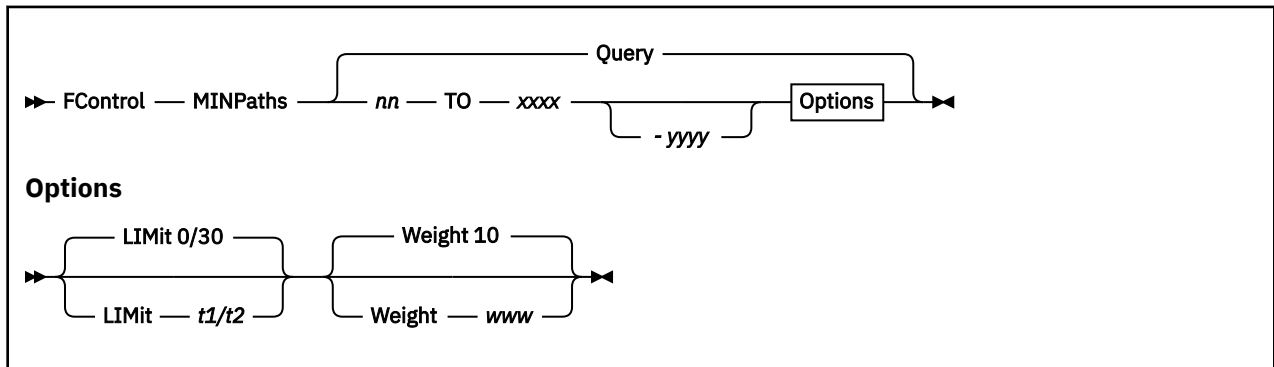
Inserts and displays the current value of MAXREC on the command line. QUERY is the default.

# FCONTROL MINPATHS

## Purpose

The FCONTROL MINPATHS subcommand can define the minimum number of channel paths which are to be active to selected I/O devices. Messages will be issued when less than the minimum number of paths are found to be active to inform you of the problem so the necessary actions can be taken. MINPATHS monitoring requires that permanent performance data collection of CP monitor data be active.

## Format



## Parameters

### `nn`

Is the minimum number of paths which should be active.

Message number 321 will be shown when less than the specified number of paths are active to the device(s) specified.

### `xxxx{-yyyy}`

Defines the device number(s) of the device(s) to which the path limit `nn` applies.

You can specify either a single device number `xxxx` or a range of devices, where `xxxx` is the first device number and `yyyy` the last device number in the range (typically all the I/O devices in a string).

### `LIMit t1/t2`

Defines the period during which the number of active channel paths must have been too low before message 321 is issued, where

'`t1`' is the initial period (in minutes) during which the number of active paths must have been too low before a message is issued.

'`t2`' is the time interval (also in minutes) after which the message is to be repeated until the number of active channel paths is no longer below the set minimum. The minimum value that can be set for `t2` is 1 minute.

The maximum value that can be set for '`t1`' or '`t2`' is 255 minutes, the values default to a setting of '0/30' if omitted.

Note that the data collection interval for the CP monitor data (upon which the MINPATHS analysis is based) can be set considerably higher than 1 minute. MINPATHS analysis will occur after each SAMPLE interval, and message 321 will be generated just after the first monitor sample where the `t1` or `t2` interval is found to have been exceeded.

### `Weight www`

Sets the 'weight' `www` of the exception message when the set threshold has been exceeded for the selected variable, where `www` can be any number in the range 0 to 255. A default value of 10 will be assumed when the WEIGHT argument has not been specified.



The WEIGHT specification is relevant only when a machine's exception data are sent to a central monitor machine (activated via REMSEND argument of the FC MONCOLL command). All applicable weights will then be added, and the result divided by 10 to obtain the exception severity code which will be used to set the color of the load bar in the central monitoring machine.

Note that the MINPATHS weights will not be added together if several devices have an insufficient number of active channel paths (e.g. a whole DASD string if a channel to the control unit is down). The maximum 'weight' value for any of the affected I/O devices is taken instead.

### **Query**


Lets the current MINPATHS definitions be displayed. 'Query' is also the default.

# FCONTROL MODEL

## Purpose

The FCONTROL MODEL subcommand allows a 3278 model 5 terminal to be switched between model 5 ('wideline') mode and model 2 mode. The command is valid on model 5 terminals only.

## Format



## Parameters

- 2** Causes the screen to be switched from model 5 mode back to model 2 mode.
- 5** Switches the screen from model 2 mode (the initial setting) to model 5 (wideline) mode.


# FCONTROL MONCACHE

## Purpose

The FCONTROL MONCACHE subcommand defines the expected cache status for selected disks.

Messages will be issued when the current status does not match the expected one so that the necessary actions can be taken, and the corresponding status field on the CACHEXT display will be highlighted (shown in red with reverse video). MONCACHE monitoring requires that permanent performance data collection of CP monitor data be active.

## Format



## Parameters

### xxxx{-yyyy}

Defines the device number(s) of the device(s) to which the cache status *cf**d* applies.

You can specify either a single device number *xxxx* or a range of devices, where *xxxx* is the first device number and *yyyy* the last device number in the range.

### *cf**d*

Are three characters which define the expected cache status, for:

**c**

General cache status

**f**

DASD fast write status

**d**

Dual copy status

The status definition characters to be entered for *c*, *f* and *d* must be one of the following:

**A**

Indicating that the function should be active

**D**

Indicating that the function should be deactivated

**\***

Indicating that the status is not to be monitored

Message 326 will be generated if the actual status does not match the one defined for the disk.

### LIMit t1/t2

Defines the period during which an unexpected cache status must have persisted before message 326 is issued, where

't1' is the initial period (in minutes) during which the unexpected status must have been found before a message is issued

't2' is the time interval (also in minutes) after which the message is to be repeated until the cache status is correct again. The minimum value that can be set for *t2* is 1 minute.

The maximum value that can be set for 't1' or 't2' is 255 minutes, the values default to a setting of '0/30' if omitted.

Note that the data collection interval for the CP monitor data (upon which the MONCACHE analysis is based) can be set considerably higher than 1 minute. MONCACHE analysis will occur after each SAMPLE interval, and message 326 will be generated just after the first monitor sample where the *t1* or *t2* interval is found to have been exceeded.

#### **Weight *www***

Sets the 'weight' *www* of the exception message when an unexpected cache status has been found, where *www* can be any number in the range 0 to 255. A default value of 10 will be assumed when the WEIGHT argument has not been specified.

The WEIGHT specification is relevant only when a machine's exception data are sent to a central monitor machine (activated via REMSEND argument of the FC MONCOLL command). All applicable weights will then be added, and the result divided by 10 to obtain the exception severity code which will be used to set the color of the load bar in the central monitoring machine.

Note that the MONCACHE weights will not be added together if several devices have been found with an incorrect cache status (e.g. a whole DASD string if there is a problem with the cache on the control unit level). The maximum 'weight' value for any of the affected disks is taken instead.

#### **Query**

Lets the current MONCACHE definitions be displayed. 'Query' is also the default.

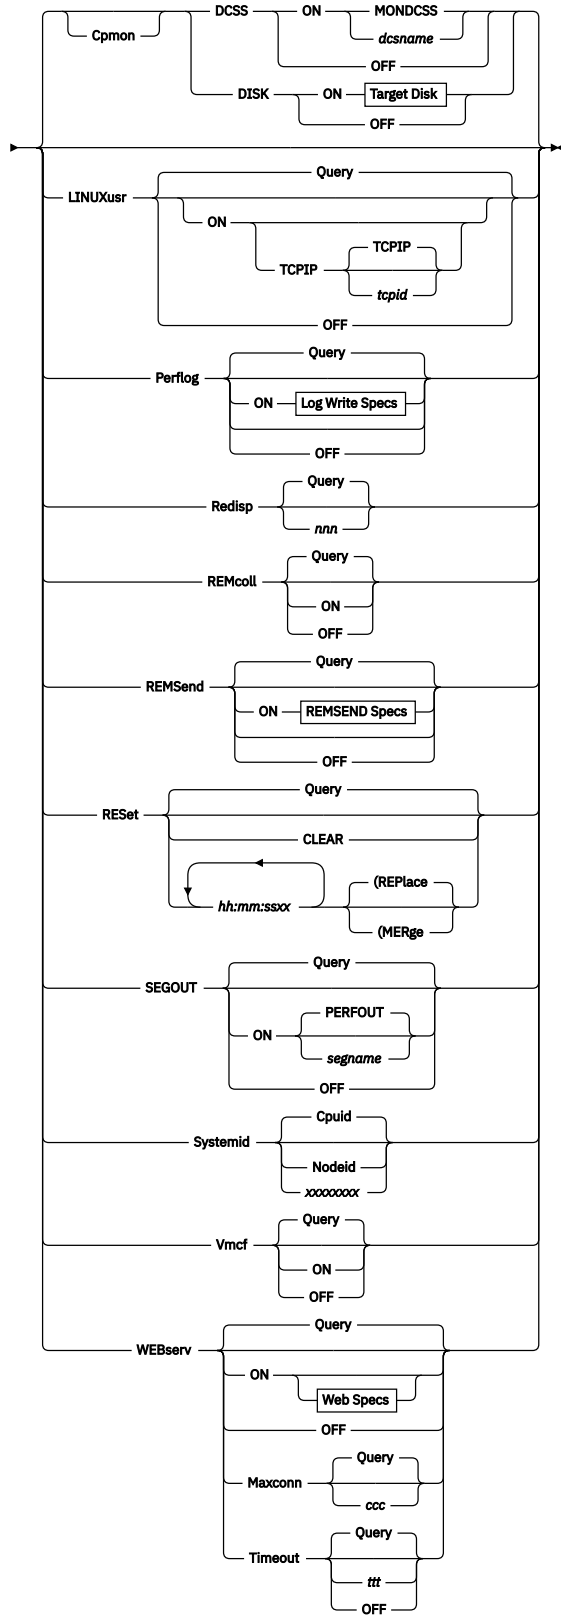
## **FCONTROL MONCOLL**

### **Purpose**

The FCONTROL MONCOLL subcommand allows you to control data collection for performance monitoring. With the exception of the REMCOLL argument, this command is valid only if the user ID is authorized to execute DIAGNOSE codes X'04' and X'90', which usually requires command privilege class E.

## Format

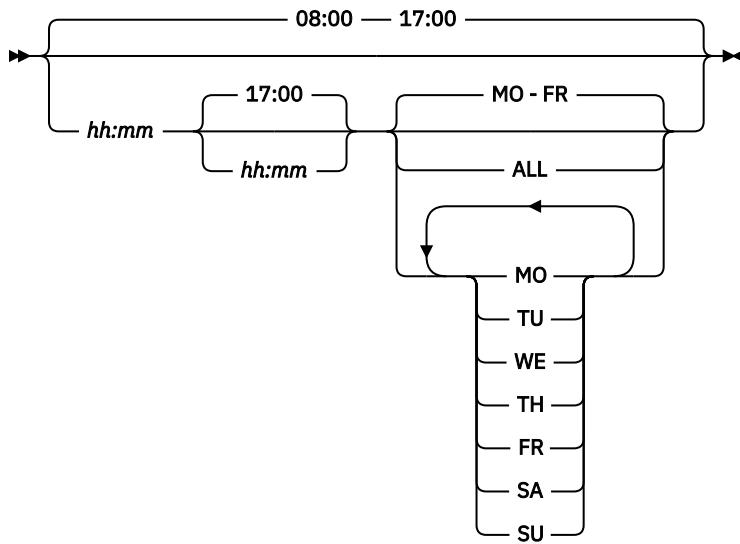
→ FControl — MONcoll →



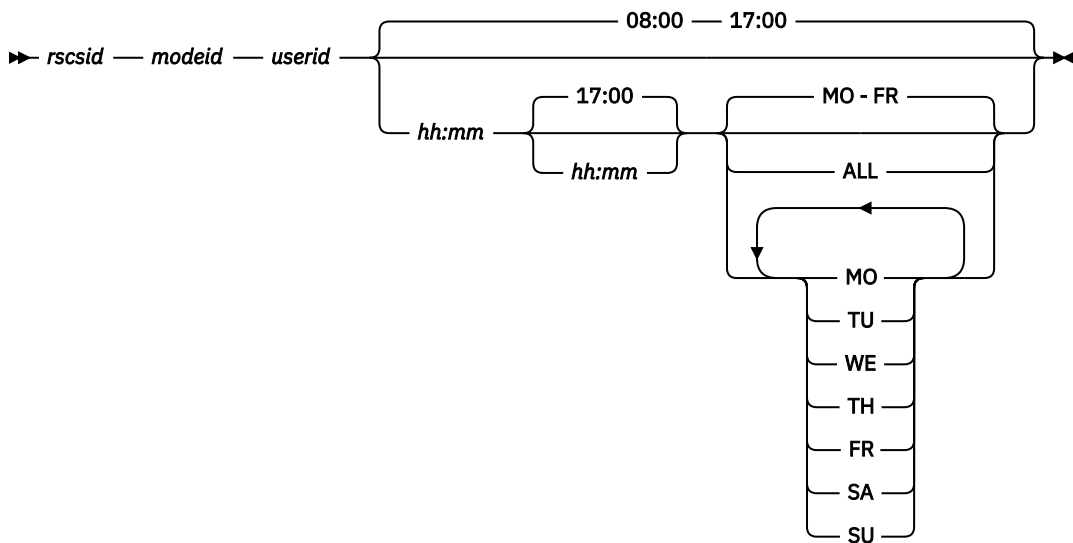
Target Disk

➤ *monw-ID* — *disknr* — *diskmode* — *fn* — *ft* ➤


### Log Write Specs



### REMSSEND Specs



### Web Specs



## Parameters

### Cpmon

controls the use of CP MONITOR data for performance analysis

### DCSS

Indicates that data extraction from the MONITOR shared segment is to be controlled

**ON *dcssname***

Indicates that data extraction from the CP MONITOR segment *dcssname* is to be activated, where *dcssname* is the segment name. The default DCSS name MONDCSS will be used if no segment name is specified.

Note that:

- Performance Toolkit will attempt to CONNECT to the segment using IUCV, the virtual machine where Performance Toolkit is running must previously have been authorized, therefore, for accessing that segment by means of a NAMESAVE *dcssname* statement in the directory, exactly as you would have to make the definition for the machine if it were to run the MONWRITE data collector module.
- Data extraction from the MONITOR shared segment requires that this segment can be loaded in the machine where Performance Toolkit is running. It may be necessary to increase the virtual storage size of that machine in order to do this.
- Performance Toolkit will automatically execute a MONITOR START command when data extraction from a shared segment is activated. When this command is successful (return code of 0), Performance Toolkit will also STOP the monitor when you either deactivate \*MONITOR data collection, or when you exit from Performance Toolkit to basic CMS and no other users are connected to the \*MONITOR facility.

Activating CP monitor data collection will let performance data be collected permanently. This mode allows automatic threshold monitoring (see “FCONTROL LIMIT” on page 42), but it will also cause some additional permanent overhead. Permanent data collection should, therefore, be activated only in one virtual machine per z/VM system (the system operator's or system programmer's machine, for example).

**OFF**

Indicates that data extraction from the MONITOR shared segment is to be stopped.

**DISK**

Indicates that MONITOR data extraction from a disk file written by the MONWRITE module is to be controlled.

**ON *monw-id disknr diskmode fn ft***

Indicates that data extraction from a MONITOR disk file is to be activated, where

***monw-id***

Is the userid of the machine which owns the disk where the MONWRITE program writes the CP MONITOR data

***disknr***

Is the virtual device number of the minidisk which contains the MONITOR data file.

Note that Performance Toolkit will attempt to LINK to this disk as virtual device number '01F1'; *the virtual machine where Performance Toolkit is running must previously have been authorized, therefore, for linking to that disk.*

***fm***

Is the file mode to be used by Performance Toolkit for accessing the disk

***fn***

Is the file name of the MONITOR data file

***ft***

Is the file type of the MONITOR data file.

**OFF**

Indicates that data extraction from the MONITOR disk file is to be stopped, and the disk released and detached.

**LINUXusr**

Controls the activation/deactivation of the TCP/IP request interface for performance data retrieval from Linux® RMF DDS interfaces.



**Query**

Displays the current interface setting on the command line. This is the default.

**ON**

Specifies that the request interface is to be activated. This is also the default. Further arguments can specify details, if the defaults are not acceptable:

**TCP/IP *tcpid***

Defines the userid of the TCP/IP service machine that provides access to the TCP/IP network, where *tcpid* is the user identification of the TCP/IP server.

**OFF**

Indicates that the TCP/IP request interface is to be deactivated.

**Perflog**

Controls the disposition of general performance data from the REDISP or REDHIST screens (see also [“FCONTROL SETTINGS”](#) on page 83).

**ON**

Specifies that all performance data inserted into the redisplay screen must also be written to file 'mmdyy PERFLAG A', and permanent data collection will be set on even if the 'FC MONCOLL ON' command has not been previously entered.

Up to three generations of the PERFLAG file will automatically be kept on the A-disk, with filetypes of PERFLAG, PERFLAG1 and PERFLAG2. These files are intended for archiving or for additional analysis by the user. Although not used for generating the default redisplay screen, they can also be selected for display by using the 'FILE *fn ft fm*' arguments of the REDHIST command (see [“REDHIST”](#) on page 268).

***hh:mm***

Defines the start and end times of the period during which performance data are to be saved on disk. The time must be entered exactly as shown, that is, with two digits for hours and minutes each, separated by a colon. The default period is from 08:00 to 17:00.

**MO-FR**

Specifies that data are to be collected only from Monday to Friday. This is also the default.

**ALL**

Specifies that data are to be collected for the whole week.

**MO TU WE TH FR SA SU**

Specifies the days of the week for which performance data are to be collected. Any combination of days is acceptable.

**OFF**

Specifies that collection of performance data on disk is to be stopped.

**Query**

Displays the current setting on the command line. This is the default.

**Redisp**

Defines the number of lines to be set aside during initialization for use as performance redisplay buffer.

**Query**

Inserts the current log buffer setting into the command line. This is the default if no further argument is entered.

***nnn***

Defines the number of lines to be set aside.

If working with the default monitor sample interval of 60 seconds Performance Toolkit needs 60 lines for one hour's data. The minimum number is 60, the maximum 14,400 lines, with a default of 720 lines.

**Note:** The log buffer space requirements can be reduced by means of letting averages for more than one sample interval be shown on each of the 'by time' log detail lines. For details see the "By-Time Log Data" description in [“Interval”](#) on page 367.

This command must be included in the initialization profile 'FCONX \$PROFILE' to be effective. The buffer size cannot be changed once the program has been initialized.

### **REMcoll**

Controls automatic collection of performance data sent from other systems via RSCS message (generated in the remote system by means of the REMSEND argument of the 'FC MONCOLL' command).

#### **Query**

Inserts the current setting into the command line. This is the default if no further argument is entered.

#### **ON**

Specifies that this machine is to act as central data collector machine for the remote performance monitoring facility, i.e. that performance data sent from another machine are to be received and inserted for use in FCONRMT mode.

#### **OFF**

Indicates that REMSEND data received from other machines are not to be collected.

### **REMSend**

Controls the automatic sending of summary performance data from this system to another machine, for use with the remote performance monitoring facility.

#### **ON**

Specifies that summary performance information is to be sent to another machine.

#### ***rscsid***

Is the userid of the RSCS machine which is to be used for data transmission

#### ***nodeid***

Is the node ID of the system to which data are to be transmitted.

Asterisks may be entered for both the 'rscs-ID' and 'node-ID' if the target machine is on the same VM system. Data transmission occurs then directly via the CP 'MSG' command, instead of via the MSG command of RSCS, and this will save some overhead.

Use of an asterisk for the RSCS machine's user ID only, with a valid node-ID specification, will let the default RSCS userid be used, as specified in the SYSTEM NETID file.

#### ***userid***

Is the userid of the virtual machine which is to receive and collect the transmitted data.

#### ***hh:mm***

Defines the start and end times of the period during which performance data are to be sent to the other machine. The time must be entered exactly as shown, that is, with two digits for hours and minutes each, separated by a colon. The default period is from 08:00 to 17:00.

#### **MO-FR**

Specifies that data are to be sent only from Monday to Friday. This is also the default.

#### **ALL**

Specifies that data are to be sent during each day of the week.

#### **MO TU WE TH FR SA SU**

Specifies the days of the week for which performance data are to be sent. Any combination of days is acceptable.

#### **OFF**

Indicates that automatic data transmission to another machine is to be stopped.

#### **Query**

Displays the current setting. This is the default if no further argument is entered.

### **RESet**

Displays the currently defined RESET times for performance data, or it allows you to clear any existing RESET times and to specify any number of new RESET times, optionally with automatic printing and/or summary/trend file creation for selected performance data.

**Query**

Displays currently defined RESET times.

**CLEAR**

Removes *all* of the currently active RESET times. Note that the initial RESET time of '00:00' will always be set again whenever the CLEAR option is issued.

**hh:mm:ss**

Defines a new RESET time. The time can also be specified as *hh:mm* only (in which case '00' seconds will be assumed) but hours, minutes and seconds must be specified with two digits each. Multiple reset times can be entered, separated by blanks. The time entered must be between 00:00:00 and 23:59:59. Reset times following an invalid argument will be ignored.

Append suffixes *xxx* to the reset time to indicate what specific action is required:

**R\_P**

Defines the start time for a new overall report period. This is the period used for the 'average' setting.

**R\_I**

Defines the start time for a new interim report period. This is the period used for the 'interim' setting, where applicable.

**R\_S**

Defines the start time for summary record generation.

**R\_T**

Resets trend data counters.

**P**

Indicates the time at which an automatic PRINT ALL command is to be issued, to print all the reports defined in file FCONX REPORTS (or in the 'reports' file defined by a previous 'FC SET REPORTS ..' command).

**I**

Indicates the time at which an automatic print command is to be issued to print all the INTERIM reports defined in file FCONX REPORTS (or in the 'reports' file defined by a previous 'FC SET REPORTS ..' command).

Note that automatic printing of INTERIM reports in fixed intervals can also be achieved by means of the 'FC SET INTERIM ..' command.

**T**

Indicates the time at which new trend records are to be generated. The type and number of trend records to be generated can be defined in file FCONX TRENDREC, or in the TRENDREC file defined by a 'FC SET TRENDREC ..' command.

**S**

Indicates the time at which new summary records are to be generated. The type and number of trend records to be generated can be defined in file FCONX SUMREC. or in the SUMREC file defined by a 'FC SET SUMREC ..' command.

If no suffixes are appended the default action is a full reset for all periods, i.e. the equivalent of appending all of the 'reset' suffixes R\_PR\_IR\_TR\_S

**(MERge**

Indicates that the actions specified for a specific time are to be merged with any other previously specified actions for the same time, that is, the previously-specified actions and the new actions will be executed.

**Note:** When using the FC MONCOLL RESET command, the (MERGE option should *not* be specified when setting the time to 00:00:00, as this is the default RESET time. The (REPLACE option (the default) should instead be used to set/replace any RESET times to 00:00:00.

**(REPlace**

Indicates that the reset actions specified for a specific time are to replace any previously specified actions for the same time, that is, only the new actions will be executed. This is the default.

If print output has been selected together with reset, average data will be calculated just prior to resetting, and these average values will be printed. Message number 371 will be shown to indicate the switch to average displays if Performance Toolkit is in monitor mode at this time.

Multiple FC MONCOLL RESET commands can be entered: all of the reset times entered since the last CLEAR will be active. An initial RESET time of '00:00' will always be set, and this time will also be set again whenever the CLEAR option is issued.

### **SEGout**

Indicates that the formatted output collectors are to be turned on or off, where:

#### **ON**

Will turn on the formatted output collectors.

#### ***segname***

Indicates the name of the shared segment where the calculated values from the Performance Toolkit reports will be stored. PERFOUT is the default segment name.

**Note:** A default PERFOUT segment, with a default size, is defined during system installation. If you wish to define a larger size for this segment, or to define the segment with a different name, please consult the Performance Toolkit Program Directory.

#### **OFF**

Will turn off the formatted output collectors.

#### **Query**

Displays the current setting on the command line. This is the default if no further argument is entered.

### **Systemid**

Defines the type of system-identifier which is to be inserted into the header line of all performance monitor displays.

#### **Cpuuid**

Will cause the CPU-ID to be inserted. This serial number is usually the real hardware CPUID, but it can be overwritten by directory statement or by the CP SET CPUID command.

#### **Nodeid**

Specifies that the RSCS node-ID of the system is to be used as identifier. If many different systems are to be monitored by the same person, use of the RSCS node-ID allows much easier identification of the systems than the CPUID.

The RSCS node-ID is determined by looking for the corresponding entry in the SYSTEM NETID file. It cannot be set if no entry exists.

#### **xxxxxxxx**

Specifies that the string xxxxxxxx is to be used as identifier.

### **Vmcf**

Indicates that the VMCF interface is to be activated or de-activated, where

#### **ON**

Activates the VMCF interface.

#### **OFF**

Deactivates the VMCF interface.

#### **Query**

Displays the current setting on the command line. This is the default if no further argument is entered.

The VMCF argument is valid only when permanent data collection has previously been activated.

### **WEBServ**

controls the activation/deactivation of the Web Server function for performance data retrieval via the Internet and standard Web Browsers. It also allows to set the maximum allowed number of web connections and the timeout for inactive web sessions.

**ON**

Specifies that the Web Server interface is to be activated. Further arguments can specify details, if the defaults are not acceptable:

**SSL**

Specifies that a secure connection will be used.

**TCPIP *tcpid* {*port*}**

Defines the userid of the TCP/IP service machine that provides access to the TCP/IP network, and, optionally, also the port number to be used. *tcpid* is the user identification of the TCP/IP machine, and *port* is the port number.

If not specified, the default port number of 81 will be used, and if the name of the TCP/IP service machine is not specified either it is assumed that the standard name TCPIP can be used. Note that the port number argument is positional. The TCPIP userid must also be entered if a port number is to be defined.

**OFF**

Indicates that the Web Server interface is to be deactivated.

**Query**

Displays the current setting. This is the default if no further argument is entered.

**Maxconn**

Controls the maximum allowed number of web connections to Performance Toolkit. The default is 10 connections.

***ccc***

Sets this number as *ccc*, where *ccc* can be in the range of 1-999.

**Query**

Displays the current setting on the command line. This is the default if no further argument is entered.

**Timeout**

Controls the time period after which the user of an inactive session will be requested to enter his logon credentials in order to continue working. The default is 30 minutes.

***ttt***

Sets this timeout as *ttt* minutes, where *ttt* can be in the range of 1-1440.

**OFF**

Disables timeout checking.

**Query**

Displays the current setting on the command line. This is the default if no further argument is entered.

**Usage**

- For PERFLOG:

1. To generate a record at each monitor interval from 9:00 am to 5:00 pm each day, issue:

```
FC MONCOLL PERFLOG ON 09:00 17:00 ALL
```

2. To generate a record at each monitor interval throughout the day, issue:

```
FC MONCOLL PERFLOG ON 00:00 00:00 ALL
```

- For RESET:

1. All the reset and action specifications must be directly appended to the time, i.e. use the string 10:00R\_P to force a reset of the 'average' period at 10:00:00.
2. The 'P', 'I', 'T' and 'S' suffixes also imply the corresponding reset actions 'R\_P', 'R\_I', 'R\_T' and 'R\_S' but reports will be printed, and history file records built, *before* resetting the counters.

3. If multiple reset or action specifications are to be entered for the same time they can simply be appended after each other in any sequence, without any intervening blanks. You would use the string 11:00:00R\_STP to start a new 'summary' file period at 11:00:00 and to also build new trend records and print reports for the 'trend' and 'average' periods ending at 11:00:00.
4. At least one reset action per day will automatically be enforced for each of the P, I, T, and S periods. This default reset occurs at midnight for any period for which no explicit reset time has been entered.

See the "Resetting Performance Data / Defining Periods" section in the [\*z/VM: Performance Toolkit Guide\*](#) for a general description of the different periods and their uses.

- For WEBSERV:


See the "Preparing the Performance Toolkit, Internet Interface" section in the [\*z/VM: Performance Toolkit Guide\*](#) for a detailed explanation of using the Internet interface.

# FCONTROL MSGCLEAR

## Purpose

The FCONTROL MSGCLEAR subcommand specifies whether CP messages and Performance Toolkit action messages are to be kept on the display when the screen is cleared (CLEAR command, CLEAR PF-key or PA2-key).

## Format



## Parameters

### Query

Displays the current MSGCLEAR setting. This is the default if no further argument is entered.

### ALL

Specifies that *all* screen lines are to be cleared, regardless of their contents (i.e. including CP messages and Performance Toolkit 'action' messages). 'ALL' is the initial setting.

### NOCpmsg

Indicates that CP messages are *not* to be removed from the screen by the CLEAR function. CP messages can then be cleared only by means of the 'DELETE' or 'REPLY' commands.

### NONE

Specifies that neither CP messages nor Performance Toolkit action messages are to be cleared from the screen by the CLEAR function. Both CP messages and action messages can then be cleared only by means of the 'DELETE' or 'REPLY' commands.

## FCONTROL MSGWAIT

### Purpose

The FCONTROL MSGWAIT subcommand specifies how many minutes a message may be left pending on the screen before the operator gets additional notification.

### Format



### Parameters

#### ***n1***

Specifies the number of minutes that may elapse before an additional information message is displayed. Minimum is 1 and the highest value accepted is 60.

#### **OFF**

Disables the function. This is also the initial setting.

#### **Query**

Inserts and displays the current MSGWAIT setting on the command line. This is the default.



# FCONTROL NUMBERS

## Purpose

The FCONTROL NUMBERS subcommand specifies whether line numbers are to be shown on the left hand side of the screen output area or whether they are to be omitted.

## Format



## Parameters

### ON

Indicates that line numbers are to be inserted.

### OFF

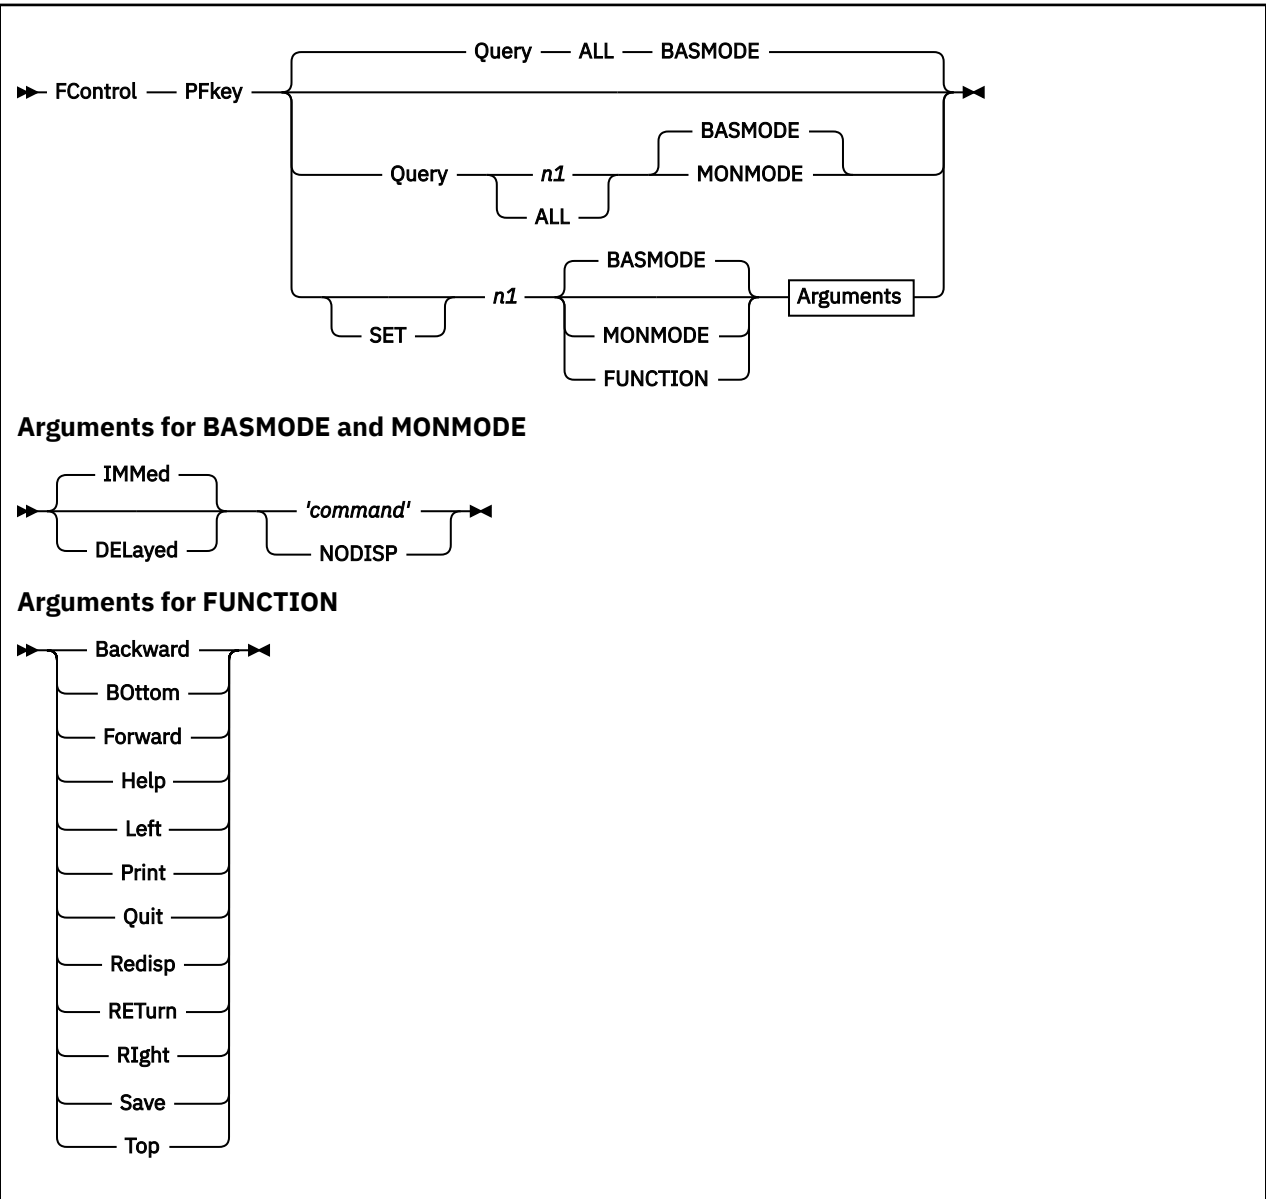
Indicates that no line numbering is desired.

# FCONTROL PFKEY

## Purpose

The FCONTROL PFKEY subcommand inserts the command text specified into the internal PF-key tables of Performance Toolkit (separate tables are kept for basic mode, re-display mode and performance monitor mode). Pressing a PF-key will cause the command text from the table corresponding to the current mode to be executed as if it had been entered on the command line.

## Format



## Parameters

### SET

Indicates that a new function is to be assigned to a PF-key. This is optional.

## **Query**

Indicates that the current PF-key settings are to be displayed. The reply will be inserted into the command line if a single PF-key is queried, or into the output area of the basic mode screen if all PF-keys are queried.

## **n1**

Specifies the number of the program function key that is to be defined or queried. Valid PF-key numbers are 1-24.

## **ALL**

Indicates that *all* PF-key settings are to be displayed. 'ALL' is valid only in conjunction with the 'QUERY' argument. Note that the output for the 'ALL' argument will be inserted into the basic mode screen; it cannot be viewed while in redisplay or performance monitor mode.

## **BASMODE**

Specifies that a PF-key function for use in 'basic' mode is to be assigned or queried.

## **FUNCTION**

Specifies that a *common* PF-key function is to be assigned (e.g. HELP, QUIT, or scrolling functions for which the same PF-keys are to be used in basic, redisplay and performance monitor mode).

PF-key FUNCTIONS will automatically be assigned to *all* PF-key tables where the corresponding function is meaningful, overlaying any commands which had previously been assigned to the corresponding PF-key. If the same function had previously been assigned to another PF-key within the same range of 1-12, or 13-24, the previous PF-key function assignment will automatically be cleared; i.e. a certain function can be assigned only to one PF-key in a 1-12 and/or 13-24 group.

## **MONMODE**

Specifies that a PF-key function for use in performance monitor mode is to be assigned or queried.

Arguments for defining basic mode or monitor mode PF-keys:

## **IMMed**

Command is to be executed immediately when the PF-key is pressed. This is also the default value. Argument valid for defining PF-keys in basic mode and performance monitor mode only.

## **DElayed**

Command text is to be inserted in the command line when PF-key 'n1' is pressed so that it can be modified before pressing the 'ENTER'-key. This argument is valid for defining PF-keys in basic mode and performance monitor mode only.

## **'command'**

The command to be executed or inserted into the command line (do not enter the apostrophes). It will be translated to upper case.

When using the 'DElAYED' option of the PFKEY command you can also insert a '¢'-sign (or whatever sign occupies the X'4A' code position on your keyboard) anywhere in the command text. The cursor will then be positioned at the location of the '¢'-sign after insertion of the text into the command line (and the '¢'-sign removed) so that the command text can easily be modified. The cursor will be positioned at the start of the command line if no '¢'-sign has been supplied.

## **NODISP**

If the command text is 'NODISP' the corresponding PF-key is set up for the NODISP-function which allows you to control the visibility of data entered after a CMS prompt ('ENTER:' in front of the command line). The NODISP-function does not influence the visibility of normal commands entered while the arrow '==>' is shown in front of the input area. After pressing the PF-key this input will remain invisible while it is entered, and it will not be written to the console log. Only the normal input area is set up for invisible input, data entered on the command line extension (bottom line with PF-key assignments) will remain visible.

The command line will automatically be reset to normal display after entering one line of data, so that the NODISP function will have to be activated again if several consecutive input lines are to remain invisible. You can also reset the command line to normal display by pressing the PF-key a second time ('flip-flop'-action of the NODISP function) when the function has erroneously been activated.

This feature allows the entering of passwords in basic mode for use, for instance, by EXEC procedures without creating a security exposure. It is not meaningful when used in performance monitor mode.

Arguments for defining FUNCTION PF-keys:

Note that the functions specified will automatically be assigned to all applicable tables, as defined by the characters 'B' (for basic mode), 'R' (for re-display mode) and 'M' (for performance monitor mode) in the following list.

**Backward**

R+M, for scrolling backward (towards beginning of data).

The BACKWARD function is initially assigned to PF-keys 7 and 19.

**Bottom**

R+M, for scrolling to bottom (to end of data).

The BOTTOM function is initially assigned to PF-keys 5 and 17.

**Forward**

R+M, for scrolling forward (towards end of data).

The FORWARD function is initially assigned to PF-keys 8 and 20.

**Help**

B+R+M, for displaying HELP information.

The HELP function is initially assigned to PF-keys 1 and 13.

**Left**

R+M, for shifting display window to the left.

The LEFT function is initially assigned to PF-keys 10 and 22.

**Print**

M, for printing current performance display data.

The PRINT function is initially assigned to PF-keys 9 and 21.

**Quit**

B, to exit from Performance Toolkit.

The QUIT function is initially assigned to PF-keys 3 and 15.

**Redisp**

B+R+M, for re-displaying console log (basic and re-display mode), or general system performance data (performance monitor mode).

The REDISP function is initially assigned to PF-keys 2 and 14.

**RIght**

R+M, for shifting display window to the right.

The RIGHT function is initially assigned to PF-keys 11 and 23.

**RETurn**

B+R+M, for returning to previous level or mode.

The RETURN function is initially assigned to PF-keys 12 and 24.

**Save**

M, for saving GDDM graphics on disk (valid on graphics displays 'GRAPHxxx' only).

The SAVE function is initially assigned to PF-keys 6 and 18.

**Top**

R+M, for scrolling to top (to beginning of data).

The TOP function is initially assigned to PF-keys 4 and 16.

## Usage

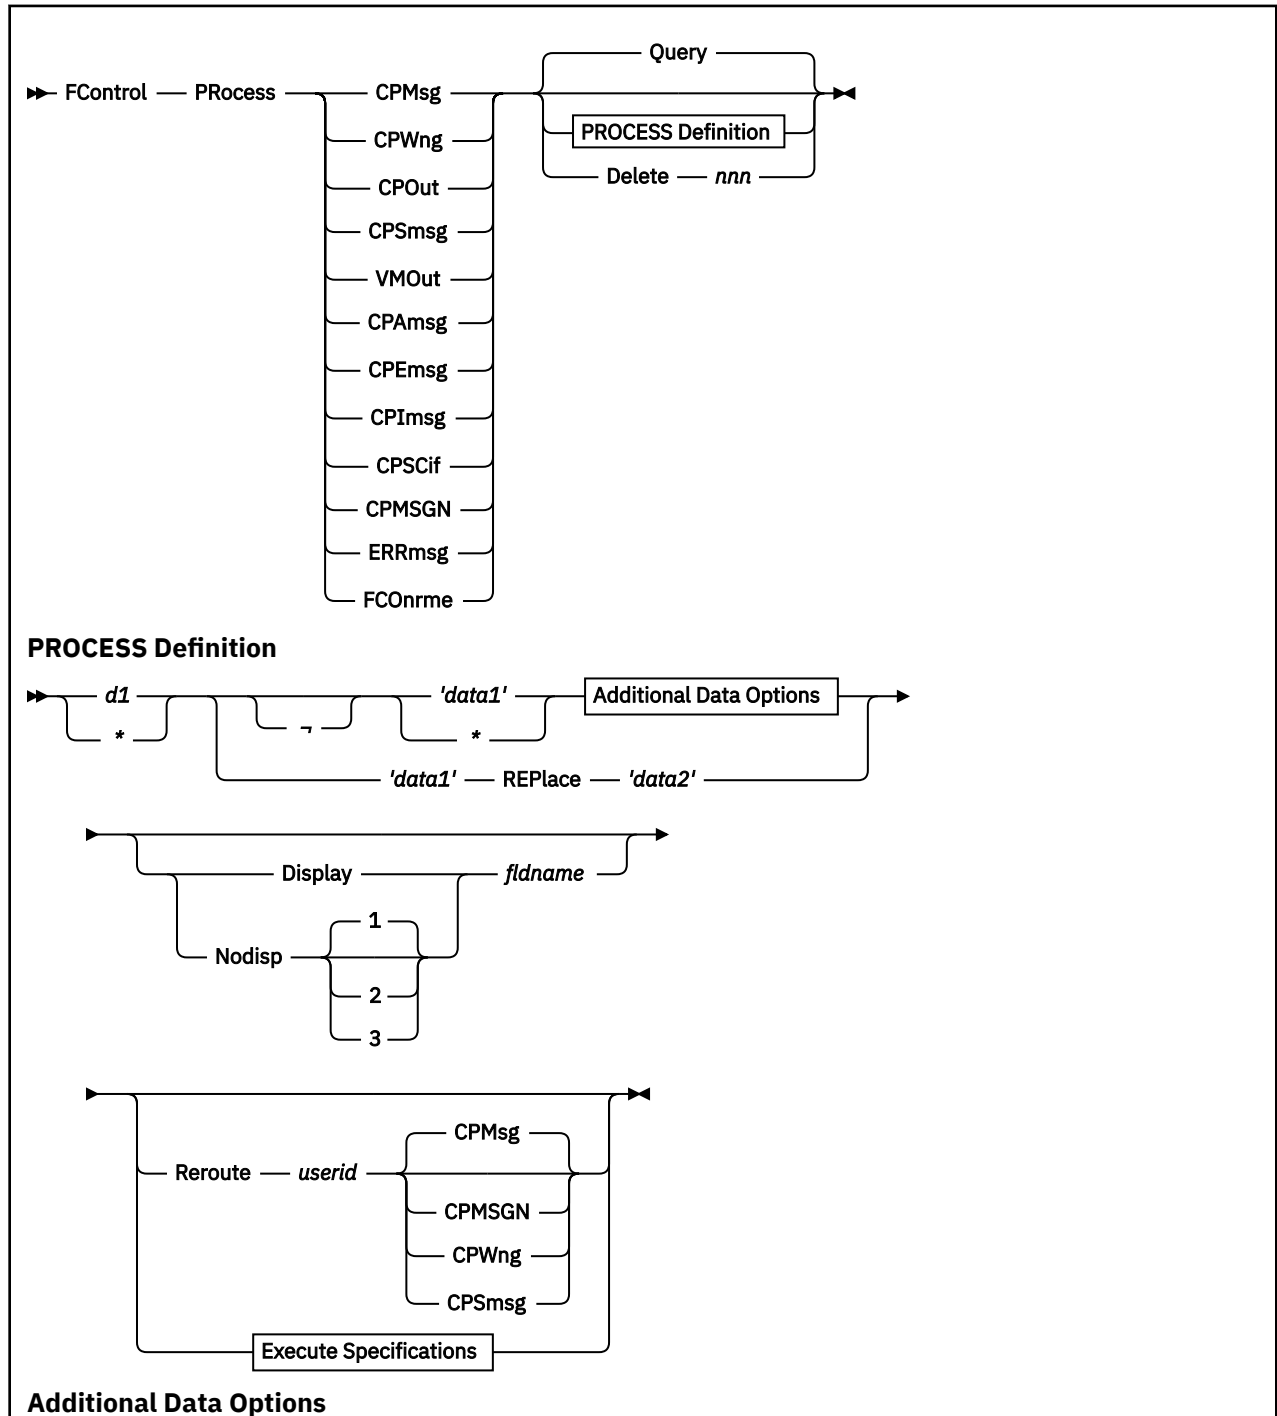
1. The PF-key definitions for the FCONXREP procedure are not affected by the 'FC PFKEY' command; you will have to modify this procedure if you want to use different PF-keys.

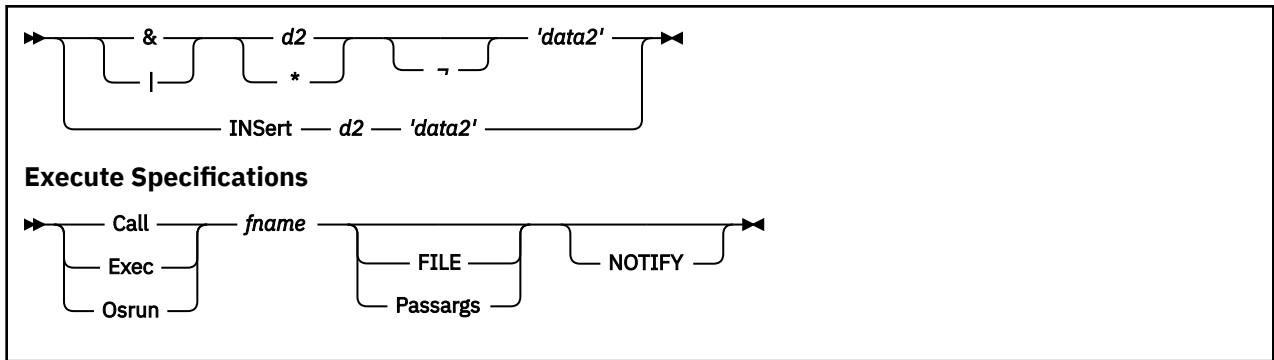
# FCONTROL PROCESS

## Purpose

This command allows definition of additional processing that is to be done when lines with certain characteristics are received.

## Format





## Parameters

The first argument indicates which type of output line is to be processed:

### CPMsg

CP message

### CPWng

CP warning

### CPOut

CP command output line

### CPSmsg

CP SMSG. Note that SMSGs will not automatically be received in IUCV format. You will have to enter the command 'SET SMSG IUCV' first if you want Performance Toolkit to receive and handle SMSGs (can also be included in the initialization file FCONX \$PROFILE).

### VMOut

Output line generated by the virtual machine (that is, CMS)

### CPAmsg

CP message line that is treated as Performance Toolkit action message (i.e. contains preceding carriage return characters).

See the "Action Messages" section in the "General System Operation" chapter of the [z/VM: Performance Toolkit Guide](#) for more information.

### CPEmsg

CP error message

### CPImsg

CP information message

### CPSCif

Output line from another machine which we received via the Single Console Image Facility (SCIF)

### CPMSGN

CP noheader message

### ERRmsg

Performance Toolkit 'A'-type message

### FCONrme

Performance Toolkit exception message from a remote system for central monitoring

**Note:** Since these exception messages are not intended to be shown on the basic mode screen, you cannot refer to it, or to the CP console log, to determine the original message layout, and the displacements to be used for trapping a message. The format used is as follows:

```
text:nodeid : xxxxx xxxxxxxx xxxxxxxxx
```

```
A A
```

```
disp.: 0 10
```

where:

***nodeid***

Is the RSCS node-ID of the originating system

***xxxxx***

Is the complete message text to be placed into the system's EXCPLOG file

that is, the format is similar to SCIF messages from another user on the same system.

**Setting new PROCESS definitions:** specify whether only certain lines should be processed which contain special data at specific locations:

***d1***

Is a number which indicates the displacement to the string. Enter an asterisk if any displacement is acceptable.

Please note that the displacement usually refers to the original line received from CP. Further processing by Performance Toolkit may insert additional data (message number, time stamp) in front of the line, it may remove data (userid of SCIF messages after 'FC SECUSER' command) from the line before it is shown on the screen, or the line may have to be split before it can be shown. The CP console log contains the original line format and should be referenced in case of doubt.

Exceptions:

- CP messages (CPMSG) and warnings (CPWNG) are written to the CP console log by Performance Toolkit preceded by a '>' sign which is to be ignored when calculating the displacement.
- Performance Toolkit 'A'-type messages are also written to the console log, but they will be preceded by a '>' sign and a time stamp. The displacement must be calculated from the start of the actual message line 'FCXxxxxnnx ...'.

↯

Indicates that additional processing is required only if the string is **not** found at the specified displacement. The '↯' operator can also be entered as 'NOT'.

***'data1'***

The data string that should be looked for (requires quotes at the beginning and end of the string and can include blanks, but it must not exceed a maximum length of 20 characters).

An asterisk can be specified to indicate that no data search is to be made, that is, that all output lines of the type specified above should be treated in the same way. Note that specifying additional logical operators (& or |) does not make sense if one of the data strings has been defined with an asterisk. Although the definition will be accepted as long as the syntax is correct, the line will always be processed if 'any data' has been specified for either the data1 or data2 field.

***& |***

Logical operators which indicate that a second string is to be searched for and that the line is to be processed if either both (for '&') or any (for '|') of the strings have been found. You can also code 'AND' for '&' and 'OR' for '|'.

***d2***

Is a number which indicates the displacement to the second data string.

↯

Indicates that the second string should not be found.

***'data2'***

The second data string that should be looked for.

**INSert**

Indicates that string 'data2' is to be inserted into the line at displacement d2. The remainder of the original line will be shifted to make room for the inserted string. Note that the '↯' or 'NOT' operators are not allowed in conjunction with the data2 string after the INSERT argument.

**REPlace**

Indicates that part of the line contents are to be replaced as follows:



- The string '*data1*' is searched for. If found, the beginning of the selected line, up to and including the string '*data1*', will be removed.
- When a second data string '*data2*' has been specified, this second string will be inserted in front of the remaining text. Otherwise just the remainder will be shown.

Note that the '-' or 'NOT' operators are not allowed in conjunction with the REPLACE argument, and that no displacement can be specified for the second data string.

This form of the FC PROCESS command is especially useful for removing message prologues, such as the normal message header:

```
'hh:mm:ss MSG FROM userid: '
```

from messages generated by the RSCS machine, or the even longer prologues received when working in a machine which acts as logical operator for remote VM systems.

### Display

The output line is to be displayed. Since this is the default anyway, you have to specify it only if the line is to be displayed as some other line type (see '*fldname*').

### Nodisp

The output line is not to be displayed. It will, however still be written to the console log file of Performance Toolkit so that it can be re-displayed. Such lines will receive a '>' prefix character on the redisplay screen to indicate that they were suppressed on the basic mode screen.

### 1 | 2 | 3

are optional 'suppression' classes which you can assign to a suppressed line. The SUPPRESS option of the REDISP command then allows you to selectively suppress one or more of these classes even in redisplay mode.

A class of 1 will be set by default if no explicit class is specified with the NODISP argument.

### fldname

Can be any of the selectable output field names CPMSG, CPWNG, CPOUT, CPSMSG, VMOUT, CPAMSG, CPMSG, CPIMSG, CPSCIF and CPMSGN (but CPSMSG will be treated like CPMSG for output). If omitted, the original output line type is assumed. Specification of a certain linetype will cause the output line to be displayed with the color, highlighting and scrolling attributes of this new linetype. Even when specified together with the 'NODISP' option, the color and highlighting will be changed accordingly in redisplay mode.

### Reroute *userid*:

The line received is to be sent to another virtual machine *userid* which must be logged on to the same VM system. The line will be sent to that machine in the format defined by the next parameter.

### CPMsg

The line is to be sent as CP message.

### CPWng

The line is to be sent as CP warning (requires special privilege class!).

### CPSmsg

The line is to be sent as CP SMSG.

### CPMSGN

The line is to be sent as CP noheader message (requires special privilege class!).

### Call *fname*

Command *fname* is to be stacked for execution ('*fname*' must be the file name of an EXEC procedure or of a CMS module to be called).

### Exec *fname*

Command 'EXEC *fname*' is to be stacked for execution ('*fname*' must be the file name of an EXEC procedure).

Note that EXEC procedures can also be called by specifying the Call argument; the Exec argument has been retained mainly for compatibility with previous FCON/ESA levels.

**Osrn *fname***

The load module *fname* is to be executed using the 'OSRUN' command (*fname* must be the name of a load module in a CMS load library for which you have previously issued a 'GLOBAL LOADLIB ...' command).

**FILE**

The output line being processed is also to be written to disk as file 'FILnnnnn PROCESS A3' with a max. line length of 160. The file name is variable (the 'nnnnn'-part is numeric and will be incremented for each new file written in order to prevent duplicate names). It will also be passed to the module or EXEC procedure being executed as first and only parameter, thus allowing the program to read the output line for further detailed analysis. The FILE option will not be executed if Performance Toolkit has previously found the A-disk to be full.

Please note that the file will automatically be erased when read for the first time (file mode 'A3'), unless the same line is to be processed by multiple commands, in which case the file mode will be set to 'A1', that is, you will have to delete it yourself after use.

Rename the file to another file mode number before reading it if it should be kept for later reference. Do not, on the other hand, use the FILE option unless you really intend to read the line later on, otherwise there may soon be dozens of unused files lying around on the A-disk.

**Passargs**

Indicates that the complete message line is to be passed to the called command as arguments, up to a maximum length of 132 characters. Any characters in excess of this maximum length will be truncated without notification.

Nothing will be passed on if neither the FILE nor the Passargs arguments have been entered.

**NOTIFY**

Indicates that the information line

Processing CMD '*fname* xxxxx xxxxx'

Is to be displayed when the conditions for execution are met, and the corresponding command is stacked for later execution. The string xxxxx xxxxx consists of the arguments passed to the command during execution.

The command will be stacked and executed without an additional information line if the NOTIFY argument is omitted, and the short form of the CMS 'Ready' message will not be shown either.

Querying PROCESS definitions:

**Query**

Indicates that all PROCESS definitions for the selected line type are to be displayed. They will be shown together with a number prefix which you will need if you want to delete a specific definition. This is the default operation if no further arguments are entered.

Deleting PROCESS definitions:

**Delete *nnn***

Indicates that PROCESS definition *nnn* for the selected line type is to be removed. It will then no longer be executed.

# FCONTROL RELOAD

## Purpose

The FCONTROL RELOAD subcommand allows refreshing the in-storage copy of the data contained in the matching disk file.

## Format



## Parameters

### AUTHoriz

Indicates that the remote data retrieval authorization file FCONRMT AUTHORIZ is to be reloaded.

Changes made to an FCONRMT AUTHORIZ file after starting PERFKIT will not become active before restarting PERFKIT unless the file is explicitly reloaded.

### LINUXusr

Indicates that the Linux system definition file FCONX LINUXUSR is to be reloaded.

Changes made to an FCONX LINUXUSR file after starting PERFKIT will not become active before restarting PERFKIT unless the file is explicitly reloaded.

### SYSTEMs

Indicates that the FCONRMT SYSTEMS file is to be reloaded.

Changes made to an FCONRMT SYSTEMS file after starting PERFKIT will not become active before restarting PERFKIT unless the file is explicitly reloaded. Definitions for systems for which no connections currently exist can be added, removed or altered. Connections for systems with active connections will persist.


No attempt will be made to automatically activate the S&F interface. Use the FC MONCOLL VMCF OFF and FC MONCOLL VMCF ON commands to attempt that.

# FCONTROL RMTLINES

## Purpose

The FCONTROL RMTLINES subcommand allows setting a default value for the number of performance data lines which are to be retrieved in remote performance data retrieval mode.

## Format



## Parameters

### *nnn*

Specifies the number of lines which are to be retrieved

### ALL

Indicates that *all* lines are to be retrieved. This is also the initial setting.

### Query

Displays the current setting on the command line. This is the default if no further argument is entered.

### SCRNSize

Sets the number of lines to be returned equal to the number of lines needed to fill the current screen.

## Usage

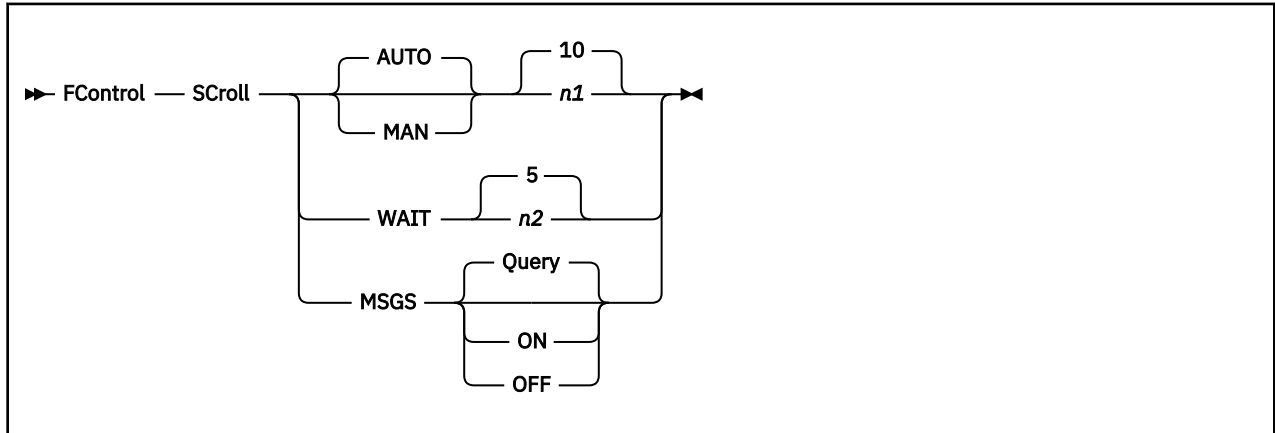
1. This command must be entered in the Performance Toolkit session of the *retrieving* machine to be effective. Every user can specify the number which is best suited for their personal needs. Entering this command in the actual performance data collection machine will have no effect on the other users which retrieve data from this machine.

# FCONTROL SCROLL

## Purpose

The FCONTROL SCROLL subcommand allows changing the current basic mode scroll setting which is displayed in the top right hand corner of the Basic Mode screen.

## Format



## Parameters

### AUTO

Indicates that scrolling is to be done automatically once the screen becomes full.

### MAN

Indicates that the screen should be left locked when it becomes full until the 'ENTER'-key is pressed.

### n1

Indicates the number of lines by which the screen should be shifted when it is scrolled (default 10 lines). The minimum is 1 and the maximum number that can be set is equal to the size of the scrollable output area - 1.

### WAIT

Means that the scroll wait time (the delay from the moment the screen has become full until it is shifted) is to be changed.

### n2

Specifies the number of seconds for the scroll wait time. The number specified must be between 0 and 60 seconds (default 5 seconds). Setting a scroll wait time of zero seconds causes immediate scrolling of most output lines, but a minimum scroll wait time of one second will still be observed when a command has been entered and the command line has reached the top of the screen. This slight inconsistency in the handling of scroll wait times has been introduced in order to allow you to freeze the screen temporarily and have a closer look at the output generated by the command, even when normal scrolling is immediate.

Only the delay for future scroll waits will be changed, a pending scroll wait will not be affected.

### MSGS

Controls the autoscrolling of the priority CP messages. The default setting – used when FCONTROL SCROLL MSGS has not been issued – is OFF, which indicates that none of the CP messages will be automatically shifted out of the screen when it becomes full.

### Query

Displays the current MSGS setting on the command line. QUERY is the default when you issue FCONTROL SCROLL MSGS with no other parameters.

**ON**

Indicates that the priority CP messages should be automatically shifted out of the screen when it becomes full. FCONTROL SCROLL MSGS ON can be used to prevent screen locking due to the screen filling up with CP messages.

**OFF**

Indicates that the autoscrolling of the priority CP messages should be disabled. If this setting is in effect and the screen fills up with the CP messages, scrolling is disabled and the message Screen locked - Delete some messages is shown in the center of the top line. When this happens, you need to delete some messages so normal operation can continue.

# FCONTROL SEARCH

## Purpose

The FCONTROL SEARCH subcommand defines the command search order to be used when a non-Performance Toolkit command is entered without a 'CP' or 'CMS' prefix.

## Format



## Parameters

### Query

Displays the current search order setting in the command line. This is the default if no further argument is entered.

### CMSCP

Indicates that the default command search order to be used is PERFKIT ==> CMS ==> CP. This search order is similar to the search order that is normally used in CMS machines.

### CPCMS

Indicates that the default command search order to be used is PERFKIT ==> CP ==> CMS. It is to be preferred, for performance reasons, in machines where CP commands are predominant. This is also the initial setting.

# FCONTROL SECUSER


## Purpose

The FCONTROL SECUSER subcommand tells Performance Toolkit that it is to act as a secondary console for another machine. Performance Toolkit will then:

- Indicate the name of the machine in the top line of the screen
- Automatically insert the command "SEND *userid* " at the beginning of the command line, so that you no longer have to enter it yourself. The "SEND *userid* " prefix can be overtyped or deleted if other commands are to be entered which should be executed directly.
- Remove all of the "*userid* : " prefixes from SCIF output lines so that the output format is the same as on the original machine.

These functions make sense only if the machine where you enter the SECUSER command is more or less dedicated to operating the other virtual machine.

## Format



## Parameters

### *userid*

Is the user identification of the other machine for which we act as a secondary console.

### *prefix*

Is an additional command prefix that is to be inserted into the command line.

## Usage

1. When working with RSCS Version 2 or 3, you might specify RSCS for the prefix. This would cause the SEND *userid* RSCS command string to be inserted and thus eliminate the need to manually type in the prefix whenever an RSCS command has to be entered. OFF disables the function.

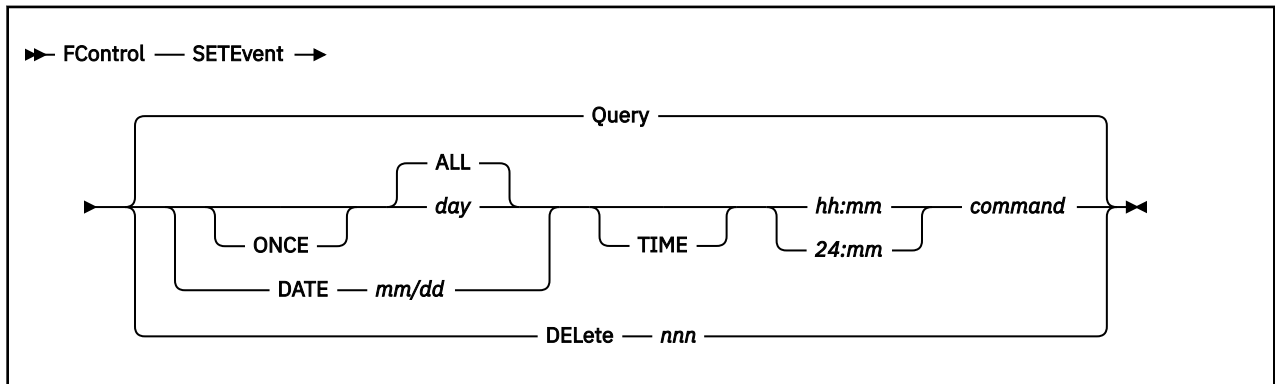


# FCONTROL SETEVENT

## Purpose

The FCONTROL SETEVENT subcommand allows defining the weekday(s) and the time at which specific events are to take place. It is intended to help in automating certain tasks related to the PERFSVM machine, but can also be used for general automation tasks on systems where no other tool is available for the automatic execution of timer driven events.

## Format



## Parameters

### DATE *mm/dd*

Defines a specific day on which the event is to take place. The date format *mm/dd* must be exactly as shown. (i.e. two digits each for the month and the day, separated by a slash.) A DATE specification implies a once only execution. (i.e. the event will be removed from the table after execution.)

### ONCE

Indicates that the event is to take place one time only.

### *day*

Defines the weekday(s) on which the event is to take place. A combination of any number of the following strings can be entered, separated by blanks. Their meaning is:

#### ALL

The event is to take place on all weekdays

#### M-F

The event is to take place on the workdays Monday to Friday only

#### MON

The event is to take place on Mondays

#### TUE

The event is to take place on Tuesdays

#### WED

The event is to take place on Wednesdays

#### THU

The event is to take place on Thursdays

#### FRI

The event is to take place on Fridays

#### SAT

The event is to take place on Saturdays

#### SUN

The event is to take place on Sundays

Specifying the 'days' M-F SAT, for instance, will let the event take place on all days except Sunday.

### **DELeTe *nnn***

Removes the event with number *nnn* from the chain of events.

Use the 'Query' argument of the command for listing the currently defined events and their numbers.

### **Query**

Indicates that all currently defined events are to be displayed. This is the default.

### **TIME *hh:mm***

Is the time when the event is to take place. The TIME argument is optional. The format must be exactly as shown, i.e. two numbers each for hours and minutes, separated by a colon.

If the hour *hh* is specified as 24 the corresponding event will be scheduled at *mm* minutes past each full hour of the day.

### **command**

Is the command which is to be executed at the selected moment. Note that the command entered will be stacked for execution like any other command entered from the command line, i.e. do not forget to add 'CMS', 'EXEC' or 'CP' prefixes as required.

## **Usage**

1. You will have noted that the highest resolution is to minutes only, and that seconds cannot be entered. This is due to the implementation of the facility which does not actually set additional timer events, but checks during each of the standard internal one-minute cycles of the program whether the time has been reached to execute another one of the set events. In other words, the command may be executed with a delay of up to 60 seconds.

## **Examples**

Entering the command

```
FC SETEVENT M-F SAT 09:00 CP MSG * This is not a Sunday
```

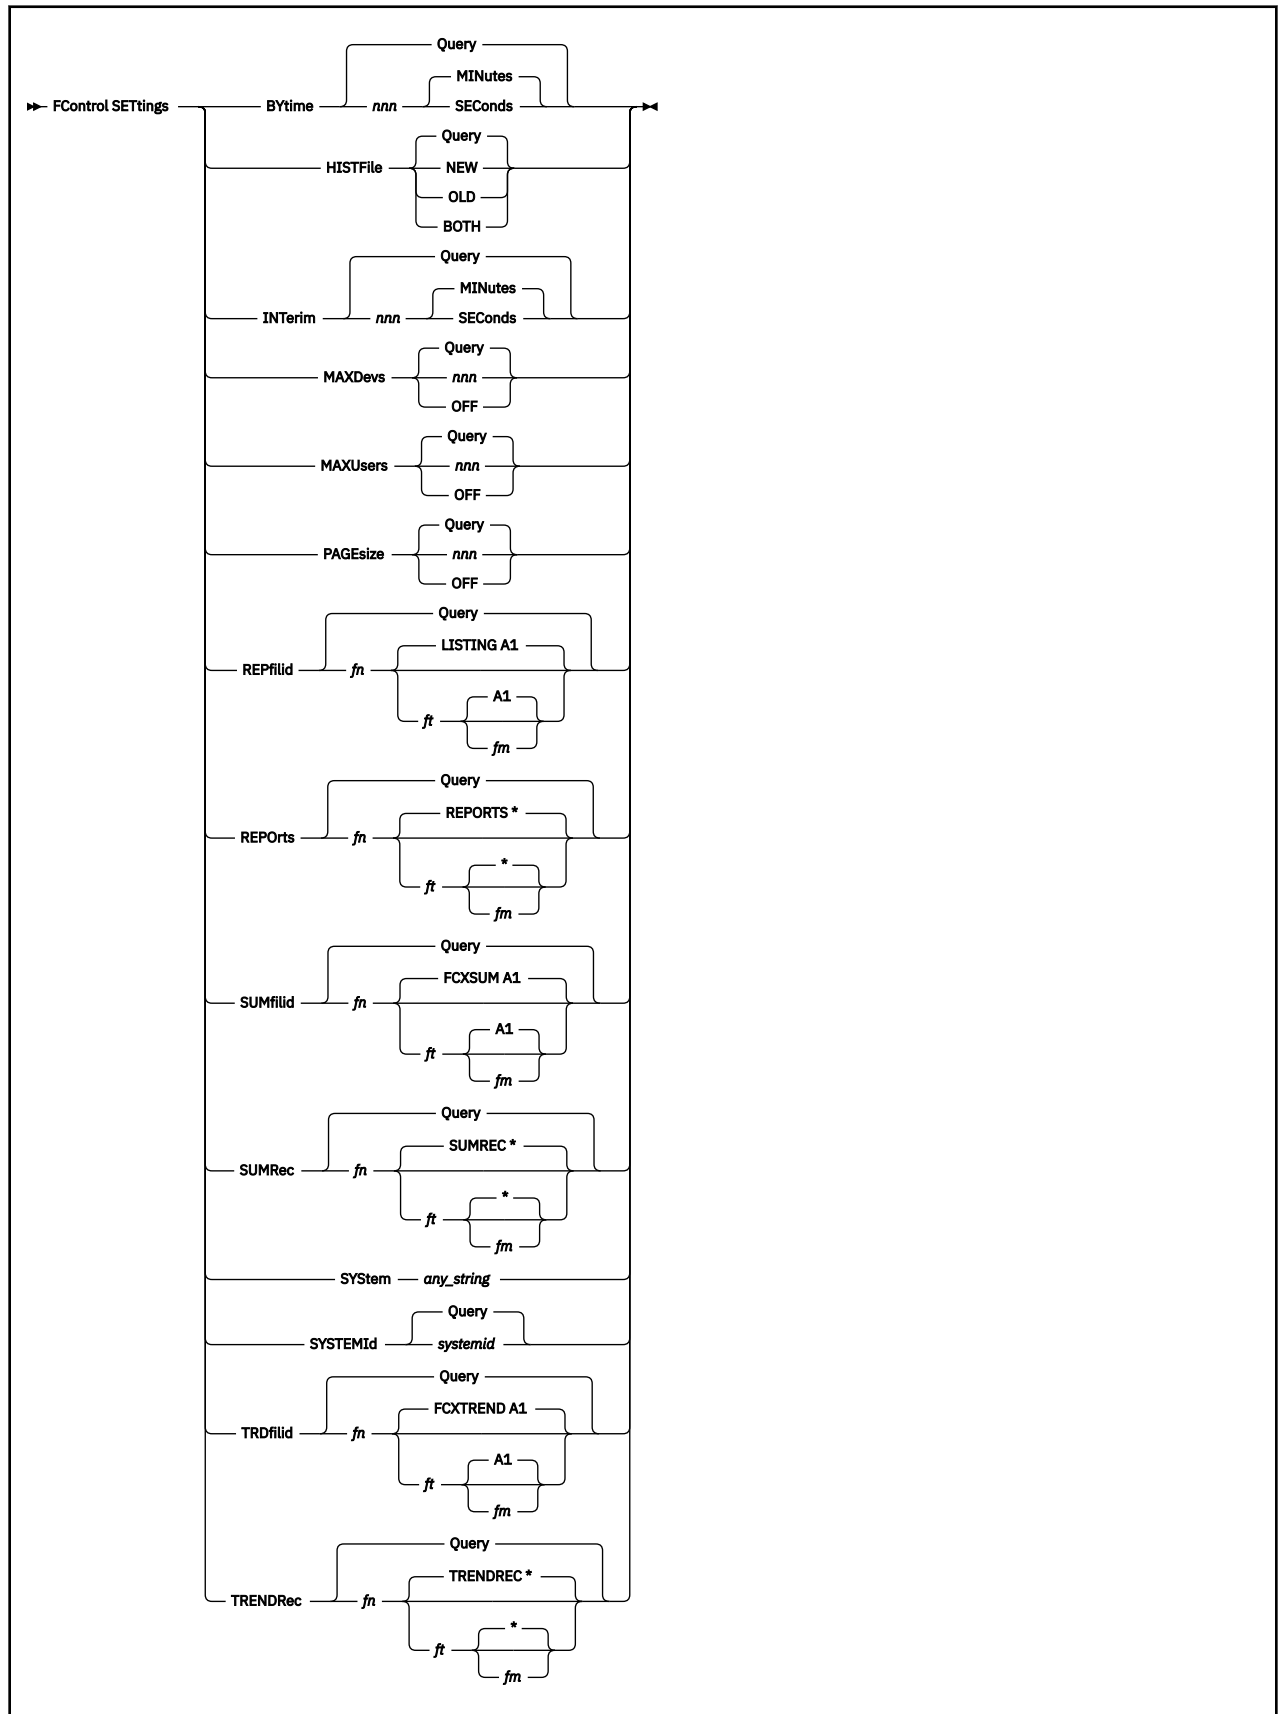
will let the message 'THIS IS NOT A SUNDAY' be displayed at 9:00 AM on all days except on Sundays.

## **FCONTROL SETTINGS**

### **Purpose**

The FCONTROL SETTINGS subcommand allows setting some general defaults.

## Format



## Parameters

### BYtime

Defines the time interval to be reported by each line of a 'by time' log. A new line will be created for each monitor sample interval if no BYTIME value has been set.

#### *nnn*

Specifies the length of time interval in seconds or in minutes. The default is in minutes.

### Query

Inserts and displays the current BYTIME setting on the command line. QUERY is the default.

### HISTFile

Indicates that the performance history log file(s) to be created will be defined, where

#### **NEW**

Indicates that *only* the new extended history log files are to be created (file types HISTLOG and HISTSUM)

#### **OLD**

Indicates that *only* the old extended history log files are to be created (file types PERFLOG and PERFHIST)

#### **BOTH**

Indicates that both the old and new history files are to be created.

BOTH has also been chosen as the initial default setting to ease migration, but you should be aware that *all* of the data available in the old PERFLOG/PERFHIST files is also available in the new files.

### Query

Inserts and displays the current HISTFILE setting on the command line. QUERY is the default.

### INTerim

Defines the length of fixed 'interim' periods, used for calculating 'interim' averages for many displays after entering the 'INTERIM' subcommand.

#### *nnn*

Specifies the length of the time interval in seconds or in minutes. Minutes is the default.

Note that time intervals smaller than the monitor sample interval cannot be reported. The 'interim' intervals are synchronized to start with the beginning of the current 'report' period, used for displaying 'average' data.

### MAXDevs

Indicates that the maximum number of I/O devices which are to be included on printed reports is to be set or queried, where

#### *nnn*

Is the maximum number of I/O devices to be included

#### **OFF**

Indicates that the number of included I/O devices should not be limited

### Query

Inserts and displays the current MAXDEVS value on the command line. QUERY is the default.

### MAXUsers

Indicates that the maximum number of users which are to be included on printed reports is to be set or queried, where

#### *nnn*

Is the maximum number of users to be included

#### **OFF**

Indicates that the number of included users should not be limited

### Query

Inserts and displays the current MAXUSERS value on the command line. QUERY is the default.

**PAGEsize**

Indicates that the default page size is to be set or queried

***nnn***

Is the maximum number of lines to be printed per page

**OFF**

Indicates that no page size limit is to be used, i.e. performance reports are to be printed with just one initial header

**Query**

Inserts and displays the current PAGESIZE value on the command line. QUERY is the default.

**REPfilid**

Indicates that the report file ID for 'printing' performance reports to disk is to be set or queried, where

***fn***

Is the disk file name

***ft***

Is the disk file type. A file type of LISTING will be used if no explicit file type is entered.

***fm***

Is the disk file mode. A file mode of A1 will be used if no explicit file mode is entered.

**Query**

Inserts and displays the current REPFILID setting on the command line. QUERY is the default.

**REPORTs**

Indicates that the file identification for the 'reports' control file referred to by the 'PRINT ALL' command is to be set or queried, where:

***fn***

Is the disk file name.

***ft***

Is the disk file type. A file type of REPORTS will be assumed if no explicit file type is entered.

***fm***

Is the disk file mode. A file mode of \* will be used if no explicit file type is entered.

**Query**

Inserts and displays the current REPORTS setting on the command line. This is the default.

File FCONX REPORTS will be looked for if no alternate print control file ID has been set.

**SUMfilid**

Indicates that the (initial) file identification for the 'summary' output file is to be set or queried, where:

***fn***

Is the disk file name.

***ft***

Is the disk file type. A file type of FCXSUM will be used if no explicit file type is entered.

***fm***

Is the disk file mode. A file mode of A1 will be used if no explicit file mode is entered.

**Query**

Inserts and displays the current SUMFILID setting on the command line. This is the default.

An initial file ID of *systemid* FCXSUM A1 will be used if no alternate file ID has been set, where the system identification of the VM system is inserted for *systemid*. Be aware that the file type may be automatically changed if an existing file with the same file ID already exists. Consecutive numbers from '01' to '99' will be either appended to the file type if it is up to 6 characters long, or to the first 6 characters of the file type if it is longer, to create unique file *identifiers*.

**SUMRec**

Indicates that the file identification for the summary records control file is to be set or queried, where:

***fn***

Is the disk file name.

***ft***

Is the disk file type. A file type of SUMREC will be used if no explicit file type is entered.

***fm***

is the disk file mode. A file mode of \* will be used if no explicit file mode is entered.

**Query**

Inserts and displays the current SUMREC setting on the command line. This is the default.

File FCONX SUMREC \* will be looked for if no alternate summary records control file ID has been set.

**SYStem**

Allows specifying a character string which describes the VM system for which reports are to be generated. The string entered will be included in the header part of all printed reports.

The maximum string length accepted is 64; the string will be truncated to this length if it is longer.

**SYSTEMID**

allows setting or querying a short system identifier which will be included on the right hand side of the print header of all printed performance reports, where

***systemid***

Is a short (maximum eight characters long) character string which is to be inserted in the print header to describe the VM system (usually the system-ID or the RSCS node-ID)

**Query**

Inserts and displays the current SYSTEMID setting on the command line. QUERY is the default.

**TRDFilid**

Indicates that the file identified for the 'trend' output file is to be set or queried, where:

***fn***

Is the disk file name.

***ft***

Is the disk file type. A file type of FCXTREND will be used if no explicit file type is entered.

***fm***

Is the disk file mode. A file mode of A1 will be used if no explicit file mode is entered.

**Query**

Inserts and displays the current TRDFILID setting on the command line. This is the default.

File *systemid* FCXTREND A1 will be created if no alternate file ID has been set, where the system identification of the VM system is inserted for *systemid*.

**TRENDRec**

Indicates that the file identified for the trend records control file is to be set or queried, where:

***fn***

Is the disk file name.

***ft***

Is the disk file type. A file type of TRENDREC will be used if no explicit file type is entered.

***fm***

Is the disk file mode. A file mode of \* will be used if no explicit file mode is entered.

**Query**

Inserts and displays the current TRENDREC setting on the command line. This is the default.

File FCONX TRENDREC \* will be looked for if no alternate trend records control file ID has been set.

# FCONTROL TIMESTAMP

## Purpose

The FCONTROL TIMESTAMP subcommand tells Performance Toolkit whether time stamps should be included in front of each line on the basic mode and redisplay screens.

## Format



## Parameters

### ON

Indicates that time stamps are to be inserted in front of *new* lines from now on.

### OFF

Indicates that no time stamps are to be inserted in front of *new* lines.

## Usage

1. Note that the time stamp setting affects only new lines; any previously entered or received lines will not be changed.
2. The user acting as a system operator (who receives all the time-stamped system messages) will usually have all command lines time-stamped by Performance Toolkit even without entering the FC TIMESTAMP ON command. The command can be used there to prevent time stamps for command lines; all other lines will not be affected.
3. The current TIMESTAMP setting affects the interpretation of *all* command lines, when a command is re-executed. This can lead to unexpected results if a command was originally entered with TIMESTAMP ON. Re-execution of the command with TIMESTAMP set OFF will cause the time-stamp prefix to be interpreted as part of the command line.




# FCONTROL UCLASS

## Purpose

The FCONTROL UCLASS subcommand allows you to define user classes for which separate averages are to be calculated and displayed on user based screens.

## Format



## Parameters

### *userid*

Specifies a user identification which is to be counted as part of the class *classname*.

The string *userid* can be

- The identification of a specific virtual machine, or
- A generic identification, containing wildcard characters '\*' and/or '%' which are to be used in the same way as for the CMS 'LISTFILE' command, to designate a set of users with similar userids.

### *classname*

Defines the user class name which is to be inserted in the 'Userid' column on the line with the group's averages. A maximum of eight characters are accepted, including blanks.

### **Query**

lists the currently active user class definitions. This is the default.

## Usage

1. UCLASS specifications become effective only for users logging on *after* they have been entered.
2. Multiple statements with the same user class name *classname* are allowed.
3. A specific virtual machine can only be part of one user class: The last matching class will be used, where the scanning sequence is identical to the sequence in which the FC UCLASS commands were entered.

This means that FC UCLASS statements should be entered in decreasing order of generality, to let the most specific selection become effective.

# FCONTROL UPDTCMS

## Purpose

The FCONTROL UPDTCMS subcommand defines the manner in which the screen is to be updated while a CMS command is being executed under Performance Toolkit.

## Format



## Parameters

### Query

Displays the current setting on the command line (default).

### Delayed

Indicates that all output lines are to be stacked and that the screen should be updated only when the CMS command has ended, or if we enter a CMS console read state. This is the initial display mode, and it also gives best performance.

### Immed

CP lines will still be stacked first, but when CMS output lines are received they will immediately be inserted into the screen, together with any stacked CP lines.

Be aware that 'immediately' can mean just that: lines may be shown immediately regardless of the current screen contents, i.e. even if it is in use by another full screen application (no 'MORE...' status).

## Usage

1. If the full screen application uses the CMS CONSOLE macro for handling its console I/O, no screen update will take place before this application has terminated.

# FCONTROL UPDTSCRN

## Purpose

The FCONTROL UPDTSCRN subcommand specifies whether console output lines are to be unstacked and inserted into the basic mode screen in any mode, or only while the screen is actually in basic mode, thus allowing you to view all output before it may be scrolled and disappear.

## Format



## Parameters

### ANYmode

Indicates that output lines are to be unstacked and inserted into the basic mode display regardless of the current operating mode of Performance Toolkit. ANYMODE is also the initial setting.

### BASmode

Indicates that output lines should be unstacked and inserted into the basic mode display only while Performance Toolkit is actually used in basic mode too. All output will remain stacked while the program is used in any non-basic mode.

### Query

Indicates that the current setting is to be queried, and the result to be inserted into the command line where you can modify it. This is the default operation if no further argument is entered.

## Usage

1. Using ANYMODE mode has the advantage that new console output does not have to remain stacked (and use a lot of virtual storage) if the screen is left in any non-basic mode (re-display mode, for example) for a long time, thus preventing potential problems when storage is exhausted.

The disadvantage of working with the ANYMODE setting is that, after a long redisplay or performance monitoring session, the user will not see output lines which have been received and already scrolled before switching the display back to basic mode (although it will always be possible, of course, to view such output in re-display mode).

2. Using the BASMODE setting will make sure that the user can see *all* output lines before they are scrolled, but it can lead to loss of output data or even an abend when the machine's virtual storage is completely filled up and new lines can no longer be stacked.


BASMODE corresponds to the standard operating mode of FCON/ESA prior to Version 2.3 where the screen update mode could not yet be selected.

# FCONTROL USERBOTL

## Purpose

The FCONTROL USERBOTL subcommand defines the contents of the bottom line for either the 'basic' mode display, or for the initial system overview display in remote monitoring mode (shown after FCONRMT command). When specified, the data entered will replace the default bottom line information on the selected display.

## Format



## Parameters

### BASIC

Indicates that the command refers to the bottom line of the basic mode display

### FCONRMT

Indicates that the command refers to the bottom line of the initial system overview display of remote monitoring mode. The other remote monitoring displays (performance data, exception log and history data) are not affected.

### 'user data'

Is the character string that is to replace the default bottom line. The apostrophes are not required, they will be assumed to be part of the data if entered. The maximum length allowed is 77 characters for the basic mode display, and 78 characters for the remote monitoring display.

If user data is not entered, the bottom line will be reset to show the original default information again.

Note that the bottom line for the basic mode display is automatically split in two parts:

- The first part, where the PF-key assignments are originally shown, can be used as an extension of the command line for entering long commands
- A second part which is defined as 'protected' and which cannot be overwritten.


The two parts are separated by a 3270 field definition attribute which will appear as a blank whose position you cannot change.

## FCONTROL USERHDR

### Purpose

The FCONTROL USERHDR subcommand defines a special user header line that is to be used instead of Performance Toolkit's default 'idle' header line. When specified, the data entered will be shown in the center of the top line as long as no special conditions (e.g. scrolling, CMS command active) arise that will overlay it temporarily.

### Format



### Parameters

#### 'user data'

Is the character string that is to replace the default header line. The apostrophes are not required, they will be assumed to be part of the data if entered. The maximum length allowed is 55 characters.

### Usage

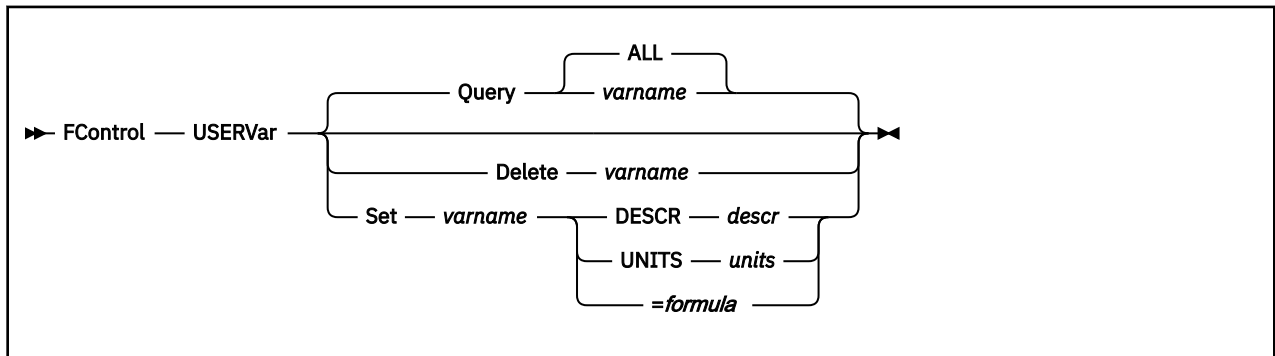
1. If no user data are entered, the header line will be reset to show Performance Toolkit's default information again.

# FCONTROL USERVAR

## Purpose

The FCONTROL USERVAR command allows the definition of 'user performance variables' for use with the PLOTxxx and GRAPHxxx commands, and they will also be included in the variable correlation coefficients display. These user variables will be calculated according to a formula you have supplied when defining them, based on the original set of performance variables available in the selected source.

## Format



## Parameters

### Delete

Indicates that a user variable is to be deleted, where *varname* is the name of the user variable

### Query

Indicates that the current definitions for one or several user variables are to be displayed, where

#### ALL

Specifies that all currently defined user variables are to be displayed

#### *varname*

Specifies the name of a single user variable which is to be displayed

Query is the default operation if no further argument is entered.

### Set

Indicates that a new user variable is to be defined, or the definitions of a previously defined user variable altered

#### *varname*

Is the name of the user performance variable that is to be defined.

#### = *formula*

Specifies the algorithm that Performance Toolkit is to use in calculating the new variable, where *formula* can be a combination of the following:

#### Performance variable names

The basis for further calculations. All variables available with the history log displays REDISP and REDHIST are eligible for use as input in a formula.

Note that such variables must be entered with leading and trailing single quotes if they contain characters which also have a mathematical meaning such as +, -, \*, /, or (). Performance Toolkit will treat these characters as mathematical operators otherwise and will, consequently, split the variable name at such points, which could cause the resulting bits not to be recognized as variables or to be treated as an entirely different variable.

#### Literals

Fixed numbers.

## Operators

The operations to be performed on the basic performance variables and literals. Valid operators are

```
+ - * /
```

## Parentheses

Specify the sequence in which the mathematical operations are to be performed.

The requested operations will be performed in the usual sequence as defined by the rules of algebra: multiplications and divisions before additions and subtractions, and from left to right within the same level of priority, unless parentheses force a different sequence.

## DESCR

Defines a short description *descr* of the variable, with a maximum length of 16 characters.

This description will be included in the header line of plots and GDDM graphics when specified; a question mark '?' will be shown instead if the description is missing.

## UNITS

Specifies the units *units* for the variable, with a maximum length of 4 characters. This specification, as the variable description above, will be used when creating graphics output. It is for information only; it does not affect the calculation of the variable at all.

## Usage

1. The specified formula will be checked for correct format. When errors are detected, you will receive error messages and the variable definition will not be accepted.

Even for a mathematically-correct formula, you might not always see the expected results when you specify the newly-defined variable name with a PLOTxxx or GRAPHxxx command, due to internal restrictions:

- All values are converted to short floating point format for the actual calculation. The result, however, is kept as a binary fixed integer, with two decimals. This means that results must remain within the range 0.01 to 20,000,000 or they are set to the respective limit of 0 or 20,000,000).

Make sure to keep your results within this range, by adding one final multiplication or division to your formula if necessary.

- Note that the basic performance variables may sometimes have a value of zero. If your formula leads to a division by zero (dividing by a performance variable whose value can be zero), the division is not executed, leaving the original value intact. This avoids a division exception, but may cause incorrect results to be shown.

You might be able to avoid divisions by zero by adapting your formula accordingly.

2. When one or more of the previously-listed conditions is found, this message is displayed:

```
FCXxxx467I nnn zero-divide(s) avoided, mmm times maximum  
exceeded
```

This always means that some values had to be changed to go on working, and it is your responsibility to evaluate the impact on the resulting graphics.

Be aware that the set of available variables is smaller when the input comes from the redisplay buffer or from old PERFLOG or PERFHIST files. User variables based on any variables that are only available with the new HISTLOG and HISTSUM file format are not valid in this case.

## Examples

Let us assume you are interested in the amount of CPU used per total number of transactions (trivial + non-trivial), which could give you some indication on workload changes. The user variable CPU/TR with the following definition would be the answer:

```
FC USERVAR SET 'CPU/TR' = CPU / (100*('TR/S'+ 'NT/S'))
```

The result should be large enough (i.e. more than 0.01 seconds) to avoid loss of data due to the minimum value mentioned in the 'Usage Notes' above. However, the value may be less than 0.1 second, resulting in very coarse granularity since only multiples of 0.01 sec. can be shown. Displaying the number of milliseconds per transaction instead of seconds would allow us to show more detail: just multiply the result by 1000. This would change the formula to:

```
FC USERVAR SET 'CPU/TR' = (CPU / (100*('TR/S'+ 'NT/S'))) * 1000
```

or

```
FC USERVAR SET 'CPU/TR' = CPU*10 / ('TR/S'+ 'NT/S')
```

and should give satisfactory values even on fast processors.



# FCONTROL USRLIMIT

## Purpose

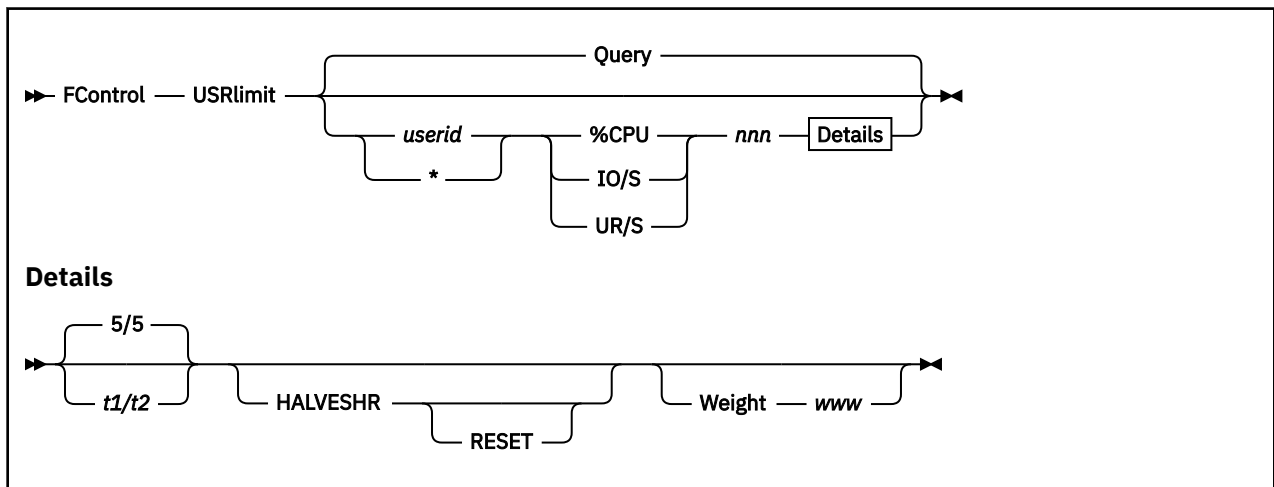
The FCONTROL USRLIMIT command allows setting thresholds for user resource consumption, and you can let the offending users' relative SHARE be automatically reduced in order to reduce their impact on general system performance. You can set:

- Individual thresholds for specific users' resource consumption, and/or
- General thresholds for all users

and specify whether or not their relative SHARE is to be changed.

Alert messages will be generated when a threshold has been exceeded for more than a predetermined time interval, and the relevant load fields on the general user display will be shown with special highlighting (red, reverse video) if the display terminal supports extended color and highlighting. The command is valid only when permanent performance data collection has been activated (command FC MONCOLL ON).

## Format



## Parameters

### *userid*

is the user identification of the virtual machine for which a threshold is to be set.

'Wildcard' characters '\*' and '%' can be used for creating generic entries which could apply to several machines; they have the same meaning as for the CMS 'LISTFILE' command. Note that because a wildcard entry could cause an override to a userid in a previous USRLIMIT entry, it is recommended that the wildcard entries be the first entries in the USRLIMIT sequence of commands, following by any non-wildcard entries.

Where an explicit entry with individual thresholds exists, *only* the individual thresholds will be applied – i.e. if you have set an individual %CPU threshold for user MAINT, but no IO/S or UR/S thresholds, then user MAINT's virtual I/O rate and UR I/O rate will not be monitored even if a general \*-entry for virtual I/O rates and UR I/O has been set too.

### **%CPU**

Indicates that a threshold for the total CPU consumption is to be set

### **IO/S**

Indicates that a threshold for the virtual I/O rate is to be set

## **UR/S**

Indicates that a threshold for the total UR I/O rate is to be set, i.e. for the total I/O rate to virtual readers, card punches and printers.

## **nnn**

Is the threshold which is to be set for the selected resource (without any decimals)

## **t1**

Is the initial period, in minutes, after which a first alert message is to be generated

## **t2**

Is the period, also in minutes, after which the alert message is to be repeated if the threshold value for the resource is still exceeded.

The actual intervals used for message generation depend also on the length of the monitor sample interval. See 'Usage Notes' below for more information.

## **HALVESH**

Specifies that a user's relative SHARE is to be set to half the previous value whenever the conditions for message generation are also met, i.e. the same *t1/t2* intervals apply as for message generation. This process is repeated, if necessary, until a relative SHARE value of 1 is set.

The argument will be effective only for users whose original relative SHARE (when Performance Toolkit was started, or when they logged on) was equal to the default of 100. Users with another relative SHARE, which have an absolute SHARE set, or which own a dedicated processor, are exempt from automatic SHARE reduction. The assumption is that the special SHARE value was assigned for a good reason, and that it should not be changed by this.

Note that the machine where Performance Toolkit is running must have been authorized to issue the CP command 'SET SHARE ...' (requires privilege class A) in order to change users' relative SHARES.

## **RESET**

Specifies that a user's relative SHARE is to be reset to the original default value of 100:

- When the machine is found to be inactive (not in-queue), *and*
- When it has not exceeded any of the applicable user thresholds in the last interval.

The machine's relative SHARE value will never be increased again if the RESET argument has not been specified.

Note that specifying RESET for one of the monitored resources (e.g. CPU) will let the users' SHARE be restored to 100 also if it had been originally reduced because the threshold for another resource (e.g. I/O/S) had been exceeded.

## **Weight www**

Sets the 'weight' *www* of the exception when one of the set thresholds is exceeded, where *www* can be any number in the range 0 to 255. If no argument is specified, a value of 10 will be set as a default for the first FC USRLimit command.

The WEIGHT specification is relevant only when a machine's exception data are sent to a central monitor machine (activated via REMSEND argument of the FC MONCOLL command, see "Remote Performance Monitoring Facility" in the *z/VM: Performance Toolkit Guide* for a description of the central monitoring function). All applicable weights will then be added, and the result divided by 10 to obtain the exception severity code which will be used to set the color of the load bar in the central monitoring machine.

Only one WEIGHT value can be set for all user thresholds: the last one specified will be used for **all** of them. The 'weight' number for an exceeded user threshold will be added once per exceeded threshold, i.e. it may be added several times if multiple thresholds have been exceeded.

## **Query**

Lists all of the currently active USRLIMIT settings. Query is the default operation if no further argument is entered

## Usage

1. For performance reasons the applicable thresholds are copied to each virtual machine's internal data buffer only once, either during PERFKIT initialization, or when the virtual machine logs on. All USRLIMIT specifications must, therefore, be included in the initialization file FCONX \$PROFILE in order to become effective for *all* users.
2. Message generation and HALVESHHR setting are coupled to the data collection interval (the monitor sample interval), i.e. a message will be generated and the user's relative share halved during the first data collection cycle after the initial or repeat intervals have elapsed. The actual intervals will thus differ from the values set with *t1* and *t2* above, and the effect will be more marked if the monitor sample interval is considerably longer than the *t1* and *t2* intervals.

## Examples

After setting a threshold of 30% for the total CPU consumption of all users on the system with command:

```
FC USRLIMIT * %CPU 30 8/10
```

the following message could be generated when USER1 has exceeded this limit for more than eight minutes:

```
User USER1 %CPU 34.8 exceeded threshold 30.0 for 8 min.
```

The message will be repeated after another 10 minutes unless the CPU consumption of user USER1 drops below the threshold value in the mean time.

No further action is taken. By specifying the command with the 'HALVESHHR' and 'RESET' arguments:

```
FC USRLIMIT * %CPU 30 8/10 HALVESHHR RESET
```

you let the users' SHARE be divided by two whenever they have exceeded a threshold for the specified interval. It will be restored to the original value of 100 once they are found to have become idle.

You can also intercept the message using the 'FC PROCESS ...' command, and use it for triggering further action (e.g. by calling an EXEC procedure which could do some additional analysis, send a warning message to the user, or even FORCE it off the system).

See also the "User Monitoring" section in the [z/VM: Performance Toolkit Guide](#) for more information on user threshold monitoring, and for hints on the use of the FC USRLIMIT command.

## Basic Mode Subcommands

---

All references to PF-key numbers pertain to the initial settings. Different PF-keys might have to be used if you reassigned them.

# CLEAR

## Purpose

The CLEAR command clears the screen. The operation of the CLEAR command for the clearing of CP messages and Performance Toolkit action messages can be controlled by the 'FC MSGCLEAR command'; *all* lines will be cleared by default.

## Format

►► CLear ◄◄

## Parameters

No parameters are allowed or tested for.

## Usage

1. The CLEAR function can also be performed by pressing the 'CLEAR'-key or the PA2 key. The PA2-key will clear only the scroll area lines and leave the command and bottom lines unchanged, while both the 'CLEAR' command and the CLEAR-key clear and reset the complete screen.

# CMS

## Purpose

The CMS prefix indicates that the standard command search order is to be bypassed and the following command executed as a CMS command. It causes the Remote CMS Command Execution Screen (FCX123) to be displayed.

## Format

➤ CMS — '*CMS command*' ➤

## Parameters

### 'CMS command'

Is any valid CMS command or EXEC name (do not enter the apostrophes).

## Usage

1. *Do not attempt to activate any full screen application.* Doing so will place the server machine in a full screen mode session which you cannot end by sending remote commands.  
  
You will have to logon to the server machine in order to exit from the full screen application if you have inadvertently started one.
2. See also “[BASMODE](#)” on [page 137](#) for a more user friendly method of executing commands remotely (for APPC/VM sessions only).

## Results

See “[FCX123, Remote CMS Command Execution Screen – CMS](#)” on [page 470](#).

# CP

## Purpose

The CP prefix indicates that the standard command search order of Performance Toolkit is to be bypassed and the following command executed as a CP command. It causes the Remote CP Command Execution Screen (FCX122) to be displayed.

Only the '#CP' form of the command is accepted while a CMS console read is pending ('ENTER:' in front of the command line), but no distinction between the two forms is made in the normal case when the command line is preceded by the standard arrow.

## Format

➤ #CP — '*CP command*' ➤

## Parameters

### '*CP command*'

Is a valid CP command (do not enter the apostrophes).

## Usage

1. CP MODE will be entered if no command follows, and the CP interface of Performance Toolkit will then be disabled so that all command input and output is shown on the console as usual. Commands entered in this mode will not be contained in the log of Performance Toolkit.
2. Enter 'BEGIN' to switch back to Performance Toolkit mode (required only if 'CP' has been entered without a command following it).
3. Please note that even the '#CP' form is not intercepted by CP but has to be read and interpreted by Performance Toolkit first. The command can, therefore, *not* be accepted while CMS programs are running which have disabled I/O interrupts from the console.

## Results

See [“FCX122, Remote CP Command Execution Screen – CP”](#) on page 469.

# DELETE

## Purpose

The DELETE command resets the message or action message lines specified so that they can be scrolled again like any other line on the screen.

## Format



## Parameters

### ***nn {- nn}***

Defines the number(s) of the line(s) to be reset. When a range of numbers has been specified, all lines whose numbers are within the defined range will be reset. The first line number of a range *must* be one of the line numbers on the screen, while the second number can be any other number.

## Usage

1. After resetting, the line's color and highlighting attributes will be changed to the ones defined for CP output (CPOUT) and the '\*' or '+' prefix removed, and the line will be scrollable.

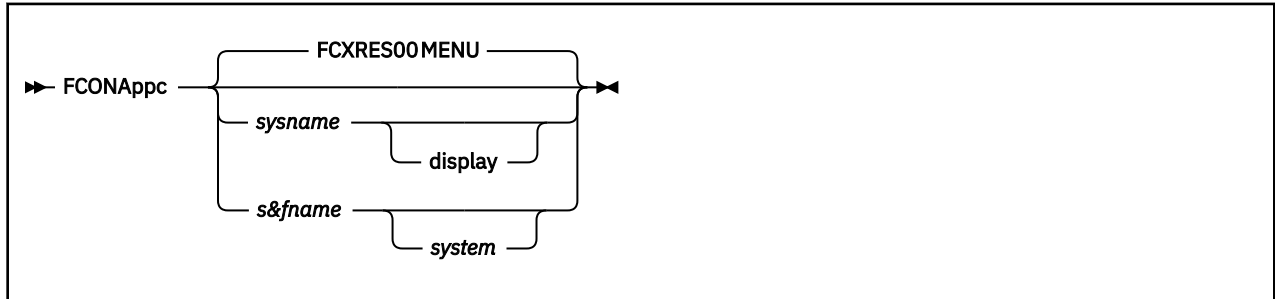


# FCONAPPC

## Purpose

The FCONAPPC command switches Performance Toolkit into APPC/VM remote performance monitoring mode.

## Format



## Parameters

### *sysname*

Is either of the following:

- A global APPC/VM resource name you want to connect to (for direct connections to a particular server), or
- A nickname in a communications directory which describes the resource you want to connect to.

A default *sysname* of FCXRES00 will be assumed when the command is entered without an argument.

### *display*

Is the selection command for a performance monitoring display which is to be retrieved. The initial performance monitoring menu (MENU display) will be retrieved, by default, if no other selection is made.

### *s&fname*

Is either of the following:

- An APPC/VM system resource name for connecting to the local PERFSVM machine's store and forward interface, or
- A nickname in a communications directory which describes the S&F system resource you want to connect to.

### *system*

Is the optional name of the system you want to connect to. The initial performance data menu for the selected system will be shown if a valid system name has been specified; otherwise a system overview menu with all accessible systems will be presented.

## Usage


1. Although it is recommended to use the system-ID or the RSCS node-ID of remote systems as nicknames, it may not always be possible to do so. Ask your systems administrator or check your xCOMDIR NAMES files to verify the *sysname* names to be used for connecting to the PERFSVM machine on a specific VM system, or to the S&F interface on the local system.
2. The FCONAPPC command can also be entered from the FCONRMT system overview display.

# FCONRMT

## Purpose

The FCONRMT command switches PERFKIT into remote performance monitoring mode, and causes the System Load Overview Screen (FCX198) to be displayed.

## Format



## Parameters

### *userid*

Is the user identification of a central data collector machine which has been set up to continuously collect performance data from other VM systems.

### *devno*

Is the virtual device number of the central data collector machine's A-disk. Device number 191 will be assumed if no other number is entered. The program will attempt to link to this disk as virtual device number 1F1.

### *fm*

Is the file mode with which the central collector machine's A-disk is to be accessed. Default file mode is 'B'.

## Usage

1. The FCONRMT command will always activate remote performance monitoring mode. Depending on the machine where the command is entered, and on the command format used, different results will be obtained:
  - If the FCONRMT command is entered without further arguments in a machine which has been set up to act as a central data collector machine for other VM systems, you will see an initial system overview display which shows the status of all remote systems, and this display will allow selection of other more detailed performance information for each of the monitored systems.
  - If the FCONRMT command is entered with the *userid* and device number of a central data collector machine, you will see a copy of the system overview display from the central collector machine, and this display will automatically be updated in one-minute intervals. You can also display all the additional more detailed data which has been collected in the central machine, and retrieve additional data from the remote systems.

Note, however, that this display is based on the reading of files from another machine which may update these files even while they are being read. The disk will be re-accessed before reading a file, but errors are still possible. Just retry the command if an error has been found.

  - If the FCONRMT command is entered without further arguments in a machine which does *not* act as central data collector machine, you will just enter remote performance monitoring mode for performance data retrieval, but no automatic updating will take place.
2. See the "Remote Performance Monitoring Facility" section in the [z/VM: Performance Toolkit Guide](#) for a detailed description of remote performance monitoring.

## Results


See the "Using Remote Performance Monitoring" section in the [\*z/VM: Performance Toolkit Guide\*](#) for more information on the System Load Overview Screen (FCX198).

# MODLEVEL

## Purpose

The MODLEVEL subcommand allows you to display service level information for one or all parts in a Performance Toolkit module.

## Format



## Parameters

### *module\_fn*

The file name of the module for which service level information is desired.

### *module\_ft*

The file type of the module for which service level information is desired. The default is MODULE.

### *module\_fm*

The file mode of the module for which service level information is desired. The default is \*.

### (SELECT *part\_name*

Use this parameter to display the service level information for one specific part in the module. If not specified, all parts in the module will be displayed.

## Usage

- When you issue this subcommand to display the service level information for all parts in a Performance Toolkit module, you will see a display as per the following example:

```
modlevel perfkit module a  
SLVL FCXINI    BASE  
SLVL FCXBAS    BASE  
SLVL FCXBAN    BASE  
SLVL FCXBAF    BASE  
.  
.  
.  
SLVL FCXOMV    BASE  
SLVL FCXOMX    VM65432
```

In this example, all of the parts in PERFKIT MODULE have the original release's service level (BASE), except FCXOMX, which was updated by APAR VM65432.

- When you issue the FC FCXLEVEL subcommand (“FCCONTROL FCXLEVEL” on page 32) from within a Performance Toolkit display screen, the Performance Toolkit program version, release, modification level, and service level will be displayed. In addition, the service level information for each part in the module used for the current invocation of Performance Toolkit (that is, as if the MODLEVEL subcommand was issued for this module) will be written to the log.

# HISTDATA

## Purpose

The HISTDATA command creates a menu with detailed and/or summary performance data history files to which you have access. The command is valid in both basic and performance monitor mode: please see the description in the performance monitor mode commands section for more details ([“HISTDATA” on page 195](#)).

# MONITOR

## Purpose

The MONITOR command switches Performance Toolkit into performance monitoring mode, and causes the Performance Data Selection Menu (FCX124) to be displayed. It is valid only in machines that are allowed to execute DIAGNOSE codes X'04' and X'90', i.e. which have been given command privilege class E.

## Format

►► MONitor ◄◄

## Parameters

No parameters are allowed or tested for.

## Usage

1. See section [“Performance Monitor Mode Subcommands”](#) on page 129 for subcommands valid in performance monitor mode.

## Results

The first screen shown in performance monitor mode is the Performance Data Selection Menu (FCX124). See the "Operation in Performance Monitor Mode" chapter in the [z/VM: Performance Toolkit Guide](#) for more information.


# MONSCAN

## Purpose


The MONSCAN command allows you to select a CP monitor file on disk or tape (created by the standard MONWRITE utility) as input for performance data analysis. When the specified file is found, a performance data scan mode is entered which looks almost identical to the normal real-time monitoring mode and which allows you to browse through the accumulated monitor data.

The MONSCAN monitoring mode can be also entered using the CMS Pipelines PERFKIT stage. To use PERFKIT as a PIPE stage, the FCXPIP MODULE must first be run to install the PERFKIT PIPE stage. This installation of the PIPE stage will last until CMS is re-IPLed.

## Format



Using PERFKIT as a PIPE stage:



## Parameters

### DISK

Indicates that a monitor data disk file is to be processed. The next few arguments describe the disk file:

***fn***

Is the file name of the MONWRITE disk file

***ft***

Is the file type of the MONWRITE disk file

***fm***

Is the file mode of the MONWRITE disk file. Default file mode is '\*', i.e. all accessed disks will be searched for the specified file if the file mode is omitted.

### TAPE

Indicates that a monitor data file on tape is to be processed, where

***device***

Describes the tape unit on which the monitor data tape has been mounted and which has been attached to the virtual machine. It can be defined as either a virtual device number or as one of the standard CMS tape unit names TAP*n*.

### FROM *hh:mm*

Allows specifying a start time which is to be searched for, where the time must be specified in exactly that format, i.e. 2 digits each for hours and minutes with a separating colon.

When found, the initial position for the data scan will be at the first batch of SAMPLE data which were created at or after the start time *hh:mm*.

***getmondatastage(s)***

The PIPE stage(s) used to get the monitor data and pass it to the primary input of the PERFKIT stage. See the "Using CMS Pipelines as Monitor Data Source" section in the [\*z/VM: Performance Toolkit Guide\*](#) for more information.



# QUIT

## Purpose

A QUIT command entered in basic mode ends Performance Toolkit operation and returns the screen to standard VM console mode.

## Format

➤ QUIT ➤

## Parameters

No parameters are allowed or tested for.

## Usage


1. The same function is performed if you press PF3 or PF15.
2. You will have to enter the QUIT command twice if permanent performance data collection has been activated (information message 192 will be shown after the first QUIT command). This is to prevent loss of accumulated performance data if the QUIT command has been accidentally entered (wrong PF-key pressed.).

# REDISP

## Purpose

The REDISP command switches the screen into 'Re-Display'-mode for browsing through the accumulated console log file (see the "Redisplay Mode" section in the [z/VM: Performance Toolkit Guide](#) as well as "Redisplay Mode Subcommands" on page 119).

## Format



## Parameters

### yyyymmdd

Is the file name of the log file. The current date is the default, which will also be assumed if an asterisk (\*) is entered for the file name.

### ft

Is the file type of the log file. The default file types are CONLOG (for the current log), or CONLOG1, CONLOG2, or CONLOG3 for previous log generations. They will be looked for in this order if no file type is specified or if it has been entered as an asterisk.

### fm

Is the file mode of the log file. The default file mode is an asterisk (\*), that is, all accessed disks will be searched for the log file.

Only actual CONLOG files should be displayed in this way. Selection of other files is possible, but will lead to unpredictable results.

### Caution:

An explicit file mode is often specified because another user's Performance Toolkit log file is to be re-displayed (the system operator's log, for example). If such a log is still active, the corresponding minidisk should be re-accessed just before entering re-display mode, so that an updated version of the minidisk directory is used. Even so, you may expect read errors after the log has been updated again, and Performance Toolkit might even ABEND.

### (SUPpress n

Indicates that one or more of the originally suppressed output line classes (controlled by FC PROCESS) are to be suppressed in redisplay mode. Specify the classes which are *not* to be shown. Any combination of the classes 1, 2, and 3 is accepted; they can be separated by blanks or entered as a single string. *None* of the data lines in the console log will be suppressed if the option is omitted.

## Usage

1. Re-display mode for the current log can also be entered by pressing PF2 or PF14.

# REPLY

## Purpose

The REPLY command resets the message or action message line specified so that it will be scrolled again (similar to the DELETE command) and then executes the command text specified.

## Format

►► Reply — <i>n1n1 - n2</i> — ' <i>command text</i> ' ►◄
--

## Parameters

### *n1 - n2*

Are the line numbers that are to be reset (maximum 2 characters long).

When a range of numbers *n1 - n2* has been specified, all lines whose numbers are within the defined range will be reset. The first line number of a range *must* be one of the line numbers on the screen, while the second number can be any other number.

### *'command text'*

Is any valid command to be executed after resetting the screen lines specified (to be entered without apostrophes.). The command text is assumed to begin with the first argument that is not a valid screen line number.

# RETURN

## Purpose

A RETURN command entered in basic mode ends Performance Toolkit operation and returns the screen to standard VM console mode. It works exactly like the QUIT command.

## Format

➤ RETURN ➤

## Parameters

No parameters are allowed or tested for.

## Usage

1. The same function is performed if you press PF12.


**Note:** You will have to enter the RETURN command twice if permanent performance data collection has been activated (information message 192 will be shown after the first RETURN command). This is to prevent loss of accumulated performance data if the RETURN command has been accidentally entered (wrong PF-key pressed.).

# TRNDSCAN

## Purpose

The TRNDSCAN command allows selecting a Performance Toolkit extended trend file on disk as input for performance data analysis. When the specified file is found, a trend file scan mode is entered which looks very similar to the normal real-time monitoring mode and which allows you to browse through the accumulated trend file history data.

## Format



## Parameters

### *fn*

Is the file name of the extended trend data file (originally the system identification of the VM system where the trend file was created).

### *ft*

Is the file type of the extended trend data file. Default file type is 'FCXTREND', the original file type used by Performance Toolkit when writing the trend file.

### *fm*

Is the file mode of the extended trend data file. Default file mode is '\*', i.e. all accessed disks will be searched for the specified file if the file mode is omitted.

### **FROM yyyyymmdd**

Allows specifying a start date which is to be searched for, where the date must be specified in exactly that format, i.e. 4 digits for the year and 2 digits each for months and days.

When found, the initial position for the data scan will be at the first batch of trend records which were created at or after the start date.

## **'ENTER'**

### **Purpose**

Pressing the 'ENTER'-key without any input in basic mode:

- Locks the screen if it was in scroll wait with 'auto scroll' set on.
- Frees and scrolls a screen which had previously been frozen either by pressing 'ENTER' in 'auto scroll' mode or because 'manual scroll' mode is active.
- Is ignored if the screen was not in scroll wait.

## Redisplay Mode Subcommands

---

Only general subcommands and the following redisplay-specific subcommands are valid in redisplay mode. Entering any other command in this mode will cause the screen to be switched back into basic mode, and the command to be executed as CP/CMS command


All references to PF-key assignments are based on the initial default settings. Different PF-keys may have to be used if the initial settings have been changed.

# BACKWARD

## Purpose

The BACKWARD command scrolls the console log backwards. i.e. towards the beginning of the file. UP is accepted as a synonym for the BACKWARD command.

## Format



## Parameters

### *n1*

Specifies the number of lines to be shifted. The default scroll value is screen size minus one.

## Usage

1. Pressing PF7 or PF19 will also scroll the log backwards for the default number of lines.



# BOTTOM

## Purpose

The BOTTOM command causes the end of the console log to be shown, which is the same as the initial screen upon entering redisplay mode.

## Format

➤ BOTTOM ➤

## Usage

1. The end of the log file will also be shown after pressing PF5 or PF17.

# FORWARD

## Purpose

The FORWARD command scrolls the screen forward, i.e. towards the end of the console log, for the number of lines specified. NEXT and DOWN are accepted as synonyms for the FORWARD command.

## Format



## Parameters

### *n1*

Indicates the number of lines by which the console log is to be moved. The default scroll value is screen size minus one.

## Usage

1. Pressing PF8 or PF20 will also scroll the log forward for the default number of lines.

# LEFT

## Purpose

The LEFT command displays the left side of the console log. This is also the initial setting after switching into redisplay mode.

## Format

▶▶ Left ▶▶

## Usage

1. The same result can be achieved by pressing PF10.

# RETURN

## Purpose

The RETURN command, when entered in redisplay mode, exits redisplay mode and returns to basic mode.

## Format



## Usage

1. 'QUIT' is accepted as a synonym for the RETURN command.
2. The same function can also be performed by pressing PF12.

# RIGHT

## Purpose

The RIGHT command displays the right side of the console log.

## Format

➤ RIGht ➤

## Usage

1. The same function can also be performed by pressing PF11.

## TOP

### Purpose

The TOP command displays the beginning of the console log.

### Format

➤ Top ➤

### Usage

1. The beginning of the log will also be shown after pressing PF4 or PF16.

## / (LOCATE)

### Purpose

The LOCATE command searches the console log for the character string specified.

### Format



### Parameters

/

Indicates a forward search. The search for the specified string starts with the current line + 1 and continues until either the string is found or the end of the log file is reached. 'Current line' is the top line for the first locate command, it is the line where the last string has been located for further locate commands that immediately follow a previous successful locate command.

-/

Indicates a backward search. The search starts with the current line - 1 and continues until either the string is found or the top of the log file is reached.

'string'

Is the character string to be located (do not enter the apostrophes unless they are part of the string.). The string and the console log records will be translated to upper case for the search. The output will be scrolled, if necessary, to show the first record containing this string which will then be placed at the top of the screen. The first string found in this record will be highlighted. It will be displayed with reverse video on screens that support extended highlighting, on other displays the display intensity will be changed. The log shown will be shifted right or left, if necessary, so that the beginning of the located string is always visible.

### Usage

1. The search may be continued by just pressing the 'ENTER'-key again.

## **'ENTER'**

### **Purpose**

Pressing the 'ENTER'-key without any input will:

- Restart the search for the character string specified with the LOCATE command if the last command entered was a 'LOCATE', or it will
- Scroll the console log backward or forward for the default number of lines, depending upon whether a BACKWARD or FORWARD command (or the corresponding PF-key function) was executed last.



## Performance Monitor Mode Subcommands

---

Only general Performance Toolkit subcommands and the following performance monitor-specific subcommands are valid in performance monitor mode.

Note that most of the selectable performance displays are based on CP monitor data: data can be calculated, and such screens selected, only if CP monitor data collection has previously been activated, and if Performance Toolkit has been set up to retrieve monitor data.

All references to PF-key assignments are based on the initial default settings. Different PF-keys may have to be used if the initial settings have been changed.

# AGELLOG

## Purpose

The AGELLOG subcommand causes the Age List Log screen (FCX297) to be displayed.

## Format

➤ AGELlog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX297, Age List Log Screen – AGELLOG”](#) on page 809.

# AUXLOG

## Purpose

The AUXLOG subcommand causes the Auxiliary Storage Log Screen (FCX146) to be displayed.

## Format

➤ AUXlog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX146, Auxiliary Storage Log Screen – AUXLOG”](#) on page 509.

# AVAILLOG

## Purpose

The AVAILLOG command causes the Available List Log Screen (FCX254) to be displayed.

This command applies only when reducing MONWRITE data gathered when running Performance Toolkit against z/VM 6.2 or earlier data.

## Format

➤ AVAILlog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX254, Available List Log Screen – AVAILLOG”](#) on page 728.

# AVERAGE

## Purpose

The AVERAGE subcommand allows many performance screens to be switched from a display of the last 'interim' interval's or the last sample's data to the display of averages based on all measurements made since the last RESET.

## Format

➤ Average ➤

## Parameters

No parameters are allowed or tested for.

## Usage

1. The AVERAGE subcommand affects *all* other performance displays that can show both current and average data.  
Average data will be shown immediately on all screens which are based on CP monitor data, but average data for the other displays will be shown only when the next sample has been taken, so you may see no immediate effect of the subcommand having been executed.
2. See the CURRENT subcommand ([“CURRENT” on page 165](#)) for an explanation of how to set the screens back to their initial status, displaying 'current' data or the INTERIM subcommand [“INTERIM” on page 200](#) for displaying 'interim' data.

## AVLA2GLG

### Purpose

The AVLA2GLG subcommand causes the Available List Data Above 2G screen (FCX295) to be displayed.

### Format

➤ AVLA2glg ➤
--------------

### Parameters

No parameters are allowed or tested for.

### Results

See [“FCX295, Available List Data Above 2G Screen – AVLA2GLG ”](#) on page 806.

## AVLB2GLG

### Purpose

The AVLB2GLG subcommand causes the Available List Data Below 2G screen (FCX294) to be displayed.

### Format

➤ AVLB2glg ➤

### Parameters

No parameters are allowed or tested for.

### Results

See [“FCX294, Available List Data Below 2G Screen – AVLB2GLG ”](#) on page 804.





# BASMODE

## Purpose

The BASMODE subcommand causes the Remote Basic Mode Screen (FCX186) to be displayed. It is valid only for remote performance monitoring sessions using APPC/VM, and will be accepted only if the user has been given CMD authorization on the remote system.

## Format

➤ BASMode ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX186, Remote Basic Mode Screen – BASMODE”](#) on page 593.

# BENCHMRK

## Purpose

The BENCHMRK subcommand causes the Benchmark Overview Menu (FCX173) to be displayed.

## Format

➤ BENCHmrk ➤

## Parameters

No parameters are allowed or tested for.

## Results

See the "Benchmarking Objects" section in the [z/VM: Performance Toolkit Guide](#) for more information on this screen.

# BFSFILES

## Purpose

The BFSFILES subcommand causes the Byte File System File Activity Log (FCX190) to be displayed.

## Format

➤ BFSFiles ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX190, Byte File System File Activity Log Screen – BFSFILES”](#) on page 599.

# BFSLOCKS

## Purpose

The BFSLOCKS subcommand causes the Byte File System Locking Activity Log (FCX192) to be displayed.

## Format

➤ BFSLocks ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX192, Byte File System Locking Activity Log Screen – BFSLOCKS”](#) on page 604.

# BFSPIPES

## Purpose

The BFSPIPES subcommand causes the Byte File System PIPEs and Linking Activity Log (FCX191) to be displayed.

## Format

➤ BFSPipes ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX191, Byte File System PIPEs and Linking Activity Log Screen – BFSPIPES”](#) on page 601.

# BOTTOM

## Purpose

The BOTTOM subcommand causes the last part of the selected performance monitor data to be shown.

## Format

➤- B0ttom ->◀

## Usage

1. The bottom part of the current display's data will also be shown after pressing PF5 or PF17.

# CACHDBSE

## Purpose

The CACHDBSE subcommand causes the CU-Cached Disks Screen (FCX111) to be displayed.

## Format

➡ CACHdbse ⬅

## Parameters

No parameters are allowed or tested for.

## Usage

1. The whole analysis is based on counters which are updated within the cached control units. These counters reflect the load and behavior of specific **disks**, but they do not tell where the I/O came from. If several systems share a device, the **total** of all activity on the disk is displayed, and **not** just the I/Os from the system which is doing the analysis.  
  
The same problem exists for systems which are running second level – an analysis by the second level system will include any I/O activity to the same devices from the first level system.
2. Second level systems may face an additional problem: since they only have one path to each device, they will retrieve data only from one cached control unit director. If a device is connected to the first level system over two paths and two different control units (e.g. 3880-23), the first level system may pass the data retrieval request randomly to either of the two control unit directors. Figures obtained by the second level system in this way will tend to be reasonably correct in the long run, but cannot be relied on for single intervals.

## Results

See [“FCX111, CU-Cached Disks Screen – CACHDBSE” on page 428](#).

# CACHDLOG

## Purpose

The CACHDLOG subcommand causes the Control Unit Cache Log Screen (FCX169) to be displayed for the selected disk.

## Format

➤ CACHDLog — <i>devno</i> ➤
-----------------------------

## Parameters

### *devno*

Is the device number of the disk for which the Control Unit Cache Log Screen (FCX169) is to be displayed.

## Results

See [“FCX169, Control Unit Cache Log Screen – CACHDLOG”](#) on page 552.



# CACHELOG

## Purpose

The CACHELOG subcommand causes the Cache Extended Functions Log Screen (FCX184) to be displayed for the selected disk.

## Format

➤ CACHELog — *devno* ➤

## Parameters

### *devno*

Is the device number of the disk for which the Cache Extended Functions Log Screen (FCX184) is to be displayed.

## Results

See [“FCX184, Cache Extended Functions Log Screen – CACHELOG”](#) on page 587.

# CACHEXT

## Purpose

The CACHEXT subcommand causes the Cache Extended Functions Screen (FCX177) to be displayed.

## Format

➤ CACHExt ➤

## Parameters

No parameters are allowed or tested for.

## Usage

1. The whole analysis is based on counters which are updated within the cached control units. These counters reflect the load and behaviour of specific **disks**, but they do not tell where the I/O came from. If several systems share a device, the **total** of all activity on the disk is displayed, and **not** just the I/Os from the system which is doing the analysis.
2. The same problem exists for systems which are running second level. An analysis by the second level system will include any I/O activity to the same devices from the first level system.

## Results

See [“FCX177, Cache Extended Functions Screen – CACHEXT” on page 569](#).

# CHANNEL

## Purpose

The CHANNEL subcommand causes the Channel Load Screen (FCX107) to be displayed.

## Format

➤ CHANNEL ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX107, Channel Load Screen – CHANNEL”](#) on page 410.

# CMS

## Purpose

The CMS subcommand prefix must be used when entering a CMS command in performance monitor mode.

## Format

➤ *'CMS command'* ➤

## Parameters

### 'CMS command'

Is any valid CMS command or EXEC name (do not enter the apostrophes).

## Usage

1. The CMS subcommand prefix will cause an immediate switch back to basic mode where the command will be executed.

## Results

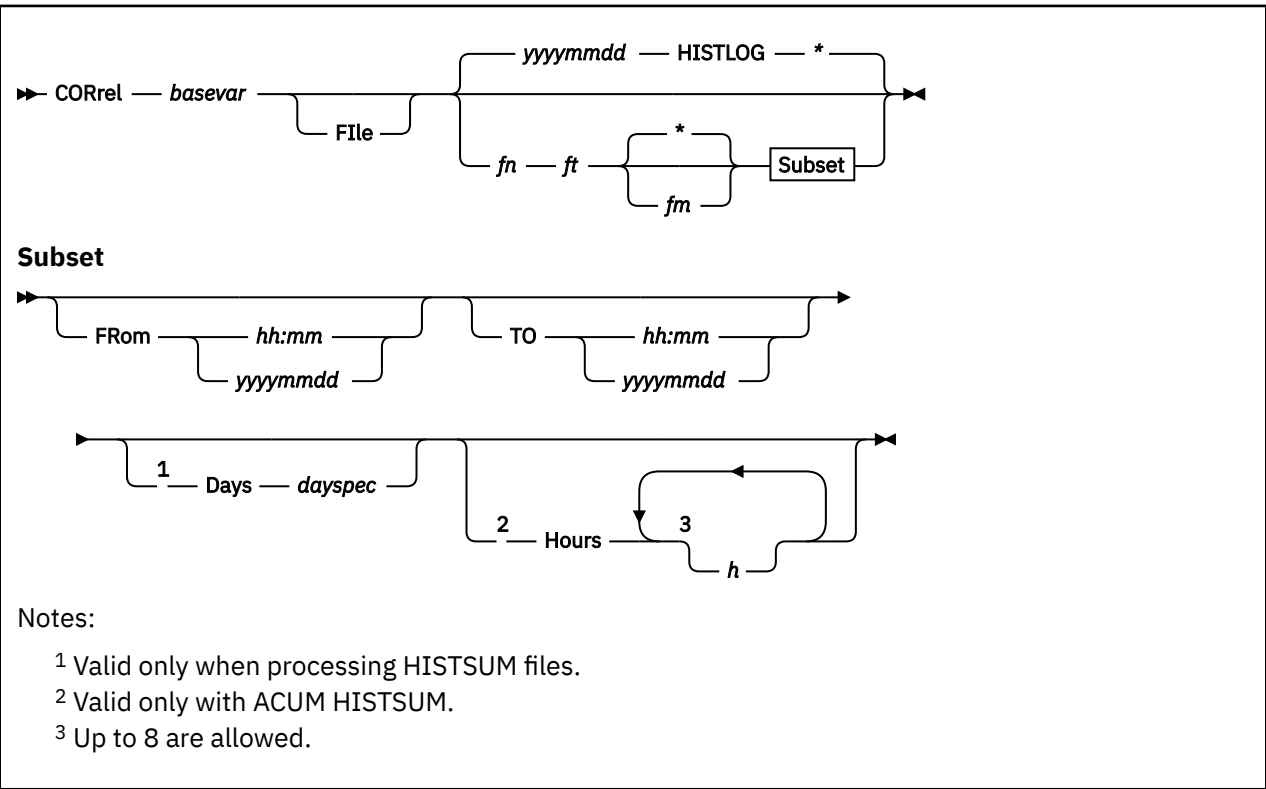
See the discussion of the CMS subcommand in basic mode ([“CMS” on page 102](#)).

# CORREL

## Purpose

The CORREL subcommand selects the Variable Correlation Screen (FCX194), with correlation coefficients based on the values of the variables found in the selected history log file.

## Format



## Parameters

### **basevar**

Is the name of the variable that is to be used as the 'base' variable when calculating correlation coefficients.

### **File fn ft fm**

Specifies the file ID of the HISTLOG/HISTSUM file to be processed. The current day's HISTLOG file will be used as input if the file specification is omitted.

### **fn**

Is the file name of the performance history file. The file name usually contains the file's creation date, in format **yyyyymmdd**.

### **ft**

Is the file type of the log file. There is no default file type: If a file name has been entered, the file type is required too.

### **fm**

Is the file mode of the log file. Default file mode is **'\*'**, i.e. all accessed disks will be searched for the file.

**FRom *hh:mm***

Defines the time stamp of the first record to be analyzed when a detailed HISTLOG*n* file is used as input. The time entered must have exactly this format, i.e. two numbers each for hours and minutes, separated by a colon.

**FRom *yyyymmdd***

Defines the date of the first record (year, month and day) to be analyzed when a summary HISTSUM file is used as input. The date entered must have exactly this format, i.e. four digits for the year, and two digits each for month and day, without any intervening blanks. The format *mmdyy* used in previous releases of the program is also accepted.

Analysis will begin with the first data record found if no 'FROM' date has been entered.

**TO *hh:mm***

Defines the time stamp of the last record to be analyzed when a detailed HISTLOG*n* file is used as input.

**TO *yyyymmdd***

Defines the last day for which data are to be included in the analysis when data from a condensed HISTSUM data file are to be analyzed.

Analysis will continue to the end of the file if the 'TO' argument is omitted.

**Days *dayspec***

Defines a selection of weekdays for which data are to be analyzed, where *dayspec* can be a combination of the following strings:

**ALL**

If all weekdays are to contribute to the analysis

**M-F**

If only the normal workdays Monday to Friday are to contribute data

**MON**

If all Mondays are to contribute data

**TUE**

If all Tuesdays are to contribute data

**WED**

If all Wednesdays are to contribute data

**THU**

If all Thursdays are to contribute data

**FRI**

If all Fridays are to contribute data

**SAT**

If all Saturdays are to contribute data

**SUN**

If all Sundays are to contribute data

The argument is valid only when analyzing a condensed summary data file ACUM HISTSUM; *all* data will be analyzed when it is omitted. When specified, only performance data collected for the selected weekdays are used in the correlation analysis. For example, specifying 'DAYS M-F SAT' will cause data to be selected for the weekdays Monday to Saturday, excluding any data for Sundays.

**HOURS *h1 h2 ... h8***

Defines a selection of up to eight one-hour intervals for which data are to be analyzed. The argument is valid only when analyzing a condensed summary data file ACUM HISTSUM. When specified, only performance data collected for the selected hours are used in the correlation analysis. For example, specifying 'HOURS 9 10 14 15' will cause data to be selected for the periods of 09:00-10:59 and 14:00-15:59, for each of the days previously selected with the 'FROM' and 'TO' and/or 'DAYS' arguments.

All hours for which data are available will be included when the HOURS argument is not specified.

## Usage

1. Only actual performance history files should be specified as input. Selection of other files may be possible if they pass some tests, but will lead to unpredictable results.

## Results


See [“FCX194, Variable Correlation Screen – CORREL” on page 607](#).

# CP

## Purpose

The CP subcommand prefix must be used when entering a CP command in performance monitor mode.

## Format



## Parameters

### 'CP command'

Is any valid CP command (do not enter the apostrophes). When entered in performance monitoring mode, the CP subcommand prefix will cause an immediate switch back to basic mode where the command will be executed. Up to 4096 bytes of the command output (if any) will be returned when CP commands are executed from remote performance monitor mode.

## Usage

1. CP MODE will be entered if no command follows, and VM's CP interface will then be disabled so that all command input and output is shown on the console as usual. Commands entered in this mode will not be contained in the toolkit's log.
2. Enter 'BEGIN' to switch back to Performance Toolkit-mode (required only if 'CP' has been entered without a command following it).

## Results

See the discussion of the CP subcommand in basic mode ([“CP” on page 103](#)).



# CPLACT

## Purpose

The CPLACT subcommand causes the CPU Pool Activity Screen (FCX309) to be displayed.

## Format

➤ CPLAct ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX309, CPU Pool Activity Screen – CPLACT”](#) on page 835.

# CPLCONF

## Purpose

The CPLCONF subcommand causes the CPU Pool Configuration Screen (FCX308) to be displayed.

## Format

➤ CPLConf ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX308, CPU Pool Configuration Screen – CPLCONF”](#) on page 833.

# CPLMENU

## Purpose

The CPLMENU subcommand causes the CPU Pool Menu Screen (FCX324) to be displayed.

## Format

➤ CPLMenu ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX324, CPU Pool Menu Screen – CPLMENU”](#) on page 857.

# CPOWNLOG

## Purpose

The CPOWNLOG subcommand causes the CPOWNED Device Log Screen (FCX170) to be displayed for the selected disk.

## Format

►► CPOWNLog — *devno* ◄◄

## Parameters

### *devno*

Is the number of the disk for which the CPOWNED Device Log Screen (FCX170) is to be displayed.

## Results

See [“FCX170, CPOWNED Device Log Screen – CPOWNLOG”](#) on page 553.

# CPSALL

## Purpose

The CPSALL subcommand causes the CP IUCV Services Log Screen (FCX139) for all communication activity to be displayed.

## Format

➤ CPSALL ➤

## Parameters

No parameters are allowed or tested for.

## Results

The CP IUCV Services Log Screen (FCX139/140/141/142) is shown after entering the CPSALL, CPSFAIL, CPSTOCP, or CPSTOVM subcommand. Note that the layout of these four displays is almost identical. See [“FCX139/140/141/142, CP IUCV Services Log Screen – CPSALL, CPSTOCP, CPSTOVM, CPSFAIL”](#) on page 500 for the sample CPSALL layout.

# CPSERV

## Purpose

The CPSERV subcommand causes the CP IUCV Services Screen (FCX106) to be displayed, with counts for IUCV and VMCF services.

## Format

➤ CPServ ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX106, CP IUCV Services Screen – CPSERV”](#) on page 408.

# CPSFAILS

## Purpose

The CPSFAILS subcommand causes the CP IUCV Services Log Screen (FCX142) for failed communications to be displayed.

## Format

➤ CPSFails ➤

## Parameters

No parameters are allowed or tested for.

## Results

The CP IUCV Services Log Screen (FCX139/140/141/142) is shown after entering the CPSALL, CPSFAIL, CPSTOCP, or CPSTOVM subcommand. Note that the layout of these four displays is almost identical. See [“FCX139/140/141/142, CP IUCV Services Log Screen – CPSALL, CPSTOCP, CPSTOVM, CPSFAIL”](#) on page 500 for the sample CPSALL layout.

# CPSTOCP

## Purpose

The CPSTOCP subcommand causes the CP IUCV Services Log Screen (FCX140) for communications from a virtual machine to CP to be displayed.

## Format

➤ CPSTOCP ➤

## Parameters

No parameters are allowed or tested for.

## Results

The CP IUCV Services Log Screen (FCX139/140/141/142) is shown after entering the CPSALL, CPSFAIL, CPSTOCP, or CPSTOVM subcommand. Note that the layout of these four displays is almost identical. See [“FCX139/140/141/142, CP IUCV Services Log Screen – CPSALL, CPSTOCP, CPSTOVM, CPSFAIL”](#) on page 500 for the sample CPSALL layout.



# CPSTOVM

## Purpose

The CPSTOVM subcommand causes the CP IUCV Services Log Screen (FCX141) for communications from CP to a virtual machine to be displayed.

## Format

➤ CPSTovm ➤

## Parameters

No parameters are allowed or tested for.

## Results

The CP IUCV Services Log Screen (FCX139/140/141/142) is shown after entering the CPSALL, CPSFAIL, CPSTOCP, or CPSTOVM subcommand. Note that the layout of these four displays is almost identical. See [“FCX139/140/141/142, CP IUCV Services Log Screen – CPSALL, CPSTOCP, CPSTOVM, CPSFAIL”](#) on page 500 for the sample CPSALL layout.

# CPU

## Purpose

The CPU subcommand causes the CPU Load and Transactions Screen (FCX100) to be displayed.

## Format

➤ CPU ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX100, CPU Load and Transactions – CPU”](#) on page 371.

# CPUMENU

## Purpose

The CPUMENU subcommand causes the CPU Activity Selection Menu Screen (FCX325) to be displayed.

## Format

➤ CPUMenu ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX325, CPU Activity Selection Menu Screen – CPUMENU”](#) on page 858.

# CTLUNIT

## Purpose

The CTLUNIT subcommand causes the Cached Control Units Screen (FCX176) to be displayed.

## Format

➤ CTLUnit ➤

## Usage

1. The whole analysis is based on counters which are updated within the cached control units. These counters reflect the load and behaviour of specific **disks**, but they do not tell where the I/O came from. If several systems share a control unit, the **total** of all activity on the control unit is displayed, and **not** just the I/Os from the system which is doing the analysis.
2. The same problem exists for systems which are running second level. An analysis by the second level system will include any I/O activity to the same control units from the first level system.
3. This report works correctly only if every logical control unit (LCU) to which your system is attached has a unique subsystem ID (SSID).

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX176, Cached Control Units Screen – CTLUNIT”](#) on page 566.

# CURRENT

## Purpose

The CURRENT subcommand causes data from the last measurement interval to be displayed. It is valid for all performance screens where the 'AVERAGE' subcommand can also be entered.

## Format

➡ CCurrent ➡

## Parameters

No parameters are allowed or tested for.

## Usage

1. The CURRENT subcommand affects *all* other performance displays that can show both current and average data.  
  
Current data will be shown immediately on all screens which are based on CP monitor data, but current data for the other displays will be shown only when the next sample has been taken, so you may see no immediate effect of the subcommand having been executed.
2. See the AVERAGE ([“AVERAGE” on page 133](#)) and INTERIM ([“INTERIM” on page 200](#)) subcommands for an explanation of how to select averages for other periods.

# DASDLOG

## Purpose

The DASDLOG subcommand causes the DASD Performance Log Screen (FCX183) to be displayed.

## Format

➤ DASDlog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX183, DASD Performance Log Screen – DASDLOG”](#) on page 586.

# DEMNDLOG

## Purpose

This command applies only when reducing MONWRITE data gathered when running Performance Toolkit against z/VM 6.2 or earlier data.

## Format

➤ DEMNDlog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX259, Demand Scan Log Screen – DEMNDLOG”](#) on page 736.

# DEVCONF

## Purpose

The DEVCONF subcommand causes the I/O Device Configuration Screen (FCX131) to be displayed.

## Format

➤ DEVConf ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX131, I/O Device Configuration Screen – DEVCONF”](#) on page 482.



# DEVICE

## Purpose

The DEVICE subcommand causes the I/O Device Details Screen (FCX110), the CP Owned Device Screen (FCX109), the General I/O Device Screen (FCX108), or the Device HPF Screen (FCX282) to be displayed, based on how you specify the subcommand. (See below.) 'I/O' is also accepted as a synonym for the DEVICE subcommand.

## Format



## Parameters

### **devno**

Displays the I/O Device Details Screen (FCX110) for the device with real device number *devno*.

### **Cpowned**

Displays the CP Owned Device Screen (FCX109), which shows usage of system areas (paging, spooling, directory cylinders, T-disk space) and the paging and spooling activity on these areas.

### **DASD**

Displays the General I/O Device Screen (FCX108), with only disks selected (instead of all devices) for which the monitor I/O domain has been enabled.

### **HPF**

Displays the HPF I/O Device Screen (FCX282) for all devices for which monitor I/O domain had been enabled.

If you specify the DEVICE subcommand with no parameters at all, the General I/O Device Screen (FCX108) will be displayed, for all devices.

## Results

See [“FCX108, General I/O Device Screen – DEVICE or DEVICE DASD”](#) on page 412, if you specified DEVICE or DEVICE DASD.

See [“FCX109, CP-Owned Device Screen – DEVICE CPOWNED”](#) on page 416, if you specified DEVICE CPOWNED.

See [“FCX110, I/O Device Details Screen – DEVICE”](#) on page 420, if you specified DEVICE *devno*.

See [“FCX282, HPF I/O Device Screen – DEVICE HPF”](#) on page 778, if you specified DEVICE HPF .

# DEVLOG

## Purpose

The DEVLOG subcommand causes the General I/O Device Data Log Screen (FCX168) to be displayed for the selected I/O device.

➤ DEVLog — *devno* ➤

## Format

## Parameters

### *devno*

Is the number of the I/O device for which the General I/O Device Data Log Screen (FCX168) is to be displayed.

## Results

See [“FCX168, General I/O Device Data Log Screen – DEVLOG”](#) on page 548.

# DEVMENU

## Purpose

The DEVMENU subcommand causes the Device Report Selection Menu (FCX284) to be displayed.

## Format

➤ DEVMENU ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX284, I/O Device Data Selection Menu Screen – DEVMENU”](#) on page 785.

# DSPACESH

## Purpose

The DSPACESH subcommand causes the Shared Data Spaces Screen (FCX134) to be displayed.

## Format

►► DSPACESH ◄◄

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX134, Shared Data Spaces Screen – DSPACESH” on page 489](#).

# DSVBKACT

## Purpose

The DSVBKACT subcommand causes the Dispatch Vector Activity screen (FCX301) to be displayed.

## Format

➤ DSVBKact ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX301, Dispatch Vector Activity Screen – DSVBKACT”](#) on page 818.

# DSVCLOG

## Purpose

The DSVCLLOG subcommand causes the Dispatch Vector Configuration Log screen (FCX300) to be displayed.

## Format

➡ DSVCLlog ➡

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX300, Dispatch Vector Configuration Change Log Screen – DSVCLLOG”](#) on page 816.

# DSVSLOG

## Purpose

The DSVSLOG subcommand causes the DSVBK Steals per Processor Log Screen (FCX303) to be displayed.

## Format

➤ DSVSLog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX303, DSVBK Steals per Processor Log Screen – DSVSLOG”](#) on page 822.

# EVNIC

## Purpose

The EVNIC subcommand causes the Extended Virtual Network Device Activity Screen (FCX270) to be displayed. This screen pertains mostly to the array of OSA devices.

This screen will display performance data above and beyond that displayed by the VNIC subcommand. See [“VNIC” on page 360](#) for more information.

For a summary of virtual network device definitions, see [“GVNIC” on page 192](#).

## Format

➡ EVNIC ⬅

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX270, Extended Virtual Network Device Activity Screen – EVNIC” on page 759](#).



# EVSWITCH

## Purpose

The EVSWITCH subcommand causes the Extended Virtual Switch Activity Screen (FCX267) to be displayed.

This screen will display performance data above and beyond that displayed by the VSWITCH subcommand. See [“VSWITCH” on page 363](#) for more information.

For a summary of VSWITCH definitions, see [“GVSWITCH” on page 193](#).

## Format

▶▶ EVSWItch ▶▶

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX267, Extended Virtual Switch Activity Screen – EVSWITCH” on page 753](#).

# EXCEPT

## Purpose

The EXCEPT subcommand causes the Exception Log Screen (FCX153) to be displayed.

## Format

➤ EXCEPT ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX153, Exception Log Screen – EXCEPT”](#) on page 524.

# FCHANNEL

## Purpose

The FCHANNEL subcommand causes the FICON Channel Load Screen (FCX215) to be displayed.

## Format

➤ FCHannel ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX215, FICON Channel Load Screen – FCHANNEL”](#) on page 657.

# FORWARD

## Purpose

The FORWARD subcommand scrolls the selected performance monitor display forward, that is, toward the end of the data. NEXT and DOWN are accepted as synonyms for the FORWARD subcommand.

## Format



## Parameters

### nn

Indicates the number of lines to be moved. The default scroll value is equal to the number of performance data lines (excluding heading lines) on the current monitor display minus one, that is, it depends on the performance screen being viewed.

## Usage

1. Pressing PF8 or PF20 will also scroll the data forward for the default number of lines.

# GLONACT

## Purpose

The GLONACT subcommand causes the Global Networking Object Activity Screen (FCX317) to be displayed. This report supports the monitor data that is obtained when running Performance Toolkit against data from z/VM 6.3 with APARs VM65583, PI21053, and VM65528, or later.

## Format

➤ GLONact ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX317, Global Networking Object Activity Screen – GLONACT”](#) on page 849.

# GRAPHDET / PLOTDET

## Purpose

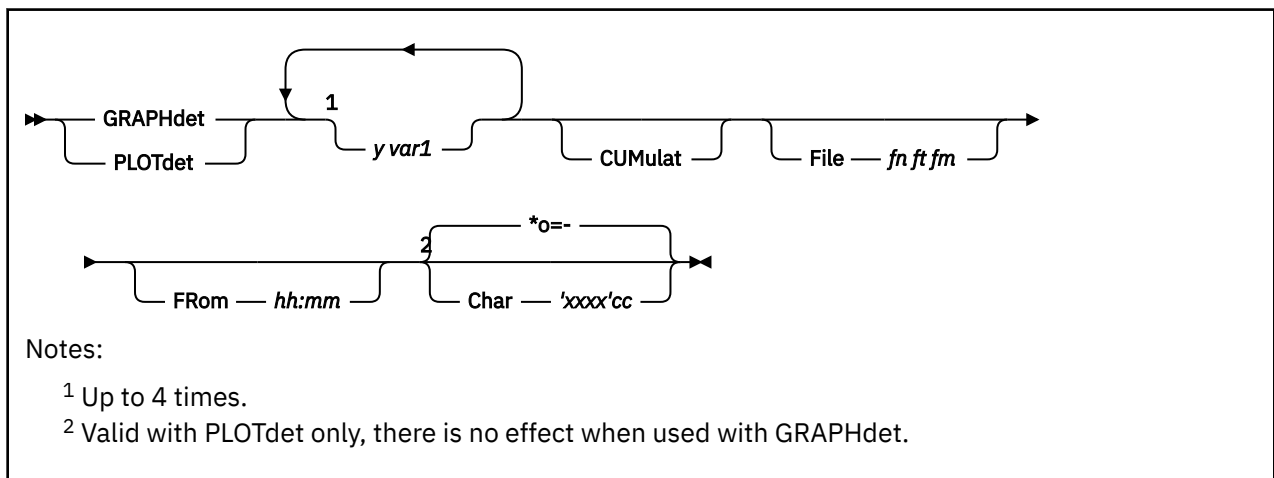
The GRAPHDET and PLOTDET subcommands allow displaying selected parts of the performance data accumulated in the monitor redisplay buffer or in one of the detailed history data logs in the form of history graphics or plots.

Both subcommands will create *detailed* history displays where elapsed time is plotted on the horizontal x-axis, while the corresponding values of up to four performance variables are plotted on the vertical y-scale. Data will be plotted for up to 140 measurements, where each plot position is determined by the average value of two consecutive measurements if the default monitor interval of 60 seconds is used. Only one measurement will be shown per plot position if the monitor interval is two minutes or more.

Two display modes are possible:

- The GRAPHDET subcommand requires GDDM services (Graphical Data Display Manager V.2 or 3) and a display terminal with graphics capability for generating and displaying real graphics output. It cannot operate without these prerequisites.
- The PLOTDET subcommand uses the same data, but it generates a simple plot using ordinary display characters for the output. It will work on any display terminal, and it has no additional software prerequisites.

## Format



## Parameters

### y-var1, y-var2, ...

define the variables whose values are to be plotted on the vertical y-axis.

All y-var names define performance variables from either the monitor redisplay screen (default), from a HISTLOG or PERFLOG file with general system performance data, or from an object's benchmark log file on disk.

See “REDISP” on page 269 (PERFxxx files), “REDHIST” on page 268 (HISTxxx files) for a list of performance variables and their description, or see the description of the selected benchmark log display for information on the selectable performance variables (note that only numerical fields with load data can be selected).

You must enter the full variable names exactly as shown (no abbreviations allowed), except for the case of the maximum user fields ('Max. CPU', 'Max. Vector', ...) of PERFLOG and HISTLOG files where the first eight characters only need be entered. Do not enter the apostrophes: they have been included only to indicate that two words are required for selecting these values.

In addition to these basic performance variables you can also specify the names of any 'user variables' which you previously defined by means of the FC USERVAR subcommand.

**CUMulat**

Indicates that the y-variable values are to be plotted cumulatively, i.e. each value added to the previous one. The string 'Cumul .' will be inserted above the y-axis when a cumulative display has been selected.

**File *fn ft fm***

Describes the input file to be analyzed, when specified. Current REDISP data in storage will be assumed as input if no input file is specified.

The input file can be a detailed system performance log file (HISTLOG, PERFLOG or RMONLOG), a user benchmark log file (such as MTUSRLG, USERLOG, etc.), or one of the supported I/O device benchmark log files (CACHELG, DEVLOG, etc.).

**Char 'xxxx'cc**

Defines the four plot characters to be used for the selected variables. This argument is intended for use with the PLOTDET subcommand only, it has no effect when used with GRAPHDET. Strings specified between apostrophes (') are interpreted as hexadecimal representation of special characters, i.e. the specification of 'f1f2f3' \* would result in the characters 123\* being used. Specification of plot characters in hex format allows the use of special graphic symbols for plotting, provided they are supported by the hardware. (Some terminal emulators used with work stations allow display of special symbols from the currently active code page, and some of these graphic symbols make very nice bar charts.) The default setting is '\*o=-'. If less than four characters are defined, the corresponding default characters will be inserted for the remaining variables.

**FFrom *hh:mm***

Defines the starting time for plotting. If the specified time is not within the time frame of the measurements currently available in the redisplay buffer, the oldest measurements will be shown (first entries in buffer).

**Results**

See "Detailed Performance History Display" on page 961 for more information on the specific displays caused by this subcommand, and [Appendix A, "Graphical Displays of Performance History Data,"](#) on page 955 for more general information.

# GRAPHICS

## Purpose

The GRAPHICS subcommand causes the Graphics Selection Menu (FCX128) to be displayed.

## Format

►► GRAPHICS ◄◄

## Parameters

No parameters are allowed or tested for.

## Results

See [“Graphics Selection Menu” on page 958](#) for more information on this specific display, and [Appendix A, “Graphical Displays of Performance History Data,” on page 955](#) for more general information.



# GRAPHSUM / PLOTSUM

## Purpose

The GRAPHSUM and PLOTSUM subcommands allow selected parts of the performance data accumulated in the monitor redisplay buffer (default) or in one of the history data files to be displayed graphically.

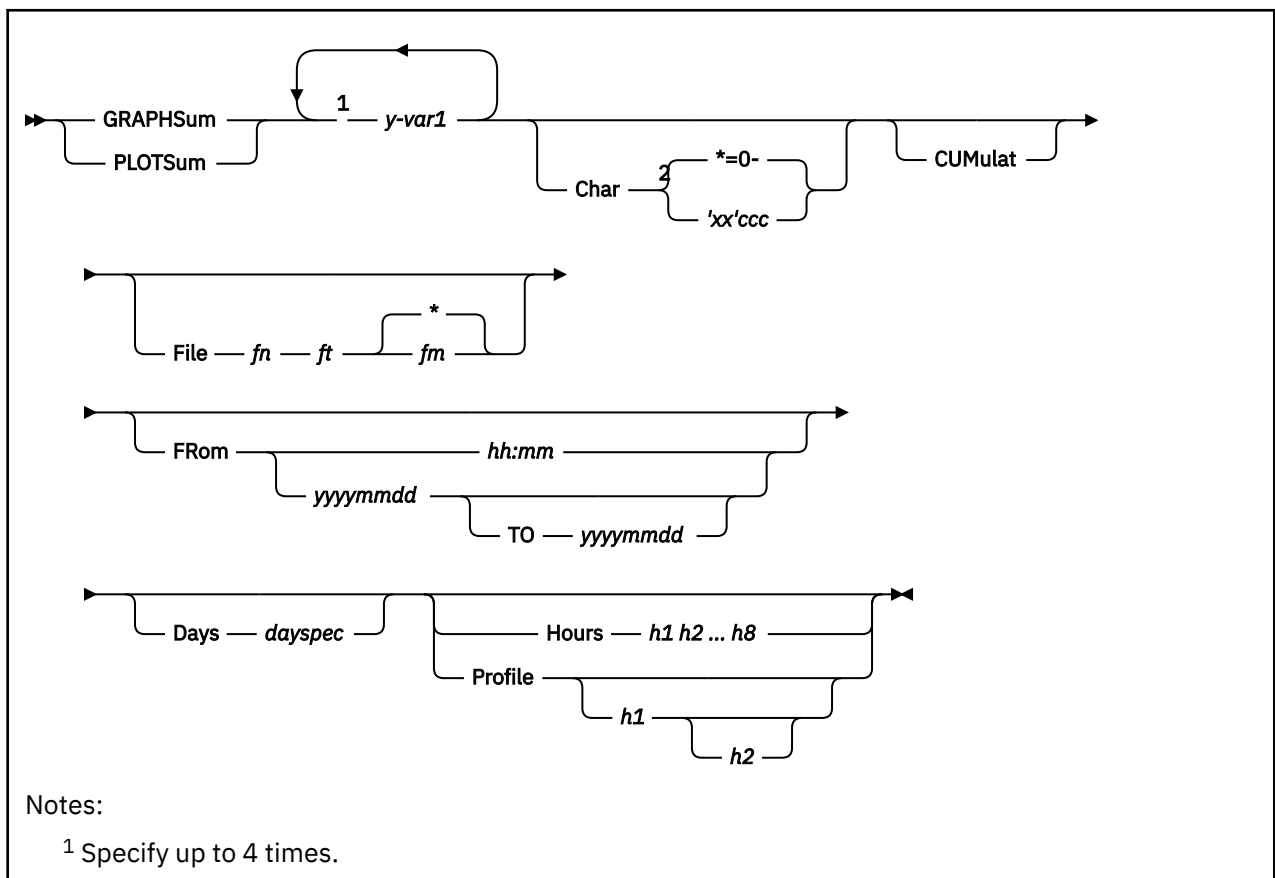
Both subcommands will create *summary* history displays where elapsed time is plotted on the horizontal x-axis, while the corresponding values of up to four performance variables are plotted on the vertical y-scale. The period for which data can be plotted depends on the source of the data:

- If data from the redisplay buffer (default), or from a detailed HISTLOG $n$ , PERFLOG $n$  or RMONLOG $n$  file, or from an object's benchmark log file are plotted, each plot position represents the average of a 10-minute interval, and data for periods of up to 12 hours can be displayed on one screen.
- If data from a HISTSUM, PERFHIST or FCXTREND file are to be displayed, any desired period can be selected, and the x-scale will automatically be adapted to the number of days to be shown. Data for at least one day will be shown per plot position.

Two display modes are possible:

- The GRAPHSUM subcommand requires GDDM services (Graphical Data Display Manager V.2 or 3) and a display terminal with graphics capability for generating and displaying real graphics displays. It cannot operate without these prerequisites.
- The PLOTSUM subcommand uses the same data, but it generates a simple plot using ordinary display characters for the output. It will work on any display terminal, and it has no additional software prerequisites.

## Format



<sup>2</sup> Valid with PLOTSum only, there is no effect when used with GRAPHSum.

## Parameters

### **y-var1**

Defines the variables whose values are to be plotted on the vertical y-axis.

All y-var names define performance variables from the monitor redisplay screen (default), from PERFLOG/PERFHIST, HISTLOG/HISTSUM or FCXTREND files with general system performance data, or from an object's benchmark log file on disk.

### **Char 'xx'ccc**

Defines the four plot characters to be used for the selected variables. This argument is intended for use with the PLOTSUM subcommand only, it has no effect when used with GRAPHSUM. Strings specified between apostrophes (') are interpreted as hexadecimal representation of special characters, i.e. the specification of 'f1f2f3' \* would result in the characters 123\* being used. Specification of plot characters in hex format allows the use of special graphic symbols for plotting, provided they are supported by the hardware. (Some terminal emulators used with work stations allow display of special symbols from the currently active code page, and some of these graphic symbols make very nice bar charts.) The default setting is '\*o=-'. If less than four characters are defined, the default characters will be inserted for the remaining variables.

### **CUMulat**

Indicates that the y-variable values are to be plotted cumulatively, i.e. each value added to the previous one. The string 'Cumul .' will be inserted above the y-axis when a cumulative display has been selected.

### **File fn ft**

Describes the input file to be analyzed. The input file can be a detailed system performance log file (HISTLOG, PERFLOG or RMONLOG), a summary system performance data file 'ACUM HISTSUM' (new format) or 'ACUM PERFHIST' (old format), with one record per hour, or a trend file 'systemid FCXTREND'. User benchmark log files (such as MTUSRLG, USERLOG, etc.), or one of the supported I/O device benchmark log files (CACHELG, DEVLOG, etc.) can also be processed. *The current redisplay buffer will be used as source if no disk file is specified as input.*

### **fm**

Defines the file mode of the source file. It is required only if several files with identical file names and file types are currently accessed on different disks.

### **FFrom hh:mm**

Defines the starting time for plotting when either a detailed performance log file is used as source data, or when data from the current redisplay buffer are to be analyzed. The time entered must have exactly this format, i.e. two numbers each for hours and minutes, separated by a colon.

### **FFrom yyyyymmdd**

Defines the starting date (year, month and day) for data analysis when data from a condensed summary data file (HISTSUM, PERFHIST or FCXTREND) are to be analyzed. The date entered must have exactly this format, i.e. four digits for the year, and two digits each for month and day, without any intervening blanks. The format *mmdyy* used in previous releases of the program is also accepted.

Analysis will begin with the first data record found if no 'FROM' date has been entered.

### **TO yyyyymmdd**

Defines the last day for which data are to be included in the analysis when data from a condensed summary data file (HISTSUM, PERFHIST or FCXTREND) are to be analyzed. Analysis will continue to the end of the file if the 'TO' argument is omitted.

### **Days dayspec**

Defines a selection of weekdays for which data are to be analyzed, where *dayspec* can be a combination of the following strings:

**ALL**

If all weekdays are to contribute to the analysis

**M-F**

If only the normal workdays Monday to Friday are to contribute data

**MON**

If all Mondays are to contribute data

**TUE**

If all Tuesdays are to contribute data

**WED**

If all Wednesdays are to contribute data

**THU**

If all Thursdays are to contribute data

**FRI**

If all Fridays are to contribute data

**SAT**

If all Saturdays are to contribute data

**SUN**

If all Sundays are to contribute data

The argument is valid only when analyzing a condensed summary data file (HISTSUM or PERFHIST) or an FCXTREND file; *all* data will be analyzed when it is omitted.

When specified, only performance data collected for the selected weekdays are used for calculating the averages which will eventually be plotted.

For example, specifying 'DAYS M-F SAT' will cause data to be selected for the weekdays Monday to Saturday, excluding any data for Sundays.

**Hours *h1 h2 ... h8***

Defines a selection of up to eight one-hour intervals for which data are to be analyzed. The argument is valid only when analyzing a condensed summary data file (HISTSUM, PERFHIST or FCXTREND).

When specified, only performance data collected for the selected hours are used for calculating the averages which will eventually be plotted.

For example, specifying 'HOURS 9 10 14 15' will cause data to be selected for the periods of 09:00-10:59 and 14:00-15:59, for each of the days previously selected with the 'FROM' and 'TO' arguments.

All hours for which data are available will be included when the HOURS argument is not specified.

**Profile {*h1 {h2}*}**

Specifies that the accumulated data are to be plotted in the form of a performance 'profile' which shows how system load varied during an average day of the period selected with the 'FROM' and 'TO' arguments. The argument is valid only when analyzing a condensed summary data file (HISTSUM, PERFHIST or FCXTREND).

***h1 and h2***

When specified, define the first and last hours of the day for which data are to be plotted. The default is 0 to 23, i.e. system load will be inserted for all 24 hours of the day.

The PROFILE and HOURS arguments are mutually exclusive.

**Usage**

1. See [“REDISP” on page 269](#) (PERFxxx files), [“REDHIST” on page 268](#) (HISTxxx files) or [Appendix G, “Trend File Variables Selection,” on page 987](#) (systemid FCXTREND files) for a list of performance variables and their description that are available with these files, and see the description of the selected benchmark log display for information on the selectable performance variables (note that only numerical fields with load data can be selected).

2. You must enter the full variable names exactly as shown (no abbreviations allowed), except for the case of the maximum user fields ('Max. CPU', 'Max. Vec', ...) of PERFLOG and PERFHIST files where the first eight characters only need be entered. Do not enter the apostrophes, they have been included only to indicate that two words are required for selecting these values.
3. In addition to these basic performance variables you can also specify the names of any 'user variables' which you previously defined by means of the FC USERVAR subcommand.

## Results

See [“Summary Performance History Display”](#) on page 963 for more information on the specific displays caused by this subcommand, and [Appendix A, “Graphical Displays of Performance History Data,”](#) on page 955 for more general information.

# GRAPHVAR / PLOTVAR

## Purpose

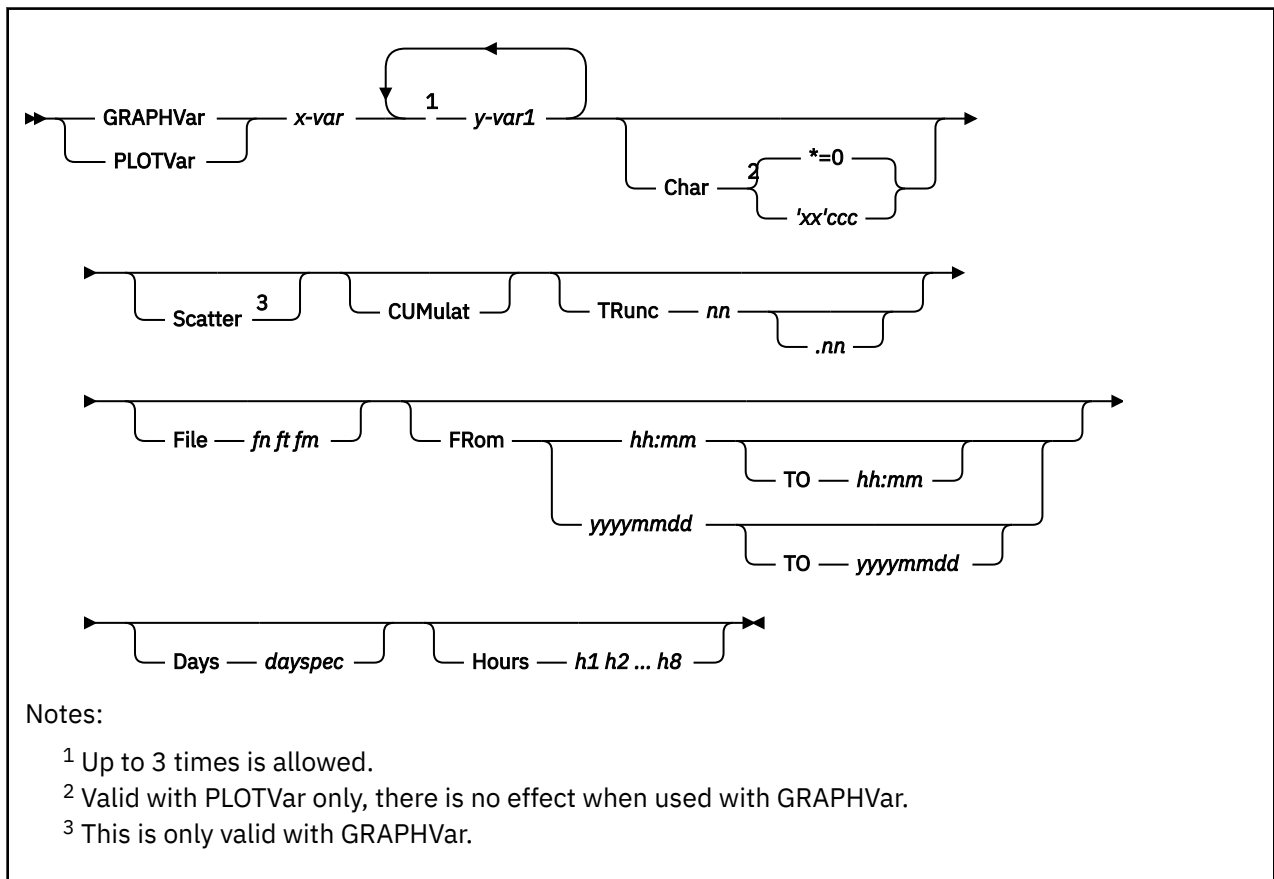
The GRAPHVAR and PLOTVAR subcommands allow selected parts of the performance data accumulated in the monitor redisplay buffer (default) or in one of the history data files on disk to be displayed graphically in the form of variable correlation graphics.

Both subcommands will generate correlation graphics which show how different performance variables depend on each other: any of the performance variables can be selected to be plotted on the horizontal x-axis, and up to three other variables can then be selected as y-variables and their average values plotted for the corresponding value of the x-variable.

Two display modes are possible:

- The GRAPHVAR subcommand requires a display terminal with graphics capability and GDDM services (Graphical Data Display Manager V.2 or 3) for generating and displaying real graphics displays. It cannot operate without these prerequisites.
- The PLOTVAR subcommand uses the same data, but it generates a simple plot using ordinary display characters for the output. It will work on any display terminal, and it has no additional software prerequisites.

## Format



## Parameters

Most of the arguments which can be used with the GRAPHVAR and PLOTVAR subcommands are identical to the ones used with GRAPHSUM and PLOTSUM. Only an abbreviated overview is given for these arguments, please refer to [“GRAPHSUM / PLOTSUM” on page 185](#) for a more detailed description.

**x-var**

Defines the variable whose values are to be plotted on the horizontal x-axis.

**y-var1**

Define the variables whose values are to be plotted on the vertical y-axis.

**Char 'xxxx'cc**

Defines the four plot characters to be used for the selected variables. This argument is intended for use with the PLOTVAR subcommand only, it has no effect when used with GRAPHVAR.

**Scatter**

Specifies that a 'scatter plot' type of graphics is to be generated, showing *all* measured values, and not just grouped averages. The argument is valid only if a single y-variable has been selected, and only for GDDM graphics (GRAPHVAR subcommand).

**CUMulat**

Indicates that the y-variable values are to be plotted cumulatively, i.e. each value added to the previous one. The string 'Cumul .' will be inserted above the y-axis when a cumulative display has been selected.

**TRunc nn{.nn}**

Defines the max. x-scale value to be used for plotting.

The x-scale is automatically adapted to show the majority of the data in a reasonable scale, with the highest few x-values ignored, if necessary. The TRUNC argument is intended for overriding this default logic where it does not yield good results, because some very high x-values would still cause the x-scale to be compressed until the bulk of the really relevant data can no longer be adequately displayed.

The actual scale used will usually still include x-values which are a bit higher than the TRUNC value.

**File fn ft**

Describes the input file to be analyzed, when specified. You can choose from the same files that can also serve as input for the GRAPHSUM subcommand.

*The current redisplay buffer will be used as source if no disk file is specified as input.*

**fm**

Defines the filemode of the source file.

**FRom hh:mm**

Defines the starting time for plotting when either a detailed performance log file is used as source data, or when data from the current redisplay buffer are to be analyzed. The time entered must have exactly this format, i.e. two numbers each for hours and minutes, separated by a colon.

**To hh:mm**

Defines the time stamp of the last record to be used for data analysis. This format is valid only when either a detailed performance log file is used as source data, or when data from the current redisplay buffer are to be analyzed.

**FRom yyyyymmdd**

Defines the starting date (year, month and day) when data from a condensed summary data file (ACUM HISTSUM or ACUM PERFHIST) or from a trend file 'systemid FCXTREND' are to be analyzed.

The date entered must have exactly this format, i.e. four digits for the year, and two digits each for month and day, without any intervening blanks. The format *mmdyy* used in previous releases of the program is also accepted. Analysis will begin with the first data record found if no 'FROM' date has been entered.

**To yyyyymmdd**

Defines the last day for which data are to be included in the analysis. This format is valid only when data from a condensed summary data file (ACUM HISTSUM or ACUM PERFHIST) or from a trend file 'systemid FCXTREND' are to be analyzed. Analysis will continue to the end of the history file if the 'TO' argument is omitted.

**Days *dayspec***

Defines a selection of weekdays for which data are to be analyzed, where *dayspec* can be a combination of the following strings:

**ALL**

If all weekdays are to contribute to the analysis

**M-F**

If only the normal workdays Monday to Friday are to contribute data

**MON**

If all Mondays are to contribute data

**TUE**

If all Tuesdays are to contribute data

**WED**

If all Wednesdays are to contribute data

**THU**

If all Thursdays are to contribute data

**FRI**

If all Fridays are to contribute data

**SAT**

If all Saturdays are to contribute data

**SUN**

If all Sundays are to contribute data

The argument is valid only when analyzing a condensed summary data file (HISTSUM or PERFHIST) or a trend file '*systemid* FCXTREND'; *all* data will be analyzed when it is omitted. When specified, only performance data collected for the selected weekdays are used for calculating the averages which will eventually be plotted.

For example, specifying '*DAYS M-F SAT*' will cause data to be selected for the weekdays Monday to Saturday, excluding any data for Sundays.

**Hours *h1 h2 ... h8***

Defines a selection of up to eight one-hour intervals for which data are to be analyzed. This argument is valid only when data from a condensed summary data file (ACUM HISTSUM or ACUM PERFHIST) or from a trend file '*systemid* FCXTREND' are to be analyzed.

All hours for which data are available will be included when the HOURS argument is not specified.

Note that the purpose of creating variable correlation plots is usually to show such correlations over a wide range of x-variable values. Restricting analysis to selected hours only is not normally conducive to this result, and the HOURS argument should be used only after careful evaluation (e.g. to eliminate specific periods from analysis where system load is atypical because of some special tasks, such as saves, being performed).

**Results**

See [“Variable Correlation Display with Simple Plot” on page 966](#) and [“Variable Correlation Display Using GDDM Graphics” on page 968](#) for more information on the specific displays caused by this subcommand, and [Appendix A, “Graphical Displays of Performance History Data,” on page 955](#) for more general information.

# GVNIC

## Purpose

The GVNIC subcommand causes the General Virtual Network Device Description Screen (FCX268) to be displayed.

To see performance data for virtual network devices, see [“VNIC” on page 360](#) and [“EVNIC” on page 176](#).

## Format

➤ GVNIC ➤
-----------

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX268, General Virtual Network Device Description Screen – GVNIC” on page 755](#).



# GVSWITCH

## Purpose

The GVSWITCH subcommand causes the General Virtual Switch Description Screen (FCX266) to be displayed.

To view performance data for VSWITCH, see [“VSWITCH” on page 363](#) and [“EVSWITCH” on page 177](#).

## Format

▶▶ GVSWItch ▶◀

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX266, General Virtual Switch Description Screen – GVSWITCH” on page 750](#).

# HIPSOCK

## Purpose

The HIPSOCK subcommand causes the Hipersocket Channel Activity Screen (FCX231) to be displayed.

## Format

►► HIPSock ◄◄

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX231, Hipersocket Channel Activity Screen – HIPSOCK”](#) on page 691.

# HISTDATA

## Purpose

The HISTDATA subcommand causes the History Data Selection Menu (FCX160) to be displayed. This menu will list all (or a subset) of the detailed and summary performance history files created by Performance Toolkit which you have access to. It can be entered not only in MONITOR mode but also in basic mode.

## Format



## Parameters

### *fn*

Is the file name of the performance history file(s).

### *ft*

Is the file type of the performance history file(s). Default file types which will be selected are the general history files FCXTREND, HISTLOG\*, HISTSUM, PERFLOG\*, RMONLOG\* and PERFHIST, and also any object's benchmark log files with default file types.

**Note:** Specific file types can be selected to display a subset of all history files, but you can only select files whose file types match the original naming conventions. For example, if anything other than the appended generation number of log files is changed in the file type, the file will no longer be included in the display.

### *fm*

Is the file mode of the performance history file(s).

## Usage

1. Asterisks can be used in the same way as for the CMS LISTFILE command when specifying the file name, file type and file mode.

## Results

See the "Viewing Performance History Data" section in the [z/VM: Performance Toolkit Guide](#) for more information on the History Data Selection Menu.

# HPALIAS

## Purpose

The HPALIAS subcommand causes the HyperPAV Alias Activity Screen (FCX327) to be displayed.

## Format

➤ HPALIAS ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX327, HyperPAV Alias Activity Screen – HPALIAS”](#) on page 861.

# HPFLOG

## Purpose

The HPFLOG subcommand causes HPF I/O Device Data Log Screen (FCX283) to be displayed for the selected I/O device.

## Format

➤ HPFLog — *devno* ➤

## Parameters

### *devno*

Is the real device number for which information should be displayed on the HPF I/O Device Data Log Screen (FCX283).

## Results

See [“FCX283, HPF I/O Device Performance Log Screen – HPFLOG”](#) on page 782.

# HPSHARE

## Purpose

The HPSHARE subcommand causes the HyperPAV Share Activity Screen (FCX328) to be displayed.

## Format

➤ HPSHARE ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX328, HyperPAV Share Activity Screen – HPSHARE ”](#) on page 863.

# IDLEUSER

## Purpose

The IDLEUSER subcommand causes the Idle Users Screen (FCX238) to be displayed.

## Format

➤ IDLEuser ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX238, Idle Users Screen – IDLEUSER”](#) on page 704.

# INTERIM

## Purpose

The INTERIM subcommand allows many performance screens to be switched from a display of the last intervals or overall average data to the display of 'interim' averages based on all measurements made during the current 'interim' period.

## Format

➤ INTERim ➤

## Parameters

No parameters are allowed or tested for.

See the [“CURRENT” on page 165](#) and [“AVERAGE” on page 133](#) subcommands for an explanation of how to select averages for other periods.



# IOCHANGE

## Purpose

The IOCHANGE subcommand causes the I/O Configuration Changes Screen (FCX185) to be displayed.

## Format

►► IOChange ◄◄

## Parameters

No parameters are allowed or tested for.

## Usage

1. The user identification shown for the 'attach' and 'detach' cases is not provided as part of the I/O event records. It is extracted from the next normal I/O sample record for the device – in other words:
  - The userid insertion may be delayed
  - A valid userid cannot be inserted if the I/O device is attached and detached within the same sample interval.A string of eight question marks will then be shown instead of the virtual machine's user identification.

## Results

See [“FCX185, I/O Configuration Changes Screen – IOCHANGE”](#) on page 590.

# IOPQUEUE

## Purpose

The IOPQUEUE subcommand causes the I/O Priority Queueing Screen (FCX237) to be displayed.

## Format

►► IOPqueue ◄◄

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX237, I/O Priority Queueing Screen – IOPQUEUE”](#) on page 702.

# IOPROCLG

## Purpose

The IOPROCLG subcommand causes the I/O Processor Log Screen (FCX232) to be displayed.

## Format

➤ IOPROCLG ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX232, I/O Processor Log Screen – IOPROCLG”](#) on page 693.

# ISFEACT

## Purpose

The ISFEACT subcommand causes the ISFC End Point activity screen (FCX273) to be displayed.

## Format

➤ ISFEAct ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX273, ISFC End Point Activity Screen – ISFEACT”](#) on page 764.

# ISFECONF

## Purpose

The ISFECONF subcommand causes the ISFC End Point configuration screen (FCX272) to be displayed.

## Format

►► ISFEConf ◄◄

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX272, ISFC End Point Configuration Screen – ISFECONF”](#) on page 761.

# ISFELOG

## Purpose

The ISFELOG *epoint* subcommand causes the ISFC End Point Activity Log Screen (FCX279) to be displayed for a selected ISFC end point.

## Format

➤ ISFELog — *epoint* ➤

## Parameters

### *epoint*

Is the ISFC end point hexadecimal number for which the ISFC End Point Activity Log Screen (FCX279) is to be displayed.

## Results

See [“FCX279, ISFC End Point Activity Log Screen – ISFELOG”](#) on page 775.

# ISFLACT

## Purpose

The ISFLACT subcommand causes the ISFC Logical Link activity screen (FCX274) to be displayed.

## Format

➤ ISFLAct ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX274, ISFC Logical Link Activity Screen – ISFLACT”](#) on page 766.

# ISFLALOG

## Purpose

The ISFLALOG subcommand causes the ISFC Logical Link Activity Log Screen (FCX281) to be displayed.

## Format

►► ISFLALog ◄◄

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX281, ISFC Logical Link Activity By-Time Log Screen – ISFLALOG”](#) on page 777.



# ISFLCONF

## Purpose

The ISFLCONF subcommand causes the ISFC Logical Link configuration screen (FCX275) to be displayed.

## Format

➤ ISFLConf ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX275, ISFC Logical Link Configuration Screen – ISFLCONF”](#) on page 768.

# ISFLLOG

## Purpose

The ISFLLOG subcommand causes the ISFC Logical Link Activity Log Screen (FCX280) to be displayed for a selected ISFC logical link.

## Format

➤ ISFLLog — <i>partner</i> ➤
------------------------------

## Parameters

### *partner*

Is the ISFC logical link identification for which the ISFC Logical Link Activity Log Screen (FCX280) is to be displayed.

## Results

See [“FCX280, ISFC Logical Link Activity Log Screen – ISFLLOG”](#) on page 776.

# LCHANNEL

## Purpose

The LCHANNEL subcommand causes the LPAR Channel Load Screen (FCX161) to be displayed, which, on an LPAR system, shows channel load caused by your own logical partition.

## Format

➤ LChannel ➤

## Parameters

No parameters are allowed or tested for.

## Results


See [“FCX161, LPAR Channel Load Screen – LCHANNEL”](#) on page 540.

# LEFT

## Purpose

The LEFT subcommand moves the display window of the performance monitor re-display screen and of the PLOTDET screen to the left.

## Format



## Parameters

*nn*

Indicates the number of columns by which data are to be moved. The default shift value is 40 for the REDISPLAY screen and 60 measurements (approximately one hour) for the PLOTDET display.

## Usage

1. Pressing PF10 or PF22 will also shift the display window left by the default number of columns.

# LGRDATA

## Purpose

The LGRDATA subcommand causes the LGR Data Screen (FCX286) to be displayed.

## Format

➤ LGRData ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX286, LGR Data Screen – LGRDATA”](#) on page 789.

# LGRELOG

## Purpose

The LGRELOG subcommand causes the LGR Event Log Screen (FCX285) to be displayed.

## Format

➤ LGRElog ➤
-------------

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX285, LGR Event Log Screen – LGRELOG”](#) on page 786.

# LINUX

## Purpose

The LINUX subcommand causes either the Linux Details Selection Screen (FCX224) or the Linux Systems Selection Screen (FCX223) to be displayed, based on whether you specify a *linuxid*.

## Format



## Parameters

### *linuxid*

Is the identification of a Linux system for which performance data are to be displayed. The Linux Details Selection Screen (FCX224) will be shown for the selected Linux system.

The Linux Systems Selection Screen (FCX223) will be shown if you enter the LINUX subcommand with no *linuxid*.

## Usage

1. In the Linux Systems Selection Screen (FCX223), the appearance of highlighted Linux guests is based on the definitions made in the FCONX LINUXUSR file. This does *not* guarantee that the RMF PM interface required for detail data retrieval is available for these Linux guests. To prevent misleading entries, users should include only those RMF PM interface IP addresses of Linux guests whose interfaces are always up and running while the guests are logged on.

## Results

See [“FCX223, Linux Systems Selection Screen – LINUX”](#) on page 674 if you specified LINUX.

See [“FCX224, Linux Details Selection Screen – LINUX”](#) on page 675 if you specified LINUX *linuxid*.

# LINUXMEN

## Purpose

The LINUXMEN subcommand causes the Linux Data Selection Menu (FCX242) to be displayed.

## Format

 LINUXMen 

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX242, Linux Data Selection Menu Screen – LINUXMEN”](#) on page 711.



# LOCATE

## Purpose

The LOCATE subcommand causes DIAG X'90' to be issued for locating the entry point specified. When found, its address will be displayed with message FCXCMD380I.

## Format

►► LOCate — xxxxxxxx ►◄

## Parameters

**xxxxxxx**

Is the name of the CP module's entry point.

# LOCKACT

## Purpose

The LOCKACT command causes the Spin Lock Collision Activity Screen (FCX326) to be displayed.

## Format

►► LOCKact ◄◄

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX326, Spin Lock Collision Activity Screen – LOCKACT”](#) on page 859.

# LOCKLOG

## Purpose

The LOCKLOG command causes the Spin Lock Log Screen (FCX265) to be displayed.

The LOCKLOG report supports monitor data from z/VM 6.3 and earlier. For monitor data from z/VM 6.4 and later, the LOCKLOG command is not available and the LOCKACT command from CPUMENU should be used instead.

## Format

▶▶ LOCKLog ◀◀

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX265, Spin Lock Log Screen – LOCKLOG”](#) on page 748.

# LPAR

## Purpose

The LPAR subcommand causes the LPAR Load Screen (FCX126) to be displayed.

## Format

➤ LPAr ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX126, LPAR Load Screen – LPAR”](#) on page 472.

# LPARCONF

## Purpose

The LPARCONF subcommand causes the LPAR Configuration Log Screen (FCX235) to be displayed. The report displays only information about primary type cores for this logical partition.

## Format

▶▶ LPARConf ▶▶

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX235, LPAR Configuration Log Screen – LPARCONF”](#) on page 699.

# LPARLOG

## Purpose

The LPARLOG command causes the LPAR Load Log Screen (FCX202) to be displayed.

## Format

➤ LPARLog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX202, LPAR Load Log – LPARLOG”](#) on page 632.

# LPARLOGM

## Purpose

The LPARLOGM subcommand causes the Logical Partition Logs Selection Menu (FCX307) to be displayed.

## Format

➤ LPARLOGM ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX307, Logical Partition Logs Menu Screen – LPARLOGM”](#) on page 832.

# LPARMENU

## Purpose

The LPARMENU subcommand causes the LPAR Activity Data Selection Menu (FCX305) to be displayed.

## Format

➤ LPARMenu ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX305, Logical Partition Activity Menu Screen – LPARMENU”](#) on page 828.



# LSHARACT

## Purpose

The LSHARACT subcommand causes the Logical Partition Share Screen (FCX306) to be displayed.

## Format

➤ LSHaract ➤

## Parameters

No parameters are allowed or tested for.

## Results


See [“FCX306, Logical Partition Share Screen – LSHARACT”](#) on page 829.

# LXCPU

## Purpose

The LXCPU subcommand causes either the Linux CPU Utilization Screen (FCX230) or the General Linux CPU Utilization Screen (FCX243) to be displayed, based on whether you specify a *linuxid*.

## Format



## Parameters

### *linuxid*

Is the identification of a Linux system for which the detailed RMF PM data gatherer-based Linux CPU Utilization Screen (FCX230) is to be displayed.

The General Linux CPU Utilization Screen (FCX243) will be shown if you enter the LXCPU subcommand with no *linuxid*.

## Usage

1. Please be aware that Linux internal data collection is not synchronized with CP's data collect activity, i.e. the Linux internal data will usually have been obtained for a different period than that used for collecting other VM data. This may lead to discrepancies between the internal Linux values and the VM activity seen from the outside.
2. Linux systems determine CPU utilization in the form of statistical data, as the number of times a processor was found in a certain state when sampled. This sampling logic does not work correctly for current Linux kernels when running under control of VM where the time during which a guest is 'suspended' (attempting to run, but not being dispatched) is not correctly accounted for, and this can lead to an apparent 'CPU utilization' that may be far too high, especially on CPU constrained systems. See [“FCX112, User Resource Usage – USER” on page 431](#) for accurate real CPU utilization numbers, as seen by VM.
3. In the General Linux CPU Utilization Screen (FCX243), the appearance of highlighted Linux guests is based on the definitions made in the FCONX LINUXUSR file. This does *not* guarantee that the RMF PM interface required for detail data retrieval is available for these Linux guests. To prevent misleading entries, users should include only those RMF PM interface IP addresses of Linux guests whose interfaces are always up and running while the guests are logged on.

## Results

See [“FCX243, General Linux CPU Utilization Screen – LXCPU” on page 712](#) if you specified LXCPU.

See [“FCX230, Linux CPU Utilization Screen – LXCPU” on page 689](#) if you specified LXCPU *linuxid*.

# LXCPULOG

## Purpose

The LXCPULOG subcommand causes the Linux CPU Load Log (FCX246) to be displayed for the selected user.

## Format

➤ LXCPULog — *userid* ➤

## Parameters

### *userid*

Is the user identification of the Linux guest machine for which the Linux CPU Load Log (FCX246) is to be displayed.

## Results

See [“FCX246, Linux CPU Load Log Screen – LXCPULOG”](#) on page 718.

# LXFILSYS

## Purpose

The LXFILSYS subcommand causes the Linux Filesystem Usage Screen (FCX228) to be displayed.

## Format

►► LXFILsys — *linuxid* ◄◄

## Parameters

### *linuxid*

Is the identification of a Linux system for which detailed RMF PM data gatherer based Linux Filesystem Usage Screen (FCX228) is to be displayed.

## Results


See [“FCX228, Linux Filesystem Usage Screen – LXFILSYS”](#) on page 684.

# LXMEM

## Purpose

The LXMEM subcommand causes either the Linux Memory Utilization and Activity Screen (FCX229) or the General Linux Memory Utilization Screen (FCX244) to be displayed, based on whether you specify a *linuxid*.

## Format



## Parameters

### *linuxid*

Is the identification of a Linux system for which the detailed RMF PM data gatherer-based Linux Memory Utilization and Activity Screen (FCX229) is to be displayed.

The General Linux Memory Utilization Screen (FCX244) will be shown if you enter the LXMEM subcommand with no *linuxid*.

## Usage

1. Please be aware that Linux internal data collection is not synchronized with CP's data collect activity, i.e. the Linux internal data will usually have been obtained for a different period than that used for collecting other VM data. This may lead to discrepancies between the internal Linux values and the VM activity seen from the outside.
2. In the General Linux Memory Utilization Screen (FCX244), the appearance of highlighted Linux guests is based on the definitions made in the FCONX LINUXUSR file. This does *not* guarantee that the RMF PM interface required for detail data retrieval is available for these Linux guests. To prevent misleading entries, users should include only those RMF PM interface IP addresses of Linux guests whose interfaces are always up and running while the guests are logged on.

## Results

See [“FCX244, General Linux Memory Utilization Screen – LXMEM”](#) on page 714 if specified LXMEM.

See [“FCX229, Linux Memory Utilization and Activity Screen – LXMEM”](#) on page 686 if specified LXMEM *linuxid*.

# LXMEMLOG

## Purpose

The LXMEMLOG subcommand causes the Linux Memory Utilization Log Screen (FCX247) to be displayed for the selected user.

## Format

➤ LXMEMLog — *userid* ➤

## Parameters

### *userid*

Is the user identification of the Linux guest machine for which the Linux Memory Utilization Log Screen (FCX247) is to be displayed.

## Results

See [“FCX247, Linux Memory Utilization Log Screen – LXMEMLOG”](#) on page 719.

# LXNETLOG

## Purpose

The LXNETLOG subcommand causes the Linux Network Activity Log Screen (FCX248) to be displayed for the selected user.

## Format

➤ LXNETLog — *userid* ➤

## Parameters

### *userid*

Is the user identification of the Linux guest machine for which the summary network activity log is to be displayed.

## Results


See [“FCX248, Linux Network Activity Log Screen – LXNETLOG”](#) on page 720.

# LXNETWRK

## Purpose

The LXNETWRK subcommand causes either the Linux Network Activity Screen (FCX227) or the General Linux Networking Activity Screen (FCX245) to be displayed, based on whether you specify a *linuxid*.

## Format



## Parameters

### *linuxid*

Is the identification of a Linux system for which the detailed RMF PM data gatherer-based Linux Network Activity Screen (FCX227) is to be displayed.

The General Linux Networking Activity Screen (FCX245) will be shown if you enter the LXNETWRK subcommand with no *linuxid*.

## Usage

- Please be aware that Linux internal data collection is not synchronized with CP's data collect activity, i.e. the Linux internal data will usually have been obtained for a different period than that used for collecting other z/VM data. This may lead to discrepancies between the internal Linux values and the z/VM activity seen from the outside.
- In the General Linux Networking Activity Screen (FCX245), the appearance of highlighted Linux guests is based on the definitions made in the FCONX LINUXUSR file. This does *not* guarantee that the RMF PM interface required for detail data retrieval is available for these Linux guests. To prevent misleading entries, users should include only those RMF PM interface IP addresses of Linux guests whose interfaces are always up and running while the guests are logged on.

## Results

See [“FCX245, General Linux Networking Activity Screen – LXNETWRK”](#) on page 716 if you specified LXNETWRK.

See [“FCX227, Linux Network Activity Screen – LXNETWRK”](#) on page 682 if you specified LXNETWRK *linuxid*.



# MDCACHE

## Purpose

The MDCACHE subcommand causes the Minidisk Cache Log Screen (FCX138) to be displayed.

## Format

➤ MDCache ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX138, Minidisk Cache Log Screen – MDCACHE”](#) on page 498.

# MDCSTOR

## Purpose

The MDCSTOR subcommand causes the Minidisk Cache Storage Log Screen (FCX178) to be displayed.

## Format

➤ MDCStor ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX178, Minidisk Cache Storage Log Screen – MDCSTOR”](#) on page 573.

# MENU

## Purpose

The MENU subcommand causes the Performance Data Selection Menu (FCX124) to be displayed.

## Format

➤ MENU ➤

## Parameters

No parameters are allowed or tested for.

## Results

See the "Operation in Performance Monitor Mode" chapter in the [z/VM: Performance Toolkit Guide](#) for more information on the Performance Data Selection Menu.

# MONDATA

## Purpose

The MONDATA subcommand causes the Monitor Data Statistics Screen (FCX155) to be displayed.

## Format

➤ MONData ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX155, Monitor Data Statistics Screen – MONDATA” on page 530](#).

# MONSET

## Purpose

The MONSET subcommand causes the Monitor Settings Screen (FCX149) to be displayed.

## Format

➤ MONSEt ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX149, Monitor Settings – MONSET” on page 515](#).

## MTCORELG

### Purpose

The MTCORELG subcommand causes the SMT Metrics Per Core Log Screen (FCX332) to be displayed.

### Format

➤ MTCORElg ➤

### Parameters

No parameters are allowed or tested for.

### Results

See [“FCX332, SMT Metrics Per Core Log Screen – MTCORELG”](#) on page 873.

# MTCTYPLG

## Purpose

The MTCTYPLG subcommand causes the SMT Metrics Per Core Type Log Screen (FCX331) to be displayed.

## Format

➤ MTCTYPlg ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX331, SMT Metrics Per Core Type Log Screen – MTCTYPLG”](#) on page 871.

# MTUSERS

## Purpose

The MTUSERS subcommand causes the Multitasking Users Screen (FCX188) to be displayed.

## Format

➤ MTUsers ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX188, Multitasking Users Screen – MTUSERS” on page 596](#).



# MTUSRLOG

## Purpose

The MTUSRLOG subcommand causes the Multitasking User Log Screen (FCX189) to be displayed for the selected user.

## Format

**➤** MTUSRLog — *userid* **➤**

## Parameters

### *userid*

Is the user identification of the virtual machine for which the Multitasking User Log Screen (FCX189) is to be displayed.

## Results

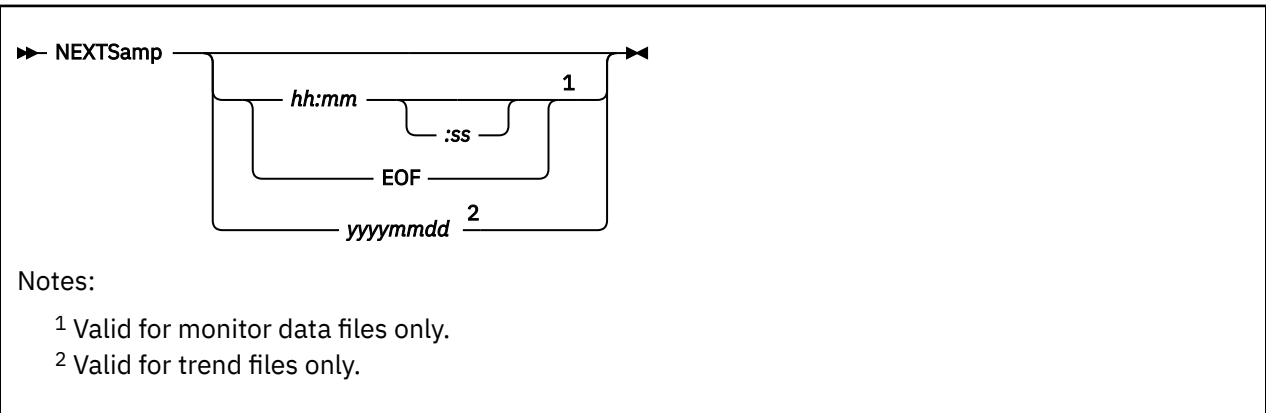
See [“FCX189, Multitasking User Log Screen – MTUSRLOG”](#) on page 598.

# NEXTSAMP

## Purpose

The NEXTSAMP subcommand is valid in monitor and trend file scan mode only. It lets records be processed which correspond to one or multiple sample intervals, up to a time or date specified.

## Format



## Parameters

### *hh:mm:ss*

Specifies the time, in hours and minutes, up to which the CP monitor disk file is to be processed. The time must be specified exactly as shown, i.e. with two digits each for hours and minutes, separated by a colon. Specifying seconds is optional.

### **EOF**

Specifies that the CP monitor disk file is to be processed up to the physical end of file.

### *yyyymmdd*

Specifies the date, in years, months and days, up to which the Performance Toolkit trend file is to be processed. The date must be specified exactly as shown, i.e. with four digits for the year and two digits each for month and days.

## Usage

1. Processing continues until data have been processed which have a time stamp equal to or higher than the specified value. Data for one interval will be processed if no end time or date is specified with the subcommand. For the EOF parameter, processing continues to the end of the CP monitor disk file.

# NSS

## Purpose

The NSS subcommand causes the Shared Segments Screen (FCX133) to be displayed.

## Format

➤ NSS ➤

## Parameters

No parameters are allowed or tested for.

## Usage

1. No information on segment members is available from monitor records; the fields shown are obtained by means of a Q NSS NAME ..... MAP subcommand to tell you what the segment spaces are used for. Because no performance data for these members are available, the member information will *always* be inserted immediately after the output line for the segment space, regardless of the sort sequence.

## Results

See [“FCX133, Shared Segments Screen – NSS” on page 487](#).

# PAGELOG

## Purpose

The PAGELOG subcommand causes the Paging Log Screen (FCX143) to be displayed.

## Format

➤ PAGElog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX143, Paging Log Screen – PAGELOG”](#) on page 502.

# PCIACT

## Purpose

The PCIACT subcommand causes the PCI Function Activity Screen, Format 0 (FCX312), PCI Function Activity Screen, Format 1 (FCX318), PCI Function Activity Screen, Format 2 (FCX320), or PCI Function Activity Screen, Format 3 (FCX322) to be displayed. The particular screen format is selected depending on the PCI function type.

## Format



## Parameters

### ROCE

Display PCI Function Activity Screen, Format 1 (FCX318) for the RoCE PCI function.

### ZEDC

Display PCI Function Activity Screen, Format 2 (FCX320) for the zEDC PCI function.

### ISM

Display PCI Function Activity Screen, Format 3 (FCX322) for the ISM PCI function.

If no parameter is specified, then the PCI Function Activity Screen, Format 0 (FCX312) will be displayed for the unclassified PCI functions (if any).

## Results

See:

- [“FCX312, PCI Function Activity Screen, Format 0 – PCIACT” on page 840](#)
- [“FCX318, PCI Function Activity Screen, Format 1 – PCIACT ROCE” on page 851](#)
- [“FCX320, PCI Function Activity Screen, Format 2 – PCIACT ZEDC” on page 853](#)
- [“FCX322, PCI Function Activity Screen, Format 3 – PCIACT ISM” on page 855](#)

# PCICONF

## Purpose

The PCICONF subcommand causes the PCI Function Configuration Screen (FCX311) to be displayed.

## Format

➤ PCIconf ➤
-------------

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX311, PCI Function Configuration Screen – PCICONF”](#) on page 838.

# PCILOG

## Purpose

The PCILOG subcommand causes the PCI Function Activity Log Screen, Format 0 (FCX313), PCI Function Activity Log Screen, Format 1 (FCX319), PCI Function Activity Log Screen, Format 2 (FCX321), or PCI Function Activity Log Screen, Format 3 (FCX323) to be displayed for the selected PCI function. The particular screen format is selected depending on the PCI function type.

## Format

➡ PCILOG — *pcifunc* ⬅

## Parameters

### *pcifunc*

is the identification of the PCI Function for which the PCI Function Activity Log Screen, Format 0 (FCX313), PCI Function Activity Log Screen, Format 1 (FCX319), PCI Function Activity Log Screen, Format 2 (FCX321), or PCI Function Activity Log screen, Format 3 (FCX323) is to be displayed.

## Results

See:

- [“FCX313, PCI Function Activity Log Screen, Format 0 – PCILOG” on page 842](#)
- [“FCX319, PCI Function Activity Log Screen, Format 1 – PCILOG” on page 852](#)
- [“FCX321, PCI Function Activity Log Screen, Format 2 – PCILOG” on page 854](#)
- [“FCX323, PCI Function Activity Log Screen, Format 3 – PCILOG” on page 856](#)

# PCIMENU

## Purpose

The PCIMENU subcommand causes the PCI Function Menu Screen (FCX310) to be displayed.

## Format

➤ PCIMenu ➤
-------------

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX310, PCI Function Menu Screen – PCIMENU”](#) on page 837.



# PHYSLOG

## Purpose

The PHYSLOG subcommand causes the Real Core Utilization Log Screen (FCX302) to be displayed.

## Format

➤ PHYSLog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX302, Real Core Utilization Log Screen – PHYSLOG”](#) on page 820.

# PINLOG

## Purpose

The PINLOG subcommand causes the Pinned Storage Log Screen (FCX314) to be displayed.

## Format

➤ PINLog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX314, Pinned Storage Log Screen – PINLOG”](#) on page 843.

# PRCLOG

## Purpose

The PRCLOG subcommand causes the Processor Log Screen (FCX304) to be displayed.

This subcommand can be used only with monitor data when running Performance Toolkit against z/VM 6.3 and later data. For z/VM 6.2 and earlier data, use the PROCLOG subcommand.

## Format

» PRCLog «

## Parameters

No parameters are allowed or tested for.

## Results

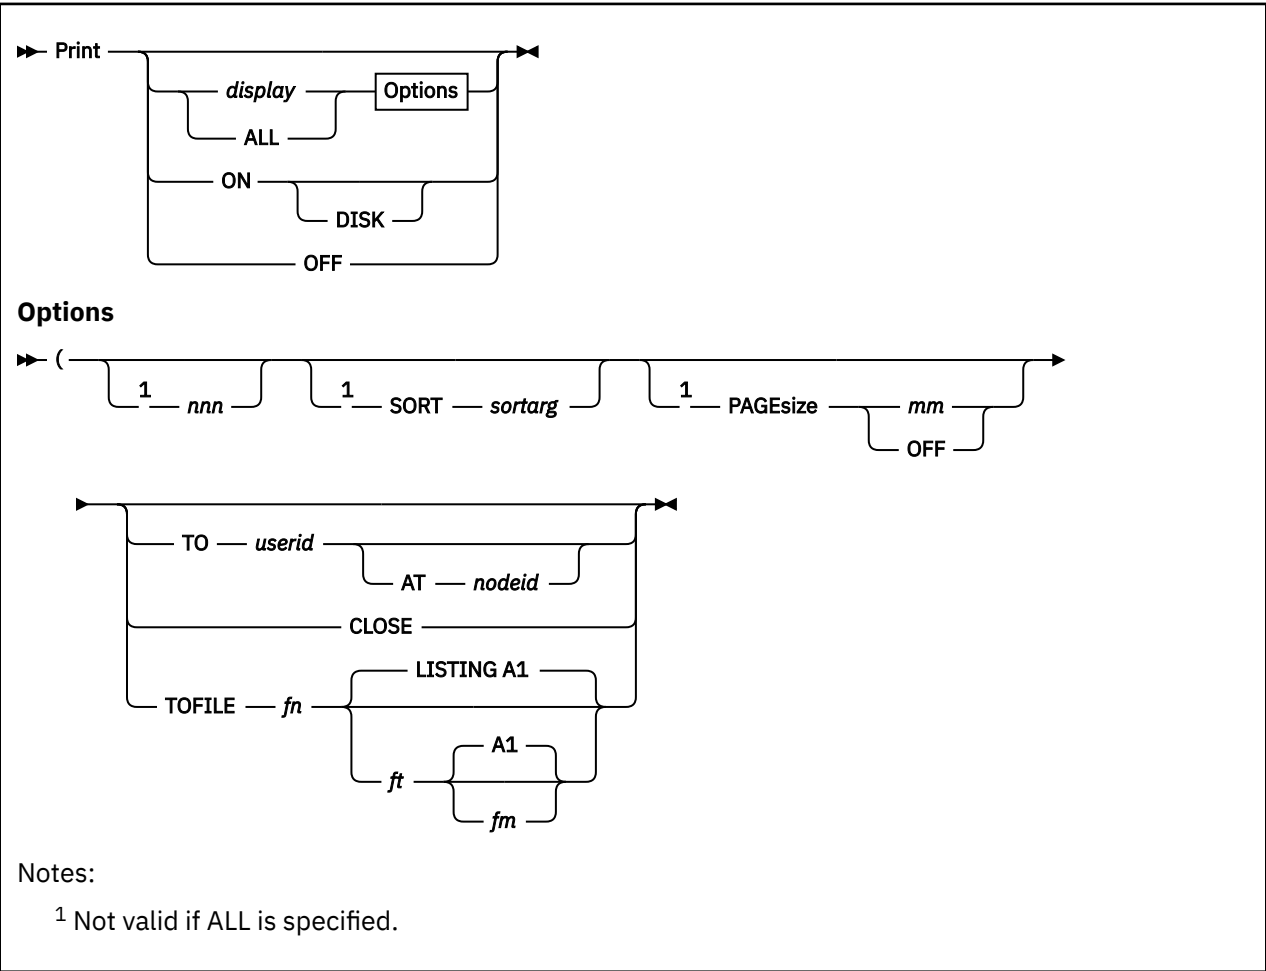
See [“FCX304, Processor Log Screen – PRCLOG”](#) on page 825.

# PRINT

## Purpose

The PRINT subcommand controls the writing of performance monitoring output to a virtual printer or to disk.

## Format



## Parameters

The PRINT subcommand without any arguments causes *all* data to be printed which are available for the *current* display. In the case of the REDISPLAY screen, each display line may extend over up to three print lines (print width limited to 130 characters, regardless of actual output device).

This mode of operation was chosen to allow printing multiple reports in a single print file.

### *display*

Can be any display that is to be printed.

### ALL

Causes a collection of different performance reports to be printed. The reports to be printed must previously be defined in file FCONX REPORTS (see "FCONX REPORTS File" in the [z/VM: Performance Toolkit Guide](#) for a detailed description).

**ON**

Causes all performance data displayed on the screen to be also written to a virtual printer for later reference. Data extending beyond the screen will *not* be included.

This type of printing remains active until a 'PRINT OFF' subcommand is entered, i.e. you will get a complete log of all the data which you looked at on the screen.

**OFF**

Causes continuous printing of monitor output to be stopped, and the virtual printer to be closed. Make sure of the printer output routing before issuing the 'PRINT OFF' subcommand!

**DISK**

Causes the performance output data to be written to disk file 'FCONMON LISTING' instead of virtual printer. A carriage control character is included in the first column of each record, so that the file may still be printed using the 'CC'-option of the CMS 'PRINT' subcommand.

Once a 'PRINT ON DISK' subcommand has been entered, *all* monitor print output will be written to disk until a 'PRINT OFF' subcommand is received.

Options:

***nnn***

Can be used to define the maximum number of print lines *nnn* to be created. It is valid only for limiting the number of output lines of actual performance data displays. This option is useful especially to reduce the output size of the USER or DEVICE displays on large systems. (Make sure the lines are SORTed so that the meaningful information is contained in the first part of the display.)

**Note:** This option is supported for single displays only, i.e. it cannot be used if ALL is specified.

**SORT *sortarg***

Indicates that the report is to be sorted as defined by the sort argument *sortarg*. The SORT option is valid only for reports that also support sorting in monitor mode, i.e. which are supported by the SORT subcommand. The same sort arguments can be used as with the SORT subcommand (see [“SORT”](#) on page 295 for a detailed description).

**Note:** This option is supported for single displays only, i.e. it cannot be used if ALL is specified.

**PAGEsize**

Defines the number of lines to be printed consecutively before a skip to a new page is forced, and the header lines printed again at the top of the page.

***mm***

The maximum number of lines to printed.

**OFF**

Forces continuous printing, without insertion of intermediate header lines even if the print output extends over several pages.

**Note:** This option is supported for single displays only, i.e. it cannot be used if ALL is specified.

**TO *userid* {AT *nodeid*}**

Sets the destination address for the print output, where

***userid***

Is the user identification of the target virtual machine where the output is to be sent

**AT *nodeid***

Specifies the node identification *nodeid* of the system where the destination user *userid* is located.

The destination ID is assumed to be defined on the same VM system if the 'AT *nodeid*' specification is omitted.

The print file will be closed automatically when a specific destination ID has been entered with the 'TO' option, i.e. the CLOSE option is implied and needs not be entered separately.

**CLOSE**

Lets the print file be closed when complete.

**TOFILE *fn* {*ft* {*fm*}}**

Indicates that the 'print' file is to be written to disk, where

***fn***

Is the file name of the disk file

***ft***

Is the file type of the disk file. A file type of LISTING will be used when no explicit file type is entered.

***fm***

Is the file mode of the disk file. A file mode of A1 will be used when no explicit file mode is entered.

**Usage**

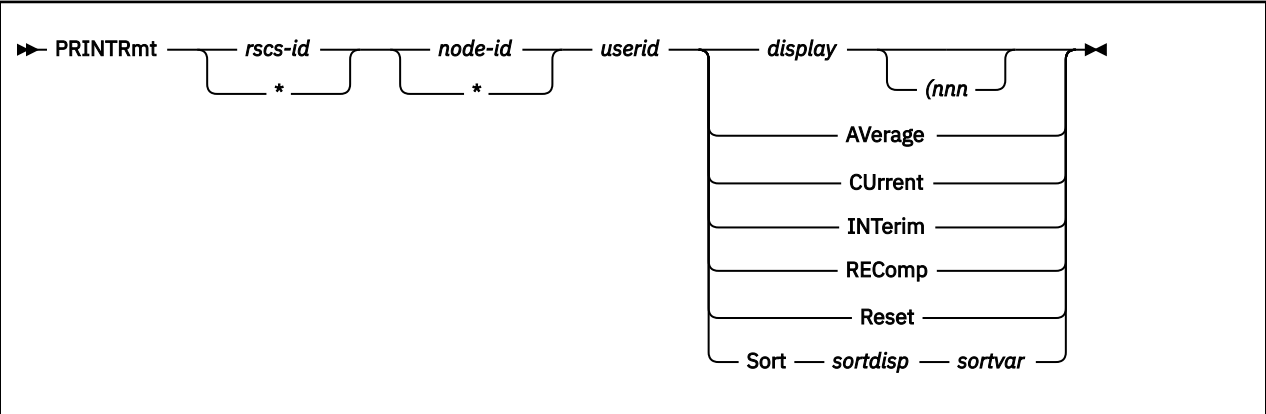
1. The data printed will always be the same as what you see on your screen. Usually you will see data from the last interval unless an 'AVERAGE' subcommand has previously been entered to force calculation of average load figures.
2. The print file will not be closed automatically when the output is directed to the virtual printer. Use any one of the following methods to close the printer when you are done:
  - 'CLOSE' option of the PRINT subcommand
  - 'TO *userid*' option of the PRINT subcommand
  - 'PRINT OFF' subcommand
  - 'CP CLOSE PRT' subcommand
3. REDHist prints one screen's worth of the selected extended history file's data, or the full benchmark log on disk if an object's benchmark log has been selected. The file name, file type and file mode of a history file on disk can be specified with additional arguments *fn ft fm*; the current day's HISTLOG file is used as input if no further arguments are entered. Since these HISTLOG files can be very large, only a small subset - enough to fill a screen - will be printed.
4. REDisp causes all CPU redisplay data in storage to be printed. The file name, file type and file mode of a history file on disk can be specified with additional arguments *fn ft fm*; the performance redisplay data in storage are displayed if no further arguments are entered.

# PRINTRMT

## Purpose

The PRINTRMT subcommand is intended for use with the remote performance data retrieval (FCONRMT mode). It will print the selected performance monitor display data in a special format which includes highlighting attribute information (as strings "'Y' for highlighting, and "'-' for display in normal intensity), and it will send the data to the node and userid specified with the subcommand. Special header and trailer lines will be added which help the receiving program to handle the data correctly.

## Format



## Parameters

The first three arguments are required to define the target user's address where the print output is to be sent.

### *rscs-id*

Is the userid of the RSCS machine where the print output is to be sent. An asterisk '\*' indicates that the PRINTRMT request is for a local user.

### *node-id*

Is the node-ID by which the target user's system is known to the RSCS machine. An asterisk '\*' indicates that the PRINTRMT request is for a local user.

### *userid*

Is the userid of the target user who is to receive the print data.

### *display*

Can be any allowable type of performance data which is to be sent to the selected user.

### *(nnn*

Can be used to define the maximum number of print lines *nnn* to be created. It is valid only for limiting the number of output lines of actual performance data displays. This option is useful especially to reduce the output size of the USER or DEVICE displays on large systems. (You obviously must have SORTed the data previously so that the meaningful information is contained in the first part of the display.)

The remaining arguments allow some additional control over the target machine:

### **Average**

Switches the performance monitor to calculate and display *average* load for the overall reporting period.

### **CUrrent**

Switches the performance monitor to calculate and display *current* load at the end of the current measuring interval.

**INTerim**

Switches the performance monitor to calculate and display *average* load for the current interim period.

**REComp**

Collects new data and recomputes the performance data for all the permanently updated general screens which are based on data extraction from CP control blocks.

Note that collection of MONITOR SAMPLE data can not be forced out of turn in this way. You will have to wait for the end of the current monitor sample interval before updated information can be shown for displays which are based on the analysis of CP MONITOR data.

**Reset**

Resets the cumulative counters to zero, and sets the current time as new start time for calculating averages.

**Sort *sortdispsortvar***

Allows changing the current sorting sequence for a selected performance display. Note that the extended format of the SORT subcommand must be used, where the first argument defines the display and the second argument the variable which is to be used for sorting. See [“SORT” on page 295](#) for a detailed description of the subcommand.

**Usage**

1. REDHist prints one screen's worth of a selected extended history file's data, or the full benchmark log on disk if an object's benchmark log has been selected. The file name, file type and file mode of a history file on disk can be specified with additional arguments *fn ft fm*; the current day's HISTLOG file is used as input if no further arguments are entered. Since HISTLOG files can be very large, only a small subset - enough to fill a screen - will be printed.
2. REDisp causes all CPU redisplay data in storage to be printed. The file name, file type and file mode of a history file on disk can be specified with additional arguments *fn ft fm*; the performance redisplay data in storage are displayed if no further arguments are entered.



# PRIVOP

## Purpose

The PRIVOP subcommand causes the Privileged Operations Screen (FCX104) to be displayed, with information on the utilization of some privileged commands and diagnose instructions.

## Format

➤ PRIVop ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX104, Privileged Operations Screen – PRIVOP”](#) on page 400.

# PROCCONF

## Purpose

The PROCCONF subcommand causes the Processor/Core Configuration Log Screen (FCX234) to be displayed.

## Format

➤ PROCCONF ➤
--------------

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX234, Processor/Core Configuration Log Screen – PROCCONF”](#) on page 697.

# PROCLOG

## Purpose

The PROCLOG subcommand causes the Processor Log Screen (FCX144) to be displayed. This subcommand supports the monitor data from z/VM 6.2 and earlier. For z/VM 6.3 and later, this subcommand is available for compatibility purposes only, and PRCLOG should be used instead. (See [“PRCLOG” on page 251.](#))

## Format

➤ PROCLog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX144, Processor Log Screen – PROCLOG” on page 504.](#)

# PROCMENU

## Purpose

The PROCMENU subcommand causes the Processor Load and Configuration Logs Menu (FCX236) to be displayed.

## Format

➤ PROCmenu ➤
--------------

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX236, Processor Load and Configuration Logs Menu – PROCMENU”](#) on page 701.

# PROCSUM

## Purpose

The PROCSUM subcommand causes the Processor Summary Log Screen (FCX239) to be displayed.

## Format

➤ PROCSum ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX239, Processor Summary Log Screen – PROCSUM”](#) on page 705.

# PUCFGLOG

## Purpose

The PUCFGLOG subcommand causes the Logical PU (Core and Threads) Configuration Log screen (FCX299) to be displayed.

## Format

➤ PUCfglog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX299, Processor Unit \(Core and Threads\) Configuration Log Screen – PUCFGLOG”](#) on page 813.

# PUORGLOG

## Purpose

The PUORGLOG subcommand causes the Logical Core Organization Log (FCX298) to be displayed.

## Format

➤ PUOrglog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX298, Logical Core Organization Log Screen – PUORGLOG”](#) on page 811.

# QDIO

## Purpose

The QDIO command causes the QDIO Activity Screen (FCX251) to be displayed.

## Format

➤ QDIO ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX251, QDIO Activity Screen – QDIO”](#) on page 723.



# QDIOLOG

## Purpose

The QDIOLOG command causes the QDIO Activity Log (FCX252) to be displayed for the selected QDIO device.

## Format

➤ QDIOLog *devno* ➤

## Parameters

### *devno*

Is the I/O device number of the QDIO device for which the QDIO Activity Log (FCX252) is to be displayed.

## Results

See [“FCX252, QDIO Activity Log Screen – QDIOLOG”](#) on page 725.

# QEBSM

## Purpose

The QEBSM command causes the QEBSM Activity Screen (FCX255) to be displayed.

## Format

➤ QEBSM ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX255, QEBSM Activity Screen – QEBSM”](#) on page 730.

# QEBSMLOG

## Purpose

The QEBSMLOG command causes the QEBSM Activity Log (FCX256) to be displayed for the selected device.

## Format

**➤ QEBSMLog — *device* ➤**

## Parameters

### *device*

is the I/O device number of the QEBSM device for which the QEBSM Activity Log (FCX256) is to be displayed.

## Results

See [“FCX256, QEBSM Activity Log Screen – QEBSMLOG”](#) on page 732.

# REDHIST

## Purpose

The REDHIST subcommand allows browsing through history data log files on disk. It can process the extended newer file format (HISTLOG and HISTSUM files), the older format (PERFLOG and PERFHIST files), and an object's benchmark log files on disk.

## Format

```
➡ REDHist — fn — ft — [fm*] ➡
```

## Parameters

### *fn*

Is the file name of the performance history file. The file name usually contains the file's creation date, in *yyyymmdd* format for HISTLOG and HISTSUM files or *mmdyy* format for PERFLOG and PERFHIST files.

### *ft*

Is the file type of the log file. The default file types are HISTLOG (for the current file), HISTLOG1 and HISTLOG2 for previous log generations when an extended history data log is to be browsed, or PERFLOG, PERFLOG1, and PERFLOG2, respectively, for history logs in the older format.

Enter the complete file type when browsing benchmark log files.

The condensed history files have default names of ACUM HISTSUM for the newer file format and ACUM PERFHIST for the older format.

Note that the file type of condensed history files in the older format (default name: ACUM PERFHIST) *must* be PERFHIST when viewing them with the REDHIST subcommand. You can keep backup copies of this file under another file name, but they cannot be displayed with the REDHIST subcommand if the file type was changed too.

### *fm*

Is the file mode of the log file. The default file mode is an asterisk (\*), that is, all accessed disks will be searched for the file.

## Usage


1. The current day's detailed HISTLOG file is selected for display if the subcommand is entered without any arguments.
2. Only actual performance history files should be displayed in this way. Selection of other files might be possible if they pass some tests, but will lead to unpredictable results.
3. An explicit file mode is often specified because another user's performance history file is to be redisplayed. If the current day's history file is to be viewed, the file will probably still be active, and the corresponding disk should be re-accessed just before displaying the file, so that an updated version of the minidisk directory is used. If the file has been updated again, read errors could occur or invalid data could be shown, possibly causing Performance Toolkit to end abnormally.

# REDISP

## Purpose

The REDISP subcommand enters performance monitoring redisplay mode, displaying key system performance indicators from the last measuring intervals.

## Format



## Parameters

### CPUDData

Causes the redisplay data to be positioned so that the CPU load information will be shown. This is the default.

### TRAnsact

Positions the redisplay data so that user activity and queue statistics information will be shown

### STORUtil

Positions the redisplay data so that storage utilization data will be shown

### MAXuser

Positions the redisplay data so that the area with the user performance extremes will be shown

### >nnn

Positions the redisplay data so that the area at a position *nnn* columns to the right will be shown

### *fn*

Is the file name of the performance history file. The file name usually contains the file's creation date, in format *mmdyy* for PERFLOG/PERFHIST files.

### *ft*

Is the file type of the log file. Default file types are PERFLOG, PERFLOG1 and PERFLOG2 respectively for history logs in the old format. The condensed history files have default names of 'ACUM PERFHIST' for the old format.

### *fm*

Is the file mode of the log file. Default file mode is '\*', i.e. all accessed disks will be searched for the file.

## Usage

1. The CPUDATA, TRANSACT, STORUTIL and MAXUSER arguments can also be entered separately, once the performance redisplay screen has been selected, to re-position the data shown on the screen. (The minimum abbreviations shown for these arguments are actually valid when entering them in this way; a shorter format may be accepted when entering them together with the REDISP subcommand.)
2. Performance data from your current redisplay buffer (in storage) will be displayed.

# RESET

## Purpose

The RESET subcommand causes all cumulative counters (used for calculation of average load values) to be reset to zero.

## Format

➤ RESet ➤

## Parameters

No parameters are allowed or tested for.

## Usage

1. An automatic RESET is performed at midnight, but you may choose to reset these counters also at other times, e.g. before periods of exceptionally high system load, so that the averages calculated give a better indication of bottlenecks under high load. See the RESET argument of the 'FC MONCOLL' subcommand ([“FCONTROL MONCOLL” on page 53](#)) for information on how to cause automatic resets at pre-defined times.

For more information, see also the “Resetting Performance Data/Defining Periods” section in the [z/VM: Performance Toolkit Guide](#).

# RESPALL

## Purpose

The RESPALL subcommand causes the Response Time Log (All) Screen (FCX136) to be displayed.

## Format

➤ RESPall ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX136, Response Time Log \(All\) Screen – RESPALL”](#) on page 494.

# RESPUP

## Purpose

The RESPUP subcommand causes the Response Time Log (UP) Screen (FCX137) to be displayed.

## Format

➤ RESPUp ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX137, Response Time Log \(UP\) Screen – RESPUP”](#) on page 496.




# RETURN

## Purpose

The RETURN subcommand, if entered from the performance data selection menu, exits performance monitor mode and returns to basic mode. If entered from any of the general performance screens, the performance data selection menu will be displayed again, and if entered from one of the detailed displays (e.g. user resource details, I/O device details) you will return to the screen from which the detailed display had been selected. Several RETURN subcommands may be necessary to return back to basic mode.

## Format



## Usage


1. 'QUIT' is accepted as a synonym for the RETURN subcommand.
2. The same function can also be performed by pressing PF12.

# RIGHT

## Purpose

The RIGHT subcommand moves the display window of the performance monitor re-display screen or of the PLOTDET display to the right.

## Format



## Parameters

*nn*

Indicates the number of columns by which data are to be moved. The default shift value is 40 for the REDISPLAY screen and 60 measurements (approximately one hour) for the PLOTDET display.

## Usage

1. Pressing PF11 or PF23 will also shift the window to the right for the default number of columns.

# RSKDET

## Purpose

The RSKDet subcommand causes the Reusable Server Kernel Performance Details Screen (FCX217) to be displayed.

## Format

**➤** RSKDet — *userid* **➤**

## Parameters

### *userid*

Is the user identification of the RSK server machine for the Reusable Server Kernel Performance Details Screen (FCX217) is to be displayed.

## Results

See [“FCX217, Reusable Server Kernel Performance Details Screen – RSKDET”](#) on page 660.

# RSKDISK

## Purpose

The RSKDISK subcommand causes the Reusable Server Kernel Disk Performance Log Screen (FCX219) to be displayed.

## Format

➤ RSKDisk — <i>userid</i> ➤
-----------------------------

## Parameters

### *userid*

Is the user identification of the RSK server machine for which the Reusable Server Kernel Disk Performance Log Screen (FCX219) is to be displayed.

## Results

See [“FCX219, Reusable Server Kernel Disk Performance Log Screen – RSKDISK”](#) on page 667.

# RSKLOG

## Purpose

The RSKLOG subcommand causes the Reusable Server Kernel Performance Log Screen (FCX218) to be displayed.

## Format

➤ RSKLog — *userid* ➤

## Parameters

### *userid*

Is the user identification of the RSK server machine for which the Reusable Server Kernel Performance Log Screen (FCX218) is to be displayed.

## Results

See [“FCX218, Reusable Server Kernel Performance Log Screen – RSKLOG”](#) on page 665.

# RSKMENU

## Purpose

The RSKMENU subcommand causes the Reusable Server Kernel Performance Data Menu (FCX216) to be displayed.

## Format

➤ RSKmenu ➤
-------------

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX216, Reusable Server Kernel Performance Data Menu Screen – RSKMENU”](#) on page 659.

# RSKSTOR

## Purpose

The RSKSTOR subcommand causes the Reusable Server Kernel Storage Log Screen (FCX220) to be displayed.

## Format

➤ RSKStor — *userid* ➤

## Parameters

### *userid*

Is the user identification of the RSK server machine for which the Reusable Server Kernel Storage Log Screen (FCX220) is to be displayed.

## Results

See [“FCX220, Reusable Server Kernel Storage Log Screen – RSKSTOR”](#) on page 669.

# SCHDATA

## Purpose

The SCHDATA subcommand causes the Scheduler Data Log Screen (FCX175) to be displayed.

## Format

➤ SCHData ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX175, Scheduler Data Log Screen – SCHDATA”](#) on page 564.



# SCHEDLOG

## Purpose

The SCHEDLOG subcommand causes the Scheduler Queue Log Screen (FCX145) to be displayed.

## Format

➤ SCHEDlog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX145, Scheduler Queue Log Screen – SCHEDLOG” on page 507](#).

# SCSI

## Purpose

The SCSI subcommand causes the SCSI Device Screen (FCX249) for the emulated FBA devices (EDEV) to be displayed. These devices must have been previously defined by the CP SET EDEV command.

## Format

➤ SCSI ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX249, SCSI Device – SCSI”](#) on page 721.

# SCSILOG

## Purpose

The SCSILOG subcommand causes the SCSI Resource Utilization Log Screen (FCX250) to be displayed for the selected emulated FBA device (EDEV). This device must have been previously defined by the CP SET EDEV command.

## Format

▶▶ SCSILog — *devno* ◀◀

## Parameters

### *devno*

Is the emulated FBA device number for which the SCSI Resource Utilization Log Screen (FCX250) is to be displayed.

## Results

See [“FCX250, SCSI Resource Utilization Log – SCSILOG”](#) on page 722.

# SEEKDET

## Purpose

The SEEKDET subcommand causes the DASD Seeks Details Screen (FCX158) to be displayed.

## Format

➤ SEEKDET — *devno* ➤

## Parameters

### *devno*

Specifies the device number of the disk for which the DASD Seeks Details Screen (FCX158) is to be displayed.

## Usage

1. When entered in real-time monitoring mode, a:

MONITOR EVENT ENABLE SEEKS DEVICE xxxx

command is automatically executed for the selected disk. If SEEKS domain data for that disk had not previously been enabled, message:

FCXSED451I SEEKS EVENT data collection started for device xxxx

will be shown to make you aware of the fact. *Note that seeks data collection for the disk will not be disabled again when you leave the display.* You will have to explicitly enter the subcommand:

MONITOR EVENT DISABLE SEEKS DEVICE xxxx

for disabling seeks data collection for a single disk, or the subcommand:

MONITOR EVENT DISABLE SEEKS ALL

for disabling seeks data collection for all disks.

2. The Volume serial and Disc. time (msec) information is obtained from I/O domain SAMPLE records. It can not be inserted before at least one (for Volume serial) or two (for Disc. time (msec)) sample intervals have elapsed.

## Results

See [“FCX158, DASD Seeks Details Screen – SEEKDET” on page 536.](#)

# SEEKDIST

## Purpose

The SEEKDIST subcommand causes the DASD Seeks Distances Screen (FCX157) to be displayed.

## Format

➤ SEEKdist ➤

## Parameters

No parameters are allowed or tested for.

## Usage

1. Seeks data collection and analysis for disk packs with a high I/O activity causes considerable overhead. Seeks EVENT data should, therefore, be enabled only for limited periods of time, and/or for a restricted number of disks.
2. The Volume Serial and Disc Time msec information is obtained from I/O domain SAMPLE records. It can not be inserted before at least one (for Volume Serial) or two (for Disc Time msec) sample intervals have elapsed.

## Results

See [“FCX157, DASD Seeks Distances Screen – SEEKDIST” on page 534.](#)

# SEEKDLOG

## Purpose

The SEEKDLOG subcommand causes the DASD Seeks Distances Log Screen (FCX172) to be displayed for the selected disk.

## Format

➤ SEEKDLog — <i>devno</i> ➤
-----------------------------

## Parameters

### *devno*

Is the number of the disk for which the DASD Seeks Distances Log Screen (FCX172) is to be displayed.

## Results

See [“FCX172, DASD Seeks Distances Log Screen – SEEKDLOG”](#) on page 557.

# SEEKLLOG

## Purpose

The SEEKLLOG subcommand causes the DASD Seeks Locations Log Screen (FCX171) to be displayed for the selected disk.

## Format

```
➤ SEEKLLog — devno ➤
```

## Parameters

### *devno*

Is the number of the disk for which the DASD Seeks Locations Log Screen (FCX171) is to be displayed.

## Results

See [“FCX171, DASD Seeks Locations Log Screen – SEEKLLOG”](#) on page 556.

# SEEKLOC

## Purpose

The SEEKLOC subcommand causes the DASD Seeks Locations Screen (FCX156) to be displayed.

## Format

▶▶ SEEKLoc ◀◀
---------------

## Parameters

No parameters are allowed or tested for.

## Usage

1. Seeks data collection and analysis for disk packs with a high I/O activity causes considerable overhead. Seeks EVENT data should, therefore, be enabled only for limited periods of time, and/or for a restricted number of disks.
2. The Volume Serial and Disc Time msec information is obtained from I/O domain SAMPLE records. It can not be inserted before at least one (for Volume Serial) or two (for Disc Time msec) sample intervals have elapsed.

## Results

See [“FCX156, DASD Seeks Locations Screen – SEEKLOC” on page 532.](#)




# SFS

## Purpose

The SFS subcommand causes either the Shared File System Server Screen (FCX116) or the Shared File System Server Details (FCX117) to be displayed, based on whether a specific *userid* is specified.

## Format



## Parameters

### *userid*

Is the identification of a shared file system server for which detailed performance data are to be displayed. A Shared File System Server Details Screen (FCX117) will be shown with performance information for the selected user only.

The Shared File System Server Screen (FCX116), which shows summary performance data for all Shared File System server machines, will be shown if you enter the SFS subcommand with no *userid*.

## Usage

1. The analysis of SFS server performance is based on both APPLDATA and USER domain records of the CP MONITOR. No data can be shown when the APPLDATA records are missing, but Performance Toolkit will attempt to show as much information as possible even when the user data records are not available. Dots will then be inserted into the fields which could not be calculated. This situation will *always* occur during the very first full interval. Data extraction from the relevant user records will start only once the SFS server machines have been identified by their APPLDATA records, i.e. the fields based on USER records (page wait state, CPU consumption) can be calculated only one cycle later.

## Results

See [“FCX116, Shared File System Server Screen – SFS” on page 455](#) if you specified SFS.

See [“FCX117, Shared File System Server Details Screen – SFS” on page 457](#) if you specified SFS *userid*.

# SFSIOLOG

## Purpose

The SFSIOLOG subcommand causes the Shared File System I/O Log Screen (FCX151) to be displayed.

## Format

►► SFSIolog ◄◄

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX151, Shared File System I/O Log Screen – SFSIOLOG”](#) on page 520.

# SFSLOG

## Purpose

The SFSLOG subcommand causes the Shared File System Log Screen (FCX150) to be displayed.

## Format

➤ SFSLog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX150, Shared File System Log Screen – SFSLOG”](#) on page 518.

# SFSMENU

## Purpose

The SFSMenu subcommand causes the SFS and BFS Logs Selection Menu (FCX221) to be displayed.

## Format

➤ SFSMenu ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX221, SFS and BFS Logs Selection Menu Screen – SFSMENU”](#) on page 671.

# SFSREQ

## Purpose

The SFSREQ subcommand causes the Shared File System Requests Log Screen (FCX152) to be displayed.

## Format

►► SFSReq ◄◄

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX152, Shared File System Requests Log Screen – SFSREQ” on page 522](#).

# SKIPSAMP

## Purpose

The SKIPSAMP subcommand is valid in monitor and trend file scan mode only. It will skip a number of records until a set of records is found which was collected at a specific time or date.

## Format

➡ SKIPSamp

1

hh:mm

:ss

2

yyyymmdd

Notes:

- <sup>1</sup> Valid for monitor data files only.
- <sup>2</sup> Valid for trend files only.

## Parameters

### hh:mm:ss

Specifies the time, in hours and minutes, of the next batch of monitor sample data you are interested in. The time must be specified exactly as shown, i.e. with two digits each for hours and minutes, separated by a colon. Specifying seconds is optional.

### yyyymmdd

Specifies the date, in years, months and days, up to which the intermediate trend file records are to be skipped. The date must be specified exactly as shown, i.e. with four digits for the year and two digits each for month and days.

## Usage

1. The file will be scanned until sample data are found which have been collected at a time equal to or higher than the specified value. If entered without a target time or date, only the next interval will be processed, just as if a NEXTSAMP subcommand had been entered. Note that the intermediate data will *not* be processed. The subcommand is intended to allow skipping to a specific point of time in the collected data with a minimum of overhead. Use the NEXTSAMP subcommand for skipping if all intermediate data should also be processed, thus creating a complete set of 'by time' log displays, and also allowing correct averages to be calculated for monitor data.

# **SORT**

## **Purpose**

The SORT subcommand causes output lines on certain screens to be sorted according to the selected criteria. Sorting is possible only on data columns that have a dot above the column heading. Data columns on screens with no dots above any of the headings cannot be selected for sorting.

## **Format**



## **Parameters**

### ***display***

Specifies the display that is to be sorted. The argument is required only for setting the sorting sequence on a display other than the one that is currently being viewed. It is identical to the subcommand required for selecting the targeted display for viewing.

The data column that is selected for sorting is indicated by the underscore characters above the column heading.

### ***colname***


Specifies the data column that is to be sorted. Enter the character string in the data column's bottom heading line as the sort argument.

# SPOOL

## Purpose

The SPOOL subcommand causes the SPOOL File User Details Screen (FCX121), SPOOL File User Summary Screen (FCX120), or the SPOOL File Summary Screen (FCX119) to be displayed, based on how you specify the subcommand. (See below.)

## Format



## Parameters

### *userid*

Displays the SPOOL File User Details Screen (FCX121) for the specified *userid*.

### User

(Without further argument) displays the SPOOL File User Summary Screen (FCX120), with the total SPOOL space utilization for all users.

### Nosystem

Displays the SPOOL File Summary Screen (FCX119) for the 100 largest SPOOL files, but will exclude system files (for example, NSS files) from the display.

If you specify the SPOOL subcommand with no parameters at all, the SPOOL File Summary Screen (FCX119) will be displayed for the 100 largest SPOOL files.

## Usage

**Usage Note:** Scanning a large number of SPOOL files can cause considerable overhead. The display is therefore not automatically updated in the usual 1-minute intervals.

You are shown an updated display in the following instances:

- After initially entering a SPOOL display from any other type of performance display
- After pressing the ENTER-key without input

## Results (For Selected User)

See [“FCX119, SPOOL File Summary Screen – SPOOL or SPOOL NOSYSTEM”](#) on page 461 if specified SPOOL or SPOOL NOSYSTEM.

See [“FCX120, SPOOL File User Summary Screen – SPOOL USER”](#) on page 464 if specified SPOOL USER.

See [“FCX121, SPOOL File User Details Screen – SPOOL”](#) on page 466 if specified SPOOL *userid*.



# SSICONF

## Purpose

The SSICONF subcommand causes the SSI configuration screen (FCX276) to be displayed.

## Format

➤ SSIconf ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX276, SSI Configuration Screen – SSICONF”](#) on page 770.

# SSIMENU

## Purpose

The SSIMenu subcommand causes the SSI and ISFC Selection Menu (FCX271) to be displayed.

## Format

►► SSIMenu ◄◄

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX271, SSI Data Menu Screen – SSIMENU”](#) on page 760.

# SSISCHLG

## Purpose

The SSISCHLG subcommand causes the SSI State Change Synchronization Activity log screen (FCX277) to be displayed.

## Format

➤ SSISChlg ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX277, SSI State Change Synchronization Activity Log Screen – SSISCHLG”](#) on page 772.

## SSISMILG

### Purpose

The SSISMILG subcommand causes the SSI State/Mode Information log screen (FCX278) to be displayed.

### Format

►► SSISMilg ◄◄

### Parameters

No parameters are allowed or tested for.

### Results

See [“FCX278, SSI State/Mode Information Log Screen – SSISMILG”](#) on page 773.

# STEALLOG

## Purpose

The STEALLOG subcommand causes the Steal Statistics screen (FCX296) to be displayed.

## Format

➤ STEALlog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX296, Steal Statistics Screen – STEALLOG” on page 807](#).

# STORAGE

## Purpose

The STORAGE subcommand causes the Storage Utilization Screen (FCX103) to be displayed.

## Format

►► STOrage ◄◄

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX103, Storage Utilization – STORAGE”](#) on page 392.

# STORLOG

## Purpose

The STORLOG command causes the Storage Utilization Log Screen (FCX253) to be displayed.

## Format

➤ STORLog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX253, Storage Utilization Log Screen – STORLOG”](#) on page 726.

# STORMENU

## Purpose

The STORMENU command displays the Storage Management Logs Menu (FCX260).

## Format

➤ STORMenu ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX260, Storage Management Logs Menu Screen – STORMENU” on page 739](#).



## SUBPLOG

### Purpose

The SUBPLOG subcommand causes the Subpool Storage Log Screen (FCX316) to be displayed.

### Format

➤ SUBPlog ➤

### Parameters

No parameters are allowed or tested for.

### Results

See [“FCX316, Subpool Storage Log Screen – SUBPLOG”](#) on page 848.

# SXSAVAIL

## Purpose

The SXSAVAIL command causes the SXS Available Page Queues Management Log Screen (FCX261) to be displayed.

## Format

➤ SXSAvail ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX261, SXS Available Page Queues Management Log Screen – SXSAVAIL”](#) on page 740.

# SXSDEFER

## Purpose

The SXSDEFER command causes the SXS Deferred Tasks Management Log Screen (FCX263) to be displayed.

## Format

➤ SXSDefer ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX263, SXS Deferred Tasks Management Log Screen – SXSDEFER”](#) on page 744.

# SXSPAGE

## Purpose

The SXSPAGE command causes the SXS Page Management Log Screen (FCX262) to be displayed.

## Format

►► SXSPage ◄◄

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX262, SXS Page Management Log Screen – SXSPAGE”](#) on page 742.

# SXSUTIL

## Purpose

The SXSUTIL command causes the SXS Storage Utilization Log Screen (FCX264) to be displayed.

## Format

➤ SXSUtil ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX264, SXS Storage Utilization Log Screen – SXSUTIL”](#) on page 746.

# **SYSCONF**

## **Purpose**

The SYSCONF subcommand causes the System Configuration Screen (FCX180) screen to be displayed.

## **Format**

▶▶ SYSConf ◀◀
---------------

## **Parameters**

No parameters are allowed or tested for.

## **Results**

See [“FCX180, System Configuration – SYSCONF”](#) on page 577.

# SYSLOAD

## Purpose

The SYSLOAD subcommand causes the System Load Overview Screen (FCX198) to be displayed when entered while in an APPC/VM session via an APPC/VM store and forward server that is also acting as central data collector machine.

This subcommand is not valid in native monitor mode. Note that the System Load Overview Screen (FCX198) can also be displayed by issuing the FCONRMT command in basic mode. See the "Remote Performance Monitoring Facility" section in the [“FCONRMT” on page 106](#) for a detailed description of remote performance monitoring.

## Format

➤ SYSLOAD ➤

## Parameters

No parameters are allowed or tested for.

## Results

See the "Connecting to an S&F Server for Remote Connections" section in the [z/VM: Performance Toolkit Guide](#) for more information on the SYSMENU System Selection Menu and the SYSLOAD System Load Overview Screen.

# SYSLOG

## Purpose

The SYSLOG subcommand causes the System Facilities Log Screen (FCX179) to be displayed.

## Format

➤ SYSLog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX179, System Facilities Log Screen – SYSLOG”](#) on page 575.



# SYSMENU

## Purpose

The SYSMENU subcommand causes the System Selection Menu (FCX193) be displayed when entered while in an APPC/VM session via an APPC/VM store and forward server.

This subcommand is not valid in native monitor mode.

## Format

➤ SYSMENU ➤
-------------

## Parameters

No parameters are allowed or tested for.

## Results

See the "Connecting to an S&F Server for Remote Connections" section in the [\*z/VM: Performance Toolkit Guide\*](#) for more information on the SYSMENU System Selection Menu and the SYSLOAD System Load Overview Screen.

# SYSSET

## Purpose

The SYSSET subcommand causes the System Settings Screen (FCX154) to be displayed.

## Format

➤ SYSSet ➤
------------

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX154, System Settings Screen – SYSSET”](#) on page 526.

# SYSSUMLG

## Purpose

The SYSSUMLG subcommand causes the System Summary Log Screen (FCX225) to be displayed.

## Format

➤ SYSSUmlg ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX225, System Summary Log Screen – SYSSUMLG”](#) on page 677.

# SYSTEM

## Purpose

The SYSTEM subcommand causes the System Counters Screen (FCX102) to be displayed, with information on the system's management of storage, paging, and the minidisk cache.

## Format

►► SYStem ◄◄

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX102, System Counters – SYSTEM”](#) on page 385.

# SYSTRANS

## Purpose

The SYSTRANS subcommand causes the System Transaction Statistics Screen (FCX130) to be displayed.

## Format

►► SYSTRans ◄◄

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX130, System Counters Screen – SYSTRANS”](#) on page 480.

# TCPACTLG

## Purpose

The TCPACTLG subcommand causes the General TCP/IP Activity Log Screen (FCX204) to be displayed.

## Format

► TCPActlg — *userid* ◄

## Parameters

### *userid*

Is the user identification of the TCP/IP server machine for which the general activity log is to be displayed.

## Results

See [“FCX204, General TCP/IP Activity Log Screen – TCPACTLG”](#) on page 638.

# TCPBPLOG

## Purpose

The TCPBPLOG subcommand causes the TCP/IP Buffer Pools Log Screen (FCX210) to be displayed.

## Format

```
➤ TCPBplog — userid ➤
```

## Parameters

### *userid*

Is the user identification of the TCP/IP server machine for which the TCP/IP Buffer Pools Log Screen (FCX210) is to be displayed.

## Results

See [“FCX210, TCP/IP Buffer Pools Log Screen – TCPBPLOG”](#) on page 648.

# TCPCONF

## Purpose

The TCPCONF subcommand causes the TCP/IP Server Configuration Screen (FCX212) to be displayed.

## Format

➤ TCPConf — *userid* ➤

## Parameters

### *userid*

Is the user identification of the TCP/IP server machine for which the TCP/IP Server Configuration Screen (FCX212) is to be displayed.

## Results

See [“FCX212, TCP/IP Server Configuration Screen – TCPCONF”](#) on page 650.



# TCPDATLG

## Purpose

The TCPDATLG subcommand causes the General TCP/IP Data Transfer Log Screen (FCX205) to be displayed.

## Format

➤ TCPDatlg — <i>userid</i> ➤
------------------------------

## Parameters

### *userid*

Is the user identification of the TCP/IP server machine for which the General TCP/IP Data Transfer Log Screen (FCX205) is to be displayed.

## Results

See [“FCX205, General TCP/IP Data Transfer Log Screen – TCPDATLG”](#) on page 640.

# TCPDOSLG

## Purpose

The TCPDOSLG command lets the TCP/IP Denial of Service Log Screen (FCX233) be displayed.

## Format

```
➤ TCPDoslg — userid ➤
```

## Parameters

### *userid*

Is the user identification of the TCP/IP server machine for which the TCP/IP Denial of Service Log Screen (FCX233) is to be displayed.

## Results

See [“FCX233, TCP/IP Denial of Service Log Screen – TCPDOSLG”](#) on page 695.

# TCPICMP

## Purpose

The TCPICMP subcommand causes the TCP/IP ICMP Messages Log Screen (FCX206) to be displayed.

## Format

➤ TCPICmp — *userid* ➤

## Parameters

### *userid*

Is the user identification of the TCP/IP server machine for which the TCP/IP ICMP Messages Log Screen (FCX206) is to be displayed.

## Results

See [“FCX206, TCP/IP ICMP Messages Log Screen – TCPICMP”](#) on page 642.

# TCPIOLOG

## Purpose

The TCPIOLOG command lets the TCP/IP I/O Activity Log Screen (FCX222) be displayed.

## Format

➤ TCPIOlog — *userid* ➤

## Parameters

### **userid**

Is the user identification of the TCP/IP server machine for which the TCP/IP I/O Activity Log Screen (FCX222) is to be displayed.

## Results

See [“FCX222, TCP/IP I/O Activity Log Screen – TCPIOLOG”](#) on page 672.

# TCPLINKS

## Purpose

The TCPLINKS subcommand causes the TCP/IP Links Activity Log Screen (FCX208) to be displayed.

## Format

`➤ TCPLinks — userid ➤`

## Parameters

### *userid*

Is the user identification of the TCP/IP server machine for which the TCP/IP Links Activity Log Screen (FCX208) is to be displayed.

## Results

See [“FCX208, TCP/IP Links Activity Log Screen – TCPLINKS”](#) on page 646.

# TCPMENU

## Purpose

The TCPMENU subcommand causes the TCP/IP Performance Data Menu (FCX203) to be displayed.

## Format

➤ TCPMenu ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX203, TCP/IP Performance Data Menu Screen – TCPMENU” on page 636](#).

# TCPSESS

## Purpose

The TCPSESS subcommand causes the TCP/IP TCP and UDP Sessions Log Screen (FCX207) to be displayed.

## Format

➤ TCPSESS — *userid* ➤

## Parameters

### *userid*

Is the user identification of the TCP/IP server machine for which the TCP/IP TCP and UDP Sessions Log Screen (FCX207) is to be displayed.

## Results

See [“FCX207, TCP/IP TCP and UDP Sessions Log Screen – TCPSESS”](#) on page 644.

# TCPUSERS

## Purpose

The TCPUSERS subcommand causes the TCP/IP Users Activity Log Screen (FCX214) to be displayed.

## Format

➤ TCPUsers — *userid* ➤

## Parameters

### *userid*

Is the user identification of the TCP/IP server machine for which the TCP/IP Users Activity Log Screen (FCX214) is to be displayed.

## Usage

1. Once a virtual machine has shown any TCP/IP activity, its ID will be shown in the log for each interval at least up to the next 'reset', even while inactive. At reset time users that did not show any activity (i.e. did not produce any OPEN/CLOSE records) during the previous reset interval and that do not have an open session will be removed from the internal tables, and they will no longer be included in the log.

## Results

See [“FCX214, TCP/IP TCP and UDP Sessions Log Screen – TCPUSERS”](#) on page 655.



## TOP

### Purpose

The TOP subcommand displays the beginning of the current selection of performance data.

### Format

➤ Top ➤

### Parameters

No parameters are allowed or tested for.

### Usage

1. The beginning of the log will also be shown after pressing PF4 or PF16.

# TOPOLOG

## Purpose

The TOPOLOG subcommand causes the System Topology Machine Organization Screen (FCX287) to be displayed.

## Format

➤ TOPOlog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX287, System Topology Machine Organization Screen – TOPOLOG”](#) on page 792.

# UCOMM

## Purpose

The UCOMM subcommand causes the User Communication Screen (FCX132) to be displayed.

## Format

➤ UCOmm ➤
-----------

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX132, User Communication Screen – UCOMM”](#) on page 484.

# UCOMMLOG

## Purpose

The UCOMMLOG subcommand causes either the User Communication Log Screen (FCX167) or the General User Communication Log Screen (FCX182) to be displayed, based on whether a specific *userid* is specified.

## Format



## Parameters

### *userid*

Is the user identification of the virtual machine for which the User Communication Log Screen (FCX167) is to be displayed.

The General User Communication Log Screen (FCX182), with sums and averages for all users on the system, will be shown when the subcommand is entered without an additional argument.

## Results

See [“FCX182, General User Communication Log Screen – UCOMMLOG” on page 585](#) if you specified UCOMMLOG.

See [“FCX167, User Communication Log Screen – UCOMMLOG” on page 547](#) if you specified UCOMMLOG *userid*.

# UCONF

## Purpose

The UCONF subcommand causes the User Configuration Screen (FCX226) to be displayed.

## Format

➤ UCONF ➤
-----------

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX226, User Configuration – UCONF”](#) on page 679.

# UDEFMENU

## Purpose

The UDEFMENU subcommand causes the User Defined Displays Selection Menu (FCX213) to be displayed.

## Format

➤ UDEFMenu ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX213, User-Defined Displays Selection Menu Screen – UDEFMENU”](#) on page 654.

# UPAGE

## Purpose

The UPAGE subcommand causes the User Page Data Screen (FCX113) to be displayed.

## Format

➤ UPAGE ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX113, User Page Data Screen – UPAGE”](#) on page 435.

# UPAGELOG

## Purpose

The UPAGELOG subcommand causes the User Page Data Log Screen (FCX163) to be displayed for the selected user.

## Format

➤ UPAGELog — *userid* ➤

## Parameters

### *userid*

Is the user identification of the virtual machine for which the User Page Data Log Screen (FCX163) is to be displayed.

## Results

See [“FCX163, User Page Data Log Screen – UPAGELOG”](#) on page 543.



# UPGACT

## Purpose

The UPGACT subcommand causes the User Page Activity screen (FCX290) to be displayed.

## Format

➤ UPGAct ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX290, User Page Activity Screen – UPGACT”](#) on page 798.

# UPGACTLG

## Purpose

The UPGACTLG subcommand causes the User Page Activity Log screen (FCX291) to be displayed for the selected user.

## Format

**➤** UPGACTLg — *userid* **➤**

## Parameters

### *userid*

Is the user identification of the virtual machine for which the User Page Activity Log screen (FCX291) is to be displayed.

## Results

See [“FCX291, User Page Activity Log Screen – UPGACTLG”](#) on page 800.

# UPGMENU

## Purpose

The UPGMENU subcommand causes the User Paging Menu (FCX289) to be displayed.

## Format

➤ UPGMenu ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX289, User Paging Menu Screen – UPGMENU” on page 797](#).

# UPGUTL

## Purpose

The UPGUTL subcommand causes the User Page Utilization Data screen (FCX292) to be displayed.

## Format

➤ UPGUtl ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX292, User Page Utilization Data Screen – UPGUTL”](#) on page 801.

# UPGUTLLG

## Purpose

The UPGUTLLG subcommand causes the User Page Utilization Data Log screen (FCX293) to be displayed for the selected user.

## Format

➤ UPGUTLLg — *userid* ➤

## Parameters

### *userid*

Is the user identification of the virtual machine for which the User Page Utilization Data Log screen (FCX293) is to be displayed.

## Results

See [“FCX293, User Page Utilization Data Log Screen – UPGUTLLG”](#) on page 803.

# UQDIO

## Purpose

The UQDIO command causes the User QDIO Activity Screen (FCX257) to be displayed.

## Format

➤ UQDIO ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX257, User QDIO Activity Screen – UQDIO”](#) on page 733.

# UQDIOLOG

## Purpose

The UQDIOLOG command causes the User QDIO Activity Log (FCX258) to be displayed for the selected user.

## Format

`➤ UQDIOLog — userid ➤`

## Parameters

### *userid*

Is the user identification of the virtual machine for which the User QDIO Activity Log (FCX258) is to be displayed.

## Results

See [“FCX258, User QDIO Activity Log Screen – UQDIOLOG”](#) on page 735.

# URESP

## Purpose

The URESP subcommand causes the User Response Time Screen (FCX127) to be displayed.

## Format

➤ UREsp ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX127, User Response Time Screen – URESP”](#) on page 475.



# URESPLOG

## Purpose

The URESPLOG subcommand causes the User Response Time Log Screen (FCX165) to be displayed for the selected user.

## Format

➤ URESPLog — *userid* ➤

## Parameters

### *userid*

Is the user identification of the virtual machine for which the User Response Time Log Screen (FCX165) is to be displayed.

## Results


See [“FCX165, User Response Time Log Screen – URESPLOG”](#) on page 545.

# USER

## Purpose

The USER subcommand causes either the User Resource Details Screen (FCX115) or the User Resource Usage Screen (FCX112) to be displayed, based on whether a *userid* is specified.

## Format



## Parameters

### *userid*

Is the identification of a virtual machine for which detailed performance data are to be displayed. A User Resource Details Screen (FCX115) will be shown with performance information for the selected user only.

The User Resource Usage Screen (FCX112) will be shown if you enter the USER subcommand with no *userid*.

## Results

See [“FCX112, User Resource Usage – USER” on page 431](#) if you specified USER.

See [“FCX115, User Resource Details – USER userid” on page 443](#) if you specified USER *userid*.

# USERLOG

## Purpose

The USERLOG subcommand causes the User Resource Usage Log Screen (FCX162) to be displayed for the selected user.

## Format

➤ USERLog — *userid* ➤

## Parameters

### *userid*

Is the user identification of the virtual machine for which the User Resource Usage Log Screen (FCX162) is to be displayed.

## Results

See [“FCX162, User Resource Usage Log – USERLOG userid”](#) on page 542.

# USRMPLOG

## Purpose

The USRMPLOG subcommand causes the Multiprocessor User Activity Log Screen (FCX288) to be displayed for the selected user.

## Format

➤ USRMPLog — *userid* ➤

## Parameters

### *userid*

Is the user identification of the virtual machine for which the Multiprocessor User Activity Log Screen (FCX288) is to be displayed.

## Results

See [“FCX288, Multiprocessor User Activity Log Screen – USRMPLOG”](#) on page 794.

# USRPRCTM

## Purpose

The USRPRCTM subcommand causes the User Processor Time Screen (FCX333) to be displayed.

## Format

➤ USRPrctm ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX333, User Processor Time Screen – USRPRCTM”](#) on page 875.

# USRTMLOG

## Purpose

The USRTMLOG subcommand causes the User Processor Time Log Screen (FCX334) to be displayed for the selected user.

## Format

➤ USRTmlog — *userid* ➤

## Parameters

### *userid*

Is the user identification of the virtual machine for which the User Processor Time Log Screen (FCX334) is to be displayed.

## Results


See [“FCX334, User Processor Time Log Screen – USRTMLOG” on page 877](#).

# USTAT / USTATG

## Purpose

The USTAT subcommand causes the User Wait States Screen (FCX114) to be displayed (with user activity and wait state percentages shown numerically). The USTATG subcommand displays the same information graphically in the form of bar charts. Note that the latter requires access to GDDM code and a terminal with graphics capability.

## Format



## Parameters

Additional sort arguments are valid for the USTATG subcommand only. Use the SORT subcommand for changing the sorting sequence of the normal USTAT display.

### %ACT

Causes user output data to be sorted in descending order of user activity (percentage of samples where the user was found in queue). %ACT sorting will automatically be assumed if no other sort sequence has been specified.

### %PGW

Causes user output data to be sorted in descending order of page wait (percentage of in-queue samples where the user was found waiting for pages)

### %IOW

Causes user output data to be sorted in descending order of I/O wait (percentage of in-queue samples where the user was found waiting for an I/O to complete)

### %PSW

Causes user output data to be sorted in descending order of PSW wait (percentage of in-queue samples where the user was found in PSW wait but not idle)

### %RUN

Causes user output data to be sorted in descending order of runnable state. This is the percentage of in-queue samples where the user was found to be either using CPU or waiting for CPU, or in any state other than the ones specifically displayed (e.g. instruction simulation wait, console function, or IUCV soft wait states)

### %EL

Causes user output data to be sorted in descending order of resource wait state (percentage of in-queue samples where the user was found in the eligible list)

## Results

See [“FCX114, User Wait States Screen – USTAT/USTATG”](#) on page 438.

# USTATLOG

## Purpose

The USTATLOG subcommand causes the User Wait States Log Screen (FCX164) to be displayed for the selected user.

## Format

➤ USTATLog — *userid* ➤

## Parameters

### *userid*

Is the user identification of the virtual machine for which the User Wait States Log Screen (FCX164) is to be displayed.

## Results

See [“FCX164, User Wait States Log Screen – USTATLOG”](#) on page 544.



# USTLOG

## Purpose

The USTLOG subcommand causes the User Wait State Log Screen (FCX135) to be displayed.

## Format

➤ USTLog ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX135, User Wait State Log Screen – USTLOG” on page 491.](#)

# USTMPLOG

## Purpose

The USTMPLOG subcommand causes the Multiprocessor User Wait States Log Screen (FCX315) to be displayed for the selected user.

## Format

**➤** USTMPLog — *userid* **➤**

## Parameters

### *userid*

Is the user identification of the virtual machine for which the Multiprocessor User Wait States Log Screen (FCX315) is to be displayed.

## Results

See [“FCX315, Multiprocessor User Wait States Log Screen – USTMPLOG”](#) on page 845.

# UTRANDET

## Purpose

The UTRANDET subcommand causes the User Transaction Details Screen (FCX174) to be displayed for the selected user.

## Format

**►► UTRANDET — *userid* ►►**

## Parameters

### *userid*

Is the user identification of the virtual machine for which the User Transaction Details Screen (FCX174) is to be displayed.

## Results

See [“FCX174, User Transaction Details Screen – UTRANDET”](#) on page 559.

# UTRANLOG

## Purpose

The UTRANLOG subcommand causes the User Resources per Transaction Log Screen (FCX166) to be displayed for the selected user.

## Format

**►► UTRANLog — *userid* ►►**

## Parameters

### *userid*

Is the user identification of the virtual machine for which the User Resources per Transaction Log Screen (FCX166) is to be displayed.

## Results

See [“FCX166, User Resources per Transaction Log Screen – UTRANLOG”](#) on page 546.

# UTRANS

## Purpose

The UTRANS subcommand causes the User Resources per Transaction Screen (FCX129) to be displayed.

## Format

➤ UTRans ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX129, User Resources per Transaction Screen – UTRANS”](#) on page 478.

# VDISKS

## Purpose

The VDISKS subcommand causes the Virtual Disks in Storage Screen (FCX147) to be displayed.

## Format

➤ VDisks ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX147, Virtual Disks in Storage Screen – VDISKS” on page 511](#).

# VMRM

## Purpose

The VMRM subcommand causes the VM Resource Manager Screen (FCX241) to be displayed.

## Format

➤ VMRM ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX241, VM Resource Manager Screen – VMRM” on page 709](#).

# VNIC

## Purpose

The VNIC subcommand causes the Virtual Network Device Activity Screen (FCX269) to be displayed. This screen contains performance data for virtual network devices.

To see additional performance data for virtual network devices, see [“EVNIC” on page 176](#).

For a summary of virtual network device definitions, see [“GVNIC” on page 192](#).

## Format

▶▶ VNIC ▶▶

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX269, Virtual Network Device Activity Screen – VNIC” on page 757](#).



# VOLLOG

## Purpose

The VOLLOG subcommand causes the DASD Volume Performance Log Screen (FCX330) to be displayed for the selected I/O device.

## Format

**▶▶ VOLLog — *devno* ▶▶**

## Parameters

### *devno*

Is the real device number of the base RDEV of the volume about which information should be displayed on the DASD Volume Performance Log Screen (FCX330).

## Results

See [“FCX330, DASD Volume Performance Log Screen – VOLLOG”](#) on page 868.

# VOLUME

## Purpose

The VOLUME subcommand causes the DASD Volume Performance Screen (FCX329) to be displayed.

## Format

➤ VOLUME ➤

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX329, DASD Volume Performance Screen – VOLUME”](#) on page 865.

# VSWITCH

## Purpose

The VSWITCH subcommand causes the Virtual Switch Activity Screen (FCX240) to be displayed.

To view additional performance data for VSWITCH, see [“EVSWITCH” on page 177](#).

For a summary of VSWITCH definitions, see [“GVSWITCH” on page 193](#).

## Format

 VSWItch 

## Parameters

No parameters are allowed or tested for.

## Results

See [“FCX240, Virtual Switch Activity Screen – VSWITCH” on page 707](#).

## / (LOCATE)

### Purpose

The LOCATE subcommand searches the monitor output buffer for the character string specified.

### Format



### Parameters

where

/

Indicates a forward search. The search for the string specified starts with the current line + 1 and continues until either the string is found or the end of the display buffer is reached. 'Current line' is the uppermost non-header line shown on the display terminal.

-/

Indicates a backward search. The search starts with the current line - 1 and continues until either the string is found or the top of the display buffer is reached.

#### **string**

is the character string to be located. Both the string and the monitor data to be scanned will be translated to upper case for the search. If it could be found, the record containing this string will be shown as first data record on the screen.

## **'ENTER'**

### **Purpose**

Pressing the ENTER key without any input will cause a new measuring interval to be started, that is, performance data will be recalculated using current system counter values.



## Chapter 2. Performance Toolkit Reports

See “Performance Monitor Mode Subcommands” on page 129 for a description of the specific performance monitor subcommands you will need for selecting and manipulating performance reports, and to the "Using Performance Toolkit in Performance Monitor Mode" chapter in *z/VM: Performance Toolkit Guide* for some hints on how to work with the performance monitor.

## Report Layout Overview

The general performance report layout follows:

```

Identifier                               Interval start and end times           Mode
|                                     |
|                                     |
|                                     |
|                                     |
FCXnnn      CPU nnnn   SER nnnnn  Interval HH:MM:SS - HH:MM:SS    Perf. Monitor
.....
.....
.....
.....
.....
.....
.....
.....
.....
.....
.....
.....
.....
.....
.....
... message area ...
Command ==> _ Command Line
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

```

Figure 1. General Layout of Performance Monitor Reports

The header line indicates the following:

- The identifier of the current performance data display
- The CPU type being monitored
- The CPU serial number

The corresponding RSCS node-ID, or any other string, can alternatively be inserted if desired (controlled by 'SYSTEMID' argument of the FC MONCOLL subcommand, see [“FCCONTROL MONCOLL” on page 53](#) for details and space needed). Note that the CPU type and serial number will be replaced by the creation date of the monitor data being viewed while in 'Monitor Data Scan' mode.

- The period of time for which the performance values on the report have been calculated

The **mode** field contains the string 'Perf. Monitor' to indicate the current display mode. When concurrent printing of performance data has been activated, the string `Print on` will be shown in this field. When print output is to be written to a disk file, the string shown will be `Print on Disk`.

## Interval

The time interval for which the performance data on the selected screen have been calculated is shown in the header line of most screens. The possible intervals are:

- Current data (the last interval)
- Interim data (averages since the last INTERIM RESET)
- Average data (averages since the last PRINT RESET)
- Mean data (average since the last PRINT RESET)
- By-time log data (since the oldest measurement).

## Current Data

This is the initial display mode (and sometimes the only mode possible) for most performance displays. It provides information about system behavior during the last measuring interval, that is, the start and end times shown usually indicate the time when the last two samples were taken. This is the CP monitor sample interval for most displays; it can be a 1-minute interval that is not linked to the monitor sample interval for some detailed data screens.

You can switch from a display of average or interim data to current data by entering the CURRENT subcommand (see [“CURRENT” on page 165](#)).

## Interim Data

The INTERIM subcommand (see [“INTERIM” on page 200](#)) allows you to display interim averages of most performance data fields on screens that support interim data, that is, averages will be calculated for the current interim period. The start time inserted while interim averages are displayed is the start time for the interim period, that is, the time of the last interim RESET. Sections at the end of the screen descriptions indicate whether INTERIM averages are supported.

## Average Data

The AVERAGE subcommand (see [“AVERAGE” on page 133](#)) allows you to display the averages of most performance data fields for the entire time that monitoring was active or since the last RESET subcommand (automatically generated at midnight if no other RESET time has been entered). The start time inserted while averages are displayed is the start time for the entire period since the last RESET. Averages cannot be calculated for some of the detailed performance displays.

## Mean Data

History log-type data screens usually show mean data in the first entries, which display the overall averages since the last RESET subcommand or since data collection was activated. The lines are designated by the >>Mean>> string in the Interval End Time column. The interval shown in the header line is the period for which these averages were calculated.

Detail lines below the >>Mean>> data area still show earlier performance data prior to the RESET time.

## By-Time Log Data

History log-type data screens without mean data show performance data by time.

All "by time" logs by default show detail lines that contain data only for the last monitor sample interval, that is, a new detail line is built for each new sample interval. While this high-resolution mode of operation is probably what users need for analyzing performance problems in real time, it also requires a large amount of work buffer space for a given period or, for a given work buffer space, does not accommodate detail lines for long periods. The BYTIME argument of the FC SETTINGS subcommand allows defining a different mode of operation for most logs where the average or total values for a longer period are shown in each detail line. The length of that period is defined with the BYTIME argument.

The interval for the performance redisplay screen indicates the time stamp of the oldest measurement that can still be displayed as the starting time, with the seconds set to 00. The interval shown in the header line is the period for which these averages were calculated.



**Example:**

Entering this subcommand:

```
FC SETTINGS BYTIME 15
```

will lead to detail log lines being built that refer to periods of 15 minutes, even though much shorter monitor sample intervals can be used. See [“FCONTROL SETTINGS” on page 83](#) for a description of this subcommand.

**Exceptions:**

The BYTIME setting does not affect the REDISP display and the generation of detailed HISTLOG and benchmark logs.

## Function Key Definitions

---

The function key definitions shown for all examples, and also any references in the explanatory text, correspond to the initial setting. They can be redefined using the FC PFKEY SET *nn* FUNCTION *xxxx* command (see [“FCONTROL PFKEY” on page 66](#)). You should refer to the actual numbers displayed in the bottom line for a specific function.

## Data Consistency

---

Information from many different system control blocks (hundreds, or even thousands on large systems) has to be extracted during each data collect cycle. The fields cannot all be copied at a single moment, and slight inconsistencies between values from different sources must be expected. This inconsistency will be more noticeable for very short measuring intervals and when system load is high.

Although Performance Toolkit tries to use the CP monitor sample interval for collecting the control block based data of the remaining permanently updated general screens, the interval cannot be exactly the same and inconsistencies must be expected. This is true especially when the performance toolkit machine does not have a direct connection to the \*MONITOR IUCV service but has to read the data from another machine's disk.

Note also that some detailed performance displays which are based on data collection from CP control blocks use a fixed one-minute interval for data collection. Their collection interval is not linked to the collection interval of all the general performance displays and, because of the different interval, the values shown will not usually match the equivalent values on a general screen.

## Data Field Values

---

In performance monitor mode, field values are numeric, when applicable and possible.

There are times when a field value is *temporarily* unavailable — at initial startup or if a monitor domain is disabled, for example. At these times, Performance Toolkit displays dots to indicate that the field value is unavailable. Likewise, if a field value cannot be calculated (because the values are too high or because they exceed a conversion routine's capacity, for example), the Performance Toolkit displays question marks. The question marks are displayed during the time frame in which the conditions exist. If this issue persists over a significant number of intervals, it should be investigated.

There could be conditions in which certain fields are *permanently* unavailable — the field does not exist in the current z/VM version or the field value cannot be calculated due to its specific type for the Mean/Total/System lines, for example. In these cases, Performance Toolkit displays dashes or dots.

When the value that is to be displayed exceeds the field width, a scaling suffix is used. For a decimal scale factor (based on a power of 10), a suffix of k, m, g, t, p, or e is appended. For a binary scale factor (based on a power of 2), a suffix of K, M, G, T, P, or E is appended. These suffixes stand for kilo, mega, giga, tera, peta, and exa scaling coefficients, respectively.

## Data Completeness

---

Monitor 'event' data records will be used, when available, to determine the 'session' start and end for all the monitor data based user displays, and for the NSS, DSPACESH and VDISKS displays, to capture all the data available. Some data will be lost when the event records are missing. Note also that the required monitor domains should remain enabled for the whole period while data collection is active. Disabling/enabling monitor domains during data collection is not recommended, since it can lead to incomplete and/or incorrect data being displayed.

## FCX100, CPU Load and Transactions – CPU

**Performance Toolkit Subcommand:** “CPU” on page 162.

The CPU Load and Transactions report (FCX100):

FCX100		CPU		nnnn	SER		nnnnn	Interval					HH:MM:SS	-	HH:MM:SS	Perf. Monitor	
CPU Load													Status or				
PROC	TYPE	%CPU	%CP	%EMU	%WT	%SYS	%SP	%SIC	%LOGLD	%PR	%ENT	ded.	User				
P00	CP	35	1	34	65	1	0	75	35	0	100	Master					
P01	CP	31	1	30	69	0	0	64	31	0	100	Alternate					
P02	CP	12	2	10	88	1	0	91	12	0	100	Alternate					
P03	CP	12	2	10	88	1	0	91	12	0	100	Alternate					
P04	CP	14	2	12	86	1	0	87	14	0	100	Alternate					
P05	CP	2	0	2	98	0	0	87	2	0	67	Alternate					
P06	CP	1	0	1	99	0	0	92	1	0	67	Alternate					
P07	CP	2	0	2	98	0	0	93	2	0	0	Alternate					
P08	CP	3	0	3	97	0	0	95	3	0	0	Alternate					
P09	CP	0	0	0	100	0	0	...	0	0	0	Alternate					
P0A	CP	0	0	0	100	0	0	...	0	0	0	Alternate					
P0B	CP	0	0	0	100	0	0	100	0	0	0	Alternate					
P0C	ZIIP	0	0	0	100	0	0	100	0	0	50	Alternate					
P0D	ZIIP	0	0	0	100	0	0	100	0	0	50	Alternate					
P0E	IFL	0	0	0	100	0	0	100	0	0	57	Alternate					
P0F	IFL	0	0	0	100	0	0	100	0	0	0	Alternate					
P10	IFL	0	0	0	100	0	0	100	0	0	0	Alternate					
P11	IFL	0	0	0	0	0	0	...	...	100	0	Parked					
Total SSCH/RSCH				54/s		Page rate				.0/s		Priv. instruct.		1906/s			
Virtual I/O rate				8/s		XSTORE paging				.0/s		Diagnose instr.		88/s			
Total rel. SHARE				3317		Tot. abs SHARE				0%							
Queue Statistics:				Q0		Q1		Q2		Q3		User Status:					
VMDBKs in queue				2		0		2		2		# of logged on users					
VMDBKs loading				0		0		0		0		# of dialed users					
Eligible VMDBKs						0		0		0		# of active users					
El. VMDBKs loading						0		0		0		# of in-queue users					
Tot. WS (pages)				4475		0		133353		83280		% in-Q users in PGWAIT					
Reserved												% in-Q users in IOWAIT					
85% elapsed time				4.314		.719		5.752		34.51		% elig. (resource wait)					
Transactions				Q-Disp		trivial		non-triv		User Extremes:							
Average users				.6		.0		39.5		Max. CPU %							
Trans. per sec.				2.2		.1		.9		Reserved							
Av. time (sec)				.312		.220		42.37		Max. IO/sec							
UP trans. time						.294		.000		Max. PGS/s							
MP trans. time						.000		42.37		Max. RESPG							
System ITR (trans. per sec. tot. CPU)								2.8		Max. MDCIO							
Emul. ITR (trans. per sec. emul. CPU)								3.0		Max. XSTORE							
Command ==>																	
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return																	

Figure 2. Layout of CPU Load and Transactions report (FCX100)

This report displays data of z/VM systems with or without expanded storage. The corresponding fields will remain empty where XSTORE is not available.

### Field Descriptions:

*CPU Load:*

#### PROC

The processor ID, in the format *Pnn*, where *nn* is the logical processor ID in hexadecimal format. The first processor shown is the base processor.

#### TYPE

The processor type. Possible values are:

##### CP

Central Processor

**ICF**

Internal Coupling Facility

**IFL**

Integrated Facility for Linux

**ZIIP**

IBM Z Integrated Information Processor

**%CPU**

The total CPU load as a percentage.

**%CP**

The percentage of supervisor (CPU) time. This includes supervisor time spent for specific users and supervisor time for system services, which is also shown under the %SYS heading.

**%EMU**

The percentage of emulation time, that is, the processor time used while virtual machines were running under control of Start Interpretive Execution (SIE) microcode ("virtual" CPU time).

**%WT**

The percentage of time spent in wait state. No distinction is made among these wait states: *idle wait*, *I/O wait*, and *page wait*. See the user status fields in this report or in [“FCX114, User Wait States Screen – USTAT/USTATG” on page 438](#) for indications of excessive I/O wait or page wait.

**%SYS**

The percentage of CPU time spent exclusively for system services, that is, not attributable to specific users.

**%SP**

The percentage of time spent spinning on a lock.

**Note:** All of the above percentages of CPU usage are calculated based on elapsed time, so that the load figures provide a meaningful basis for capacity planning. For systems running in an LPAR or second level, the value is not, however, a good basis for detecting a CPU bottleneck. Because CP might not always be able to use a logical processor when it is needed (due to contention from other systems on the processor complex), there could be many users waiting for CPU, even though the CPU usage shown is well below 100%.

See [“FCX114, User Wait States Screen – USTAT/USTATG” on page 438](#) and [“FCX135, User Wait State Log Screen – USTLOG” on page 491](#) for information about CPU wait percentages and the %LOGLD value or [“FCX126, LPAR Load Screen – LPAR” on page 472](#) for processor usage figures, which are based on the sum of processor active and voluntary wait time, similar to the output of the INDICATE LOAD subcommand.

**%SIC**

The percentage of SIE exits by interception (that is, SIE mode stopped because CP had to simulate an instruction). If no SIE instructions are reported on a processor, dots are displayed in this field.

**%LOGLD**

The total logical CPU load, that is, the total CPU load calculated based on the sum of processor active plus processor wait time and *not* total elapsed time. The value is calculated only for z/VM systems running in an LPAR; dots are displayed otherwise.

This value does not correctly reflect system load, but it is a better indicator of CPU bottlenecks than the %CPU value when the z/VM system runs in an LPAR.

**%PR**

The total percentage of the time when this CPU was in the parked state.

**%ENT**

The average value of a portion of a physical CPU to which this vertical CPU is entitled, as a percentage. A value of 100 indicates that the CPU was polarized vertically with high entitlement during the measured period. A value of 0 indicates that the CPU was polarized horizontally.

**Status or ded. User**

Contains the user ID of the virtual machine the processor is dedicated to, or general status information if the processor is not available for some reason. If the processor is not dedicated to a user, the status field can be:

**offline**

If the processor is in Offline status

**quiescing**

If the processor is to be quiesced

**quiesced**

If the processor is quiesced

**stopped**

If the processor has been stopped

**Master**

For the master processor (usually the IPL processor)

**Alternate**

For alternate processors

**Parked**

If the processor is in a Parked state

If there is no special status to be shown, dashes are displayed when running Performance Toolkit against z/VM 5.4 or earlier data.

*General system load fields:***Total SSCH/RSCH**

The total SSCH (Start Sub-CHannel) and RSCH (Resume Sub-CHannel) rate for the system. The value inserted is the sum of the SSCH and RSCH rates run on all processors used by the system.

**Virtual I/O rate**

The sum of the I/O request rates by virtual machines, excluding I/O to virtual UR devices.

**Total rel. SHARE**

The sum of the relative SHARE values of all users in the dispatch list.

**Page rate**

The total system page rate.

**XSTORE paging**

The total XSTORE page rate.

**Tot. abs SHARE**

The sum of the absolute SHARE values of all users in the dispatch list.

**Priv. instruct.**

The average number of privileged instructions *simulated* per second by CP during the last measuring interval.

**Diagnose instr.**

The average number of diagnose instructions run per second during the last measuring interval.

*Queue Statistics:*

Information for all z/VM queues (classes 0, 1, 2, and 3) are shown where available for the following fields:

**VMDBKs in queue**

The number of VM definition blocks (VMDBKs) in the dispatch list (that is, active or waiting for service) in a specific class. A VMDBK usually represents a virtual machine, but virtual MP users have several VMDBKs describing each of their virtual processors.

**VMDBKs loading**

The number of loading VMDBKs (a subset of the total number of VMDBKs in the dispatch list, with very high paging requirements).

**Eligible VMDBKs**

The number of VMDBKs in the eligible list (waiting for some resource - main storage, for example - to become available).

**El. VMDBKs loading**

The number of loading VMDBKs in the eligible list.

**Tot. WS (pages)**

The sum of the working sets of all users active in the corresponding class, expressed as the number of pages used.

**85% elapsed time**

The elapsed time slice that determines the maximum time a VMDBK can remain in the dispatch list before it is dropped. This value is continuously adapted for class 1 transactions so 85% of them can complete within a single class 1 elapsed time slice, that is, the value reflects the system responsiveness to user demand *and* workload characteristics. Consider this when comparing the values from different systems.

The values for class 0, 2, and 3 transactions are multiples of the class 1 value, with fixed multiplication factors of 6, 8, and 48, respectively.

*Transactions:*

Performance values are shown for three separate transaction groups:

**Q-Disp**

Values for quick-dispatch users

**trivial**

Values for trivial transactions

**non-trv**

Values for non-trivial transactions

Transactions are considered to have been *trivial* if they completed with a single drop from Q1 and no drops from any of the other queues. All other transactions are counted as *non-trivial*.

*Performance values:***Average users**

The average number of users active in the corresponding transaction group, calculated by dividing the total elapsed time spent in the group during the last interval by the length of the interval.

**Trans. per sec.**

The transaction rate per second.

**Av. time (sec)**

The average transaction time, including UP and MP transactions.

**UP trans. time**

The average transaction time for UP users.

**MP trans. time**

The average transaction time for MP users.

**System ITR**

The System Internal Throughput Ratio, that is, the number of transactions on the system per second of total CPU time used.

**Emul. ITR**

The Emulation Internal Throughput Ratio, that is, the number of transactions on the system per second of emulation CPU time used.

*User Status:*

Most of the following values are samples extracted at the end of the last measuring interval (exception: active users). They are *not* averages for this measuring interval.

**# of logged on users**

The number of logged-on users.

**# of dialed users**

The number of users dialed in to the system.

**# of active users**

The number of users that were found to have been active during the last measuring interval. This is the number of active VMDBKs that describe virtual processors (virtual machines with several processors will be counted as multiple "users"), and a VMDBK is considered to have been active if one of the following occurred during the last measuring interval:

- It has used at least 10 msec of total CPU
- It has used some virtual CPU
- It has had some I/O activity

**# of in queue users**

The number of users found in the dispatch list (or waiting to be added to the dispatch list) at the end of the measuring interval. The virtual machine where this performance monitor was active (collecting data at the end of the measuring interval) is not included in this figure, because it would bias the statistics. This value is taken from user wait state analysis, so it might not correspond exactly to the values shown for in-queue VMDBKs in the *Queue Statistics* fields in this report that come from another source (and cannot be collected at exactly the same time).

**% in-Q users in PGWAIT**

The percentage of in-queue users found in page wait state during the interval.

**% in-Q users in IOWAIT**

The percentage of in-queue users found in I/O wait state during the interval. This is calculated as the sum of the wait state percentages for:

- I/O wait (during CCW translation)
- I/O active

as displayed in [“FCX114, User Wait States Screen – USTAT/USTATG”](#) on page 438 and [“FCX135, User Wait State Log Screen – USTLOG”](#) on page 491.

This interpretation is an attempt to provide the performance analyst with a meaningful value for I/O wait, but be aware that it is not quite correct. While instruction simulation wait will be mostly I/O wait at least where CMS users are concerned, some other instruction simulation wait will also be contained in this number.

**% elig. (resource wait)**

The percentage of in-queue users found in an *eligible list* (that is, waiting for a constrained resource) during the measuring interval.

**User Extremes:**

These fields display the user ID and specific load imposed on the system by the heaviest user for the following areas:

**Max. CPU %**

The user that consumed the most CPU during the last measuring interval and the percentage of total CPU used (based on one processor).

**Max. IO/sec**

The virtual machine with the highest I/O rate (non-spooled) during the last measuring interval and its I/O rate.

**Max. PGS/s**

The virtual machine with the highest paging rate and its page rate.

**Max. RESPG**

The virtual machine with the highest number of currently-resident pages and the number of resident pages.

**Max. MDCIO**

The virtual machine with the highest number of MDCACHE inserts and the MDCACHE insert rate.

This value is an approximate indicator. Due to the collection mechanism, the numbers shown can be considerably lower than the users' actual MDC insert activity, especially for long monitor intervals.

You might see dashes in this field when running Performance Toolkit against z/VM 5.4 or earlier data.

**Max. XSTORE**

The virtual machine with the highest number of XSTORE pages and the current number of XSTORE pages used by this machine (without dedicated XSTORE).

**Effect of AVERAGE and INTERIM subcommands:**

Overall period averages are shown for all performance fields, except the User Extremes fields.



## FCX101/FCX181 Remote Performance Log – REDISP, REDHIST

**Performance Toolkit Subcommand:** “REDISP” on page 269 or “REDHIST” on page 268.

This report is shown after entering the REDISP subcommand (or pressing the corresponding function key). It allows redisplaying key performance values from previous measurements. The data shown can come from either of the following two sources:

1. *Performance redisplay buffer* (in storage) which contains data from the last measurements made in your machine during the *current* session. The default number of measurements that can be shown is 720. It can be changed using the FC MONCOLL REDISP *nnn* command to any other value between 60 and 14400.

The report identifier for this display is FCX101.

2. *Performance history files* (on disk, old format). These will exist only in machines where history data collection on disk has been activated (FC MONCOLL PERFLOG command). If available, they can be displayed for analyzing a previous day's performance data. The status message area (center part of the top line) will then look like this:

```
Viewing File fn ft fm
```

- that is, it will contain the full file identification (file name, file type and file mode) of the history file being viewed.

The identifier for the report based on history logs on disk is FCX181.

FCX101	CPU nnn				SER nnn		Interval		HH:MM:SS	Perf. Monitor							
TIME	CPU	%CP	%EM	%WT	%SY	-	-	IO/S	VIO/S	PG/S	XPG/S	DIAG	PRIV	LOGN	ACT		
14:09	163	53	110	237	11	.	.	308	270	194	522	1199	2117	1519	226		
14:10	153	51	102	247	11	.	.	322	269	242	571	1157	1957	1520	260		
14:11	222	57	165	178	12	.	.	349	323	207	600	1166	2072	1517	257		
14:12	146	51	95	254	11	.	.	307	296	141	418	1110	1981	1515	214		
14:13	176	49	127	224	11	.	.	284	328	180	422	1084	1832	1510	228		
14:14	167	54	113	233	11	.	.	354	323	177	429	1208	2069	1514	232		
14:15	225	67	158	175	15	.	.	441	345	218	517	1542	2672	1513	234		
14:16	212	66	146	188	14	.	.	403	332	223	622	1572	2847	1516	236		
14:17	187	57	130	213	12	.	.	376	364	195	527	1193	2269	1514	229		
14:18	164	52	112	236	11	.	.	340	302	203	561	1155	1953	1516	228		
14:19	275	59	216	125	14	.	.	406	397	209	521	1261	2293	1518	230		
14:20	177	51	126	223	11	.	.	319	274	203	477	1128	1939	1513	218		
14:21	180	54	126	220	11	.	.	360	313	170	492	1292	2294	1516	210		
14:23	144	42	102	256	10	.	.	269	240	133	442	938	1600	1520	213		
14:24	212	53	159	188	12	.	.	363	315	123	456	1188	2171	1522	223		
14:25	160	51	109	240	11	.	.	397	322	199	470	1165	2023	1528	266		
14:26	149	50	99	251	11	.	.	320	222	190	437	1036	1873	1526	247		
14:27	235	57	178	165	13	.	.	374	483	204	535	1359	2354	1524	238		

Command ==>

F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return

Figure 3. Layout of Remote Performance Log Report (FCX101)

You can view subsets of performance data by selecting one of the following subsets:

- 'CPU' data, with actual processor load and overall I/O and paging information.

Selected by CPUDATA-command.

- 'Transaction' data, with user activity and queue statistics.

Selected by TRANSACT-command.

- 'Storage Utilization' information, including XSTORE and minidisk cache data.

Selected by STORUTIL-command.

- 'Max. User' values which show which user caused the highest load in a specific area.

Selected by MAXUSER-command.

The commands for selecting a specific group of performance variables can be supplied as arguments with the REDISPLAY command, or they can be entered separately once performance redisplay mode has been entered.

Once entered, the LEFT/RIGHT shift position is retained even after quitting and re-entering performance redisplay mode.

Each line contains a series of performance values from the general CPU report, the storage utilization report, or the SYSTEM report, preceded by a time stamp indicating when they were calculated. The most recent measurements will be shown first when switching into re-display mode (last measurement in bottom line). Previously calculated values can be viewed by shifting the data backwards ('BACKWARD' sub-command or corresponding PF-key). All values will be shifted upwards by one line and new values inserted into the bottom line whenever a new measurement is made.

See the "Remote Performance Log Display" section in the *z/VM: Performance Toolkit Guide* for more information about this report.

## Redisplay of General CPU Performance Data

This part of the performance redisplay data can be selected

- by initially entering the command REDisp CPUData, or
- by entering the command CPUData while already viewing other performance redisplay data.

This part will be shown initially, before any left/right shift commands have been entered.

FCX101	CPU nnnn				SER nnnnn				Interval HH:MM:SS - HH:MM:SS				Perf. Monitor			
TIME	CPU	%CP	%EM	%WT	%SY	-	-	IO/S	VIO/S	PG/S	XPG/S	DIAG	PRIV	LOGN	ACT	
08:14	391	102	289	9	15	.	.	243	565	270	664	2197	3055	1025	275	
08:15	378	92	286	22	14	.	.	250	541	173	599	2474	3211	1040	268	
08:16	359	85	274	41	14	.	.	239	532	195	627	2054	2902	1062	283	
08:17	379	100	279	21	15	.	.	305	600	198	742	2482	3435	1084	303	
08:18	370	98	272	30	15	.	.	274	538	235	774	2550	3393	1095	304	
08:19	384	91	293	16	15	.	.	267	520	226	680	2183	3036	1112	288	
08:20	382	100	282	18	15	.	.	272	531	183	718	2807	3612	1127	300	
08:21	383	97	286	17	16	.	.	293	627	283	869	2574	3454	1145	302	
08:22	391	92	299	9	15	.	.	274	608	210	896	2175	3032	1160	313	
08:23	395	96	299	5	15	.	.	289	605	174	899	2310	3257	1177	317	
08:24	355	86	269	45	13	.	.	241	546	122	738	2091	2901	1185	304	
08:25	391	95	296	9	15	.	.	286	1078	210	719	3806	4750	1203	312	
08:26	400	99	301	0	15	.	.	278	630	176	854	2714	3733	1225	318	
08:27	369	88	281	31	14	.	.	249	575	167	771	3116	4007	1238	309	
08:28	337	86	251	63	13	.	.	255	527	134	787	1996	2879	1251	321	
08:29	361	91	270	39	14	.	.	263	610	202	703	2124	3080	1273	293	
08:30	305	92	213	95	15	.	.	305	593	216	797	2278	3169	1288	322	
08:31	331	100	231	69	15	.	.	312	671	194	861	2542	3484	1304	335	
Enter 'GRAPHICS' command for history graphics selection menu																
Command ==>																
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return																

Figure 4. Layout of performance redisplay screen: CPU data

### Field Description for CPU Data:

#### TIME

Time stamp. The values shown on this line are either averages calculated for the time interval between the current and the previous time stamp (usually equivalent to the CP monitor sample interval) or samples taken at the end of the measuring interval (current time stamp).

A time stamp as shown (format *hh:mm*)

will be inserted when looking at data from your current redisplay buffer or from a detailed 'mmdyy PERFLOGn' history file. When looking at a condensed history file 'ACUM PERFHIST', the time stamp has the format mmdyy hh, i.e. you'll see two fields containing the date and the hour during which the values have been collected.

**CPU**

Total CPU load in percent. On systems with more than one processor this field contains the *total* of the loads on *all* processors (applies also to the following six other CPU load indicators).

All of the processor load figures have been calculated based on elapsed time, so as to provide meaningful data for capacity planning. Note, however, that these figures may be misleading for systems running in an LPAR, or second level under another VM, unless the processors have been dedicated to this VM system, an overall load value well below the theoretical maximum may be no guarantee that you do not have a CPU bottleneck on the system, because other systems may contend for the same processors. Look at CPU wait percentages in the USTAT and USTLOG screens as indicators of a CPU bottleneck if you suspect problems in this area.

**%CP**

The percentage of supervisor (CPU) time.

**%EM**

The percentage of emulation (problem) time.

**%WT**

The percentage of time spent in wait.

**%SY**

The percentage of supervisor time spent for system services which could not be charged to any specific user.

**IO/S**

Number of SSCH or RSCH operations executed per second on all processors

**VIO/S**

Number of virtual I/O operations per second, without UR I/O

**PG/S**

Page rate (pages read/written per second)

**XPG/S**

XSTORE (expanded storage) page rate (inserted only on systems with expanded storage)

**DIAG**

The number of diagnose instructions executed per second by CP

**PRIV**

The number of privileged instructions per second simulated by CP

## Redisplay of User Transaction Data

This part of the performance redisplay data can be selected

- by initially entering the command REDISP TRANSACT, or
- by entering the command TRANSACT while already viewing other performance redisplay data.



**%IQ**

Percentage of in-queue users found in I/O wait state during the interval. This is calculated as the sum of the wait state percentages for

- I/O wait (CCW translation)
- I/O active

as reported in the user wait state screens USTAT and USTLOG.

**%LD**

Percentage of 'loading' in-queue users during the interval, i.e. of users with a very high page read activity

**%EL**

Percentage of 'in-queue' users in an eligible list (in resource wait) during the interval

**Q1**

Number of class 1 users in the dispatch list at the end of the interval

**Qx**

Number of non-class 1 users (in either of the classes 0, 2 or 3) in the dispatch list at the end of the interval

**Q1L**

Number of in-queue class 1 users in 'loading' state

**QxL**

Number of in-queue non-class 1 users (in either of the classes 0, 2 or 3) in 'loading' state

**E1**

Number of class 1 users in the eligible list at the end of the interval

**Ex**

Number of non-class 1 users (in either of the classes 0, 2 or 3) in the eligible list

**E1L**

Number of eligible class 1 users in 'loading' state

**ExL**

Number of eligible non-class 1 users (in either of the classes 0, 2 or 3) in 'loading' state

## Redisplay of Storage Utilization Data

This part of the performance redisplay data can be selected

- by initially entering the command REDISP STORUTIL, or
- by entering the command STORUTIL while already viewing other performance redisplay data.

```

FCX101      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS      Perf. Monitor

TIME > PPAG %ST ALO/S  FPGS %FR SHAR #TW ESCN %PGSL %SPSL XSTAV %XS XAL/S XAGE
08:14 52514 13 1918 .... 94 1551 0 0 13 61 1024M 99 459 516
08:15 52380 7 1689 .... 94 1539 0 0 14 61 1024M 99 378 516
08:16 52139 15 1864 .... 94 1930 0 0 15 61 1024M 99 406 508
08:17 51926 10 2124 .... 94 1540 0 0 16 61 1024M 99 455 504
08:18 51844 11 1950 .... 94 1551 1 0 16 61 1024M 99 482 496
08:19 51649 13 1767 .... 94 1489 0 0 17 61 1024M 99 437 490
08:20 51519 14 1802 .... 94 1410 0 0 18 61 1024M 99 449 493
08:21 51325 14 2133 .... 94 1861 0 0 20 61 1024M 99 548 490
08:22 51164 15 2031 .... 95 1485 0 0 20 61 1024M 99 544 500
08:23 50990 11 2250 .... 94 1943 0 0 21 61 1024M 99 534 517
08:24 50915 11 1809 .... 95 1820 0 0 21 61 1024M 99 427 540
08:25 50774 12 1856 .... 94 1920 0 0 22 61 1024M 99 454 495
08:26 50542 9 2032 .... 95 1646 0 0 23 61 1024M 99 532 455
08:27 50379 10 1923 .... 94 1686 0 0 23 61 1024M 99 475 439
08:28 50253 8 1753 .... 94 1721 0 0 24 61 1024M 99 475 430
08:29 50031 13 1942 .... 95 1616 0 0 24 61 1024M 99 464 406
08:30 49893 16 1980 .... 95 1659 1 0 25 61 1024M 99 510 391
08:31 49722 10 2220 .... 94 1734 0 0 25 61 1024M 99 547 379
Enter 'GRAPHics' command for history graphics selection menu
Command ==>
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return

```

Figure 6. Layout of performance redisplay screen: Storage Utilization data

Only part of the available storage utilization data can be shown on the initial screen. Use the RIGHT command or press the corresponding function key for viewing the remainder of the variables.

#### Field Descriptions for Storage Utilization Data:

##### TIME

The time stamp (see CPU data for details).

##### PPAG

The number of pageable pages in the DPA at the end of the measuring interval.

##### %ST

The percentage of pageable pages used by the current in-queue users' working sets.

##### ALO/S

The number of page frame requests per second, that is, the page frame allocation rate.

##### FPGS

This field is deprecated (dots are displayed).

##### %FR

The percentage of subpool free storage that is currently in use

##### SHAR

The number of shared pages.

##### #TW

The rate per second at which tasks were found to be waiting for a page.

##### ESCN

The percentage of demand scans (calls to the available list replenishment routine) that did not complete after the first scan.

Note that the first scan takes page frames from dormant users, and only unreferenced pages from active users, and is therefore harmless. If the number of page frames freed in this way is not yet sufficient, a second and, if necessary, even a third scan ('emergency scan') will be done, and referenced pages from active users, shared segments and CP itself will then be taken. This can severely impact performance and should be avoided. The number shown represents the number of scans that completes after the second scan or the emergency scan, or that were unsuccessful, as a percentage of all demand scans.

The number should be 0; high numbers indicate insufficient real storage.

##### %PGSL

The page slot utilization, in percentage of available slots.

**%SPSL**

The SPOOL slot utilization, in percentage of available slots.

**XSTAV**

XSTORE size available for CP

**%XS**

The percentage of XSTORE blocks in use.

**XAL/S**

The number of XSTORE pages allocated per second.

**XAGE**

The average age of XSTORE blocks, in seconds.

**MAGE**

The average age of XSTORE blocks at migration, in seconds.

**BKRD**

The blocking factor when reading pages from DASD,

**BKMG**

The blocking factor when migrating pages out of XSTORE,

**BKST**

The blocking factor for page steals, that is, when writing pages from main storage to DASD.

**MDUS**

The number of users which caused at least one insert operation into the minidisk cache during the last fair share interval.

**MDNE**

The number of XSTORE blocks and of central storage page frames used by the minidisk cache. See STORAGE or MDCSTOR displays for monitoring both values separately.

**MDR/S**

The number of read operations per second to the minidisk cache.

**%MDR**

The minidisk cache read block percentage, that is, the percentage of read block operations out of the total of read and write block operations.

**%MDH**

The minidisk cache read percentage, that is, the percentage of minidisk read operations which did not require an I/O operation to real DASD

## Redisplay of User Extremes

This part of the performance redisplay data can be selected

- by initially entering the command REDISP MAXUSER, or
- by entering the command MAXUSER while already viewing other performance redisplay data.

FCX101	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor	
TIME	>Max. CPU	Max. Vector	Max. I/O	Max. Paging	Max. Resident	Ma
08:14	CHBAT01	54.1 .....	... CHVTAM	48.3 RKAU	16.7 CHNETV	17
08:15	STW	49.3 .....	... SBU	57.6 FHI	9.6 CHNETV	17
08:16	STW	66.7 .....	... CHVTAM	50.6 DKN	7.3 CHNETV	17
08:17	STW	50.3 .....	... CHVTAM	51.7 PTR	8.6 CHNETV	17
08:18	STW	64.9 .....	... CHVTAM	53.2 BCH	11.0 CHNETV	15
08:19	CHBAT03	58.0 .....	... CHVTAM	55.0 TEL2	6.4 CHNETV	8
08:20	STW	46.9 .....	... CHVTAM	47.0 SWU	6.1 CHCAL	5
08:21	STW	50.2 .....	... CHVTAM	50.5 JBB	7.5 CHVSCS	4
08:22	CHBAT02	52.7 .....	... CHVTAM	51.1 MW	7.2 CHVSCS	4
08:23	CHBAT02	49.8 .....	... CHVTAM	48.4 CAST	5.7 VBO	4
08:24	CHBAT02	67.2 .....	... CHVTAM	48.5 CSTB	5.6 CHCAL	5
08:25	MAINT	21.3 .....	... ROP	520 SYSTEM	24.7 CHCAL	6
08:26	MAINT	14.1 .....	... CHVTAM	56.7 TUL	6.3 CHCAL	5
08:27	CHBAT01	29.7 .....	... CHVTAM	52.0 MEV	6.9 CHCAL	6
08:28	CHBAT03	27.6 .....	... CHVTAM	52.5 KIM	6.5 CHCAL	5
08:29	CHBAT01	32.0 .....	... CHVTAM	54.3 DRUE	9.5 CHCAL	5
08:30	GYS	24.4 .....	... CHVTAM	48.7 FJEI	9.6 CHCAL	5
08:31	MKOP	24.0 .....	... CHVTAM	51.3 ANB	8.2 CHCAL	5
Enter 'GRAPHICS' command for history graphics selection menu						
Command ==>						
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return						

Figure 7. Layout of performance redisplay screen: Max. User data

Only part of the available user extremes data can be shown on the initial screen. Use the 'RIGHT' command, or press the corresponding PF-key for viewing the remainder of the variables.

#### Field Description for Max. User Data:

##### TIME

Time stamp, see 'CPU' data for details

##### Max. CPU

User-ID which consumed most CPU during the last measuring interval, and the percentage of total CPU used

##### Max. Vector

User-ID which caused the highest load on the vector facility during the last interval, and the percentage used. On z/VM 5.2 and later, dots will be inserted here (or on earlier systems, when Vector Facility support is not available).

##### Max. I/O

User-ID which executed the highest number of non-spooled I/O operations during the last measuring interval, and its I/O rate

##### Max. Paging

User-ID which imposed the highest paging load onto the system during the last measuring interval, and its page rate

##### Max. Resident

User-ID with the highest number of pages resident at the end of the last measuring interval, and the number of resident pages

##### Max. MDC Ins.

User-ID with the highest rate of MDCACHE inserts, and the MDCACHE insert rate during the measuring interval.

Note that this value can be only a rough indicator. Due to the collect mechanism the numbers shown can be considerably lower than the user's actual MDC insert activity, especially for long monitor intervals.

##### Max. XSTORE

User-ID with the highest number of blocks in expanded storage, and the number of blocks occupied.



## FCX102, System Counters – SYSTEM

**Performance Toolkit Subcommand:** “SYSTEM” on page 316.

The System Counters report (FCX102):

FCX102	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor
<b>Operation</b>	<b>Count</b>	<b>Rate/s</b>	<b>Operation</b>	<b>Count</b>	<b>Rate/s</b>
Real SSCH instructions	9673	161	Real CSCH instructions	1	.0
Real HSCH instructions	0	.0	El. time slice drops	336	5.6
SVC instr. simulated	0	.0	SVC interrupts reflectd	0	.0
SVC 76 reflected	0	.0	Diagnose I/O requests	489	8.2
FP external call simul.	0	.0	FP partial executions	7449	124
Fast-path SIGP simulat.	0	.0	FP simul of Diag.X'44'	0	.0
Diag.X'9C', Not Ready	18	.3	Diag.X'9C', Wait State	642	10.7
Diag.X'9C', Soft Wait	7	.1	Diag.X'9C', Alr.Dispat.	9553	159
FP successful x-lates	19	.3	CCW chains not FP-elig.	0	.0
Fast-path aborts	0	.0	Total FP xlate attempts	19	.3
Nr. of SIE executions	55674	927	Nr. of SIE intercepts	30232	503
Entries to enabled wait	16548	275	CP issued Diag.X'44'	0	.0
HPF write	99999	99.9	HPF done	99999	99.9
<b>Storage Management</b>					
Subpool FREE req bk <2G	34118	569	Total FREE req bk <2G	34118	569
Subpool FREE req bk >2G	178404	2973	Total FREE req bk >2G	178407	2973
Reserved			Storage fast clears	294	4.9
Avail 1st frame req <2G	1052	17.5	Avail 1st frame req >2G	740	12.3
Available 1st empty <2G	0	.0	Available 1st empty >2G	0	.0
Demand scan 1st pass	---	---	Demand scan 2nd pass	---	---
Demand scan emergency	---	---	Demand scan not satisf.	---	---
Dispatch 1st pgs stolen	---	---	Shared stor. pgs taken	---	---
Pgs from dormant users	---	---	Eligible 1st pgs stolen	---	---
Fast PGINs from XSTORE	0	.0	Pages taken for FREE	0	.0
PGOUTs main to XSTORE	0	.0	Slow PGINs from XSTORE	0	.0
XSTORE allocations	0	.0	Reserved		
Glbl cycl list searched	0	.0	XSTORE releases	0	.0
Migr thresh buf increas	0	.0	Migr. target time reset	0	.0
Page migr. from dormant	0	.0	Migr thresh buf lowered	0	.0
Page migr. from active	0	.0	Dormant with page migr.	0	.0
Shared pages migrated	0	.0	Active with page migrat	0	.0
Blocks migrated from CP	0	.0	Shared sys with pg migr	0	.0
Blocks migrated to DASD	0	.0	PGMBKs sel. during migr	0	.0
No I/O for pg migration	0	.0	XSTORE migr invocations	0	.0
Pg not referenced (STL)	0	.0	Pg not referenced (MIG)	0	.0
Single system pg reads	0	.0	Page blocks read	0	.0
Pages read from DASD	0	.0	Single guest page reads	0	.0
Spool pages read	0	.0	Pages written to DASD	0	.0
Total pgs to/from DASD	0	.0	Spool pages written	0	.0
<b>Minidisk Cache Management</b>					
MDC page reads	9	.2	MDC page writes	0	.0
MDC wait queue inserts	0	.0	MDC fair share exceeded	0	.0
Compl. successful reads	412	6.9	MDC blocks invalidated	0	.0
Total MDC read requests	412	6.9	Total invalidate req.	0	.0
<b>Memory Constraint Relief</b>					
Redrive attmps anywhere	0	.0	Task redrives anywhere	0	.0
Redrive attempts <2G	0	.0	Task redrives <2G	0	.0
Redrive attempts >2G	0	.0	Task redrives >2G	0	.0
<b>Command</b> ==>					
<b>F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return</b>					

Figure 8. Layout of System Counters Report (FCX102)

The SYSTEM report provides additional detailed information about the system's operation, especially in the areas of storage management and minidisk caching. It is intended to allow a more detailed analysis of information shown on other reports (CPU and STORAGE, for example), if you suspect a problem in one of these areas.

### Notes:

1. When running Performance Toolkit against z/VM 5.3 or earlier data, you will see dashes in all fields on this report.
2. When it is running against z/VM 5.4 data, Performance Toolkit can use trend files created using monitor data from any release; in which case, older data might cause dashes to appear in some fields. Note, however, that Performance Toolkit *cannot* use trend files that were actually created using older data for this report.

### **Description of SYSTEM Display Fields:**

#### *Operation*

This section describes the events that are monitored.

#### **Count**

The count in the selected interval, that is, during the last measurement interval or since the last RESET, for "current" or "average" data selection, respectively.

#### **Rate/s**

The number of occurrences per second in the selected interval.

#### **Events Shown:**

##### **Real SSCH operations**

The number of real SSCH operations (Start Subchannel) run by the system.

##### **Real CSCH instructions**

The number of real CSCH instructions (Clear Subchannel) run by the system.

##### **Real HSCH instructions**

The number of real HSCH instructions (Halt Subchannel) run by the system.

##### **El. time slice drops**

The number of elapsed time slice drops in the interval. A virtual machine's (or more correctly: a VMDBK's) stay in the dispatch list ends when it has used up its allocated time (the 'elapsed time slice'), or when the virtual machine's work has ended and it becomes idle.

##### **SVC instr. simulated**

The number of SVC instructions simulated by the system for guest operating systems.

##### **SVC interrupts reflectd**

The number of SVC interrupts reflected to guest systems.

##### **SVC 76 reflected**

The number of SVC 76 reflected to guest systems.

##### **Diagnose I/O requests**

The number of virtual I/O requests via Diagnose codes X'18', X'20', X'58', X'A4', and X'A8'. Other virtual I/O activity (for example, via SIO, SSCH) is not included in this count. See [“FCX100, CPU Load and Transactions – CPU”](#) on page 371, and [“FCX101/FCX181 Remote Performance Log – REDISP, REDHIST”](#) on page 377 for the total virtual I/O activity.

##### **FP external call simul.**

The number of fast path reflections of guest external call interrupts.

##### **FP partial executions**

The number of fast path partial execution intercepts.

##### **Fast-path SIGP simulat.**

The number of fast path simulations of SIGP external call instructions.

##### **FP simul. of Diag.X'44'**

The number of fast path simulations of Diagnose X'44'.

##### **Diag.X'9C', Not Ready**

Count of times a guest issued Diagnose X'9C' and the target was in a VMDSTATE < VMDREADY.

##### **Diag.X'9C', Wait State**

Count of times a guest issued Diagnose X'9C' and the target was found in wait state.

**Diag.X'9C', Soft Wait**

Count of times a guest issued Diagnose X'9C' and the issuer had to soft-wait on the target.

**Diag.X'9C', Alr.Dispat.**

Count of times a guest issued Diagnose X'9C' and the target was already dispatched.

**FP successful x-lates**

The number of successful fast-path CCW translations.

**CCW chains not FP-elig.**

The number of times a CCW chain was found to be not eligible for fast-path translation.

**Fast-path aborts**

The number of times a fast-path translate attempt was aborted.

**Total FP xlate attempts**

The total number of fast-path CCW translate attempts.

**Nr. of SIE executions**

The number of times the SIE operation (Start Interpretive Execution) was run.

**Nr. of SIE intercepts**

The number of times an SIE operation (Start Interpretive Execution) ended due to an intercept (privileged operation intercepted for handling by CP).

**Entries to enabled wait**

The number of guest machine entries to an enabled wait state.

**CP issued Diag.X'44'**

Count of Diagnose X'44's issued by CP.

**HPF Write**

Number of times a write channel program for a DASD device was presented for HPF DCW translation.

**HPF Done**

Number of times HPF DCW translation processing completed successfully for a DASD device.

*Storage Management Fields*

This block provides many fields with information on central and expanded storage management.

**Subpool FREE req bk <2G**

The number of requests for free storage which were satisfied from subpools below 2G.

**Total FREE req bk <2G**

The total number of requests for free storage backed below 2G (sysperm, guestperm, and regular subpooled).

**Subpool FREE req bk >2G**

The number of requests for free storage which were satisfied from subpools above 2G.

**Total FREE req bk >2G**

The total number of requests for free storage backed above 2G (sysperm, guestperm, and regular subpooled).

**Storage fast clears**

The number of fast path clears of real storage (requests for a new page which could be satisfied by a zeroed page without disk I/O).

**Avail lst frame req <2G**

The number of times that a new frame was requested from the <2G available list.

**Avail lst frame req >2G**

The number of times that a new frame was requested from the >2G available list.

**Available lst empty <2G**

The number of times a <2G frame request could not be satisfied from the <2G available list because it was empty. This number should be very low in relation to the total number of <2G frame requests – the available list replenishment task ("demand scan") will usually ensure that a sufficient number of page frames are made available, well before the frames are actually needed.

**Available 1st empty >2G**

The number of times a >2G frame request could not be satisfied from the >2G available list because it was empty. This number should be very low in relation to the total number of >2G frame requests – the available list replenishment task ("demand scan") will usually ensure that a sufficient number of page frames are made available, well before the frames are actually needed.

**Demand scan 1st pass**

The number of times the available list replenishment task completed after the first pass.

This routine *must* make a certain number of real storage page frames free. During the first pass, it will take only pages from dormant (that is, inactive) users and unreferenced pages from active users, which should have a minimal impact on system performance.

When running Performance Toolkit against z/VM 6.3 or later data, you will see dashes in this field.

**Demand scan 2nd pass**

The number of times the available list replenishment task completed after the second pass.

During the second pass of demand scan, referenced pages from active users (in eligible or dispatch list) will also be taken. Since these pages may be needed again soon, the frequency of occurrence should be low. High numbers indicate insufficient real storage for the current workload.

When running Performance Toolkit against z/VM 6.3 or later data, you will see dashes in this field.

**Demand scan emergency**

The number of times the available list replenishment task completed after the emergency pass.

During the emergency pass of demand scan, referenced shared storage, system pages, and even reserved pages from active users (in eligible or dispatch lists) are taken. These pages have to be taken to keep the system running, but they were probably frequently referenced and will have to be paged in again soon (thrashing). High emergency pass numbers indicate a bad real storage bottleneck.

When running Performance Toolkit against z/VM 6.3 or later data, you will see dashes in this field.

**Demand scan not satisf.**

The number of times the available list replenishment task ran without satisfying the demand for frames.

This indicates an even more severe storage problem than the 'emergency pass' counts.

When running Performance Toolkit against z/VM 6.3 or later data, you will see dashes in this field.

**Dispatch 1st pgs stolen**

The number of pages taken from users in the dispatch list (unreferenced and referenced pages).

When running Performance Toolkit against z/VM 6.3 or later data, you will see dashes in this field.

**Shared stor. pgs taken**

The number of shared segment pages taken by the available list replenishment function. SYSTEM and shared segment pages are taken only during the emergency pass of the demand scan, that is, the counts should be low for systems with sufficient real storage.

When running Performance Toolkit against z/VM 6.3 or later data, you will see dashes in this field.

**Pgs from dormant users**

The number of pages taken from dormant users.

When running Performance Toolkit against z/VM 6.3 or later data, you will see dashes in this field.

**Eligible 1st pgs stolen**

The number of pages taken from users in the eligible list (unreferenced and referenced pages).

When running Performance Toolkit against z/VM 6.3 or later data, you will see dashes in this field.

**Pages taken for FREE**

The number of pages taken from users by frame table scan.

**Fast PGINs from XSTORE**

The number of fast path PAGE IN operations (page moves from expanded storage to main storage).

**Slow PGINs from XSTORE**

The number of PAGE IN operations that did not take the fast path.

**PGOUTs main to XSTORE**

The number of pages moved out from main storage to expanded storage.

**XSTORE allocations**

The number of allocations of expanded storage frames.

**XSTORE releases**

The number of times that an XSTORE frame was released.

**Glbl cycl list searched**

The number of times the migration task has searched through the global cyclic list of logged-on users.

The MIGRATE task searches through the cyclic list in order to find XSTORE blocks to migrate out and thus free up XSTORE. It chooses blocks to migrate out based on the migrate target time (initially set equal to the XSTORE average age). If half the cyclic list is searched without finding enough migratable blocks, the target time is lowered and an additional search is made through the cyclic list (see also next value).

**Migr. target time reset**

The number of times the migration target time was reset. See also description of previous value.

**Migr thresh buf increas**

The number of times the migration threshold buffer was increased.

The migration threshold buffer is the buffer used for block paging migration blocks. Its size is initially set equal to the number of paging devices times a multiplication factor of 10. If the low threshold is raised, the multiplication factor is incremented by 1 and thus the size of the threshold buffer increased.

This is an indication of migration not being able to keep up with XSTORE demand. A high number may indicate concern.

**Migr thresh buf lowered**

The number of times that the migration threshold buffer was lowered.

This is an indication of migration working well.

**Page migr. from dormant**

The number of pages migrated from dormant users.

**Dormant with page migr.**

The number of dormant users which had pages migrated from them.

**Page migr. from active**

The number of pages migrated from active users.

**Active with page migr.**

The number of active users which had pages migrated from them.

**Shared pages migrated**

The number of pages migrated from shared systems.

**Shared sys with pg migr**

The number of shared systems which had pages migrated from them.

**Blocks migrated from CP**

The number of blocks migrated from CP (system owned pages).

**PGMBKs sel. during migr**

The number of PGMBKs selected while migrating. This is an indication of the efficiency of the migration task: the more pages can be migrated from a single PGMBK, the more efficient page migration works, that is, you should see this number in relation to the number of pages migrated.

**Blocks migrated to DASD**

The number of pages migrated to DASD (the total of the pages migrated from dormant and active users, shared systems, and CP)

**XSTORE migr invocations**

The number of times the migration task was invoked.

**No I/O for pg migration**

The number of times no I/O was necessary for migrating a page. If a page already resides on disk, and if it was not referenced since being read for the last time, it can simply be released.

**Pg not referenced (MIG)**

The number of pages found by the migrate task which have been read in but were not referenced before the migrate task selected them for migration. Block paging may read in many additional pages when a page fault occurs. The ones not used are counted here, that is, a low number is good.

**Pg not referenced (STL)**

The number of pages found by the steal function which have been read in but were not referenced before the steal function moved them out of main memory to either XSTORE or DASD.

**Page blocks read**

The number of pages read from DASD (block paging).

**Single system pg reads**

The number of single-page reads for CP.

**Single guest page reads**

The number of single-page reads for virtual machines.

**Pages read from DASD**

The number of pages read from paging DASD. This value includes read operations in progress.

**Pages written to DASD**

The number of pages written to paging DASD. This value includes write operations in progress.

**Spool pages read**

The number of pages read from SPOOL. This value includes read operations in progress.

**Spool pages written**

The number of pages written to SPOOL. This value includes write operations in progress.

**Total pgs to/from DASD**

The count of pages serviced (reads and writes) for paging and spool. This value is updated after the I/O operation completes. When page or spool I/O is outstanding, its value will be less than the sum of the preceding four counters (pages read from and written to paging DASD and SPOOL).

*Minidisk Cache Management Fields*

This block provides detailed information on the management of the minidisk cache.

**MDC page reads**

The number of CMS blocks moved from XSTORE to main storage.

Note that this is *not* the number of CMS read operations: read misses are also read operations, but they will not result in a MDC page being moved to main storage. Calculate the number of read operations by adding the number of read hits, partial read hits, and read misses (you can find this value on the general STORAGE display).

**MDC page writes**

The number of minidisk data blocks moved from main storage to XSTORE. This includes blocks which are written to XSTORE as the result of a WRITE operation *and* blocks which are written to XSTORE after a READ operation for a block which could not be found in expanded storage (read misses).

**MDC wait queue inserts**

The number of insertions into the waiting queue. The waiting queue keeps track of users that request data which are already on the way to the cache and so do not require another I/O operation, even though the blocks are not yet available in the minidisk cache.

**MDC fair share exceeded**

The number of times minidisk data blocks were not put into the cache because the fair share limit was exceeded.

**Compl. successful reads**

The number of MDC reads which were completely successful (all blocks found).

**MDC blocks invalidated**

The number of MDC blocks which were invalidated.

**Total MDC read requests**

The number of read requests to MDC.

**Total invalidate req.**

The total number of invalidate requests.

*Memory Constraint Relief delta counters*

This block provides detailed information on the Memory Constraint Relief fields.

**Redrive attmps anywhere**

Count of times HCPPTGRD attempted to redrive tasks waiting for an available frame that can be located anywhere.

**Task redrives anywhere**

Count of individual task redrives performed by HCPPTGRD for an available frame that can be located anywhere.

**Redrive attempts <2G**

Count of times HCPPTGRD attempted to redrive tasks waiting for an available frame that can be located only <2G.

**Task redrives <2G**

Count of individual task redrives performed by HCPPTGRD for an available frame that can be located only <2G.

**Redrive attempts >2G**

Count of times HCPPTGRD attempted to redrive tasks waiting for an available frame that can be located only >2G.

**Task redrives >2G**

Count of individual task redrives performed by HCPPTGRD for an available frame that can be located only >2G.

**Effect of AVERAGE and INTERIM subcommands:**

Overall period totals since last RESET will be shown for all Count fields; overall averages will be displayed for the Rate/s fields.

## FCX103, Storage Utilization – STORAGE

**Performance Toolkit Subcommand:** “STORAGE” on page 302.

The Storage Utilization report (FCX103):

```

FCX103      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor

Main storage utilization:
Total real storage      2'048GB
Total available         2'048GB
Offline storage frames      0
SYSGEN storage size      2'048GB
Shared storage          11'636KB
FREE stor. subpools       5'804KB
Subpool stor. utilization  66%
Locked pages            11605
Reserved user storage    256'260KB
Set reserved SYSMAX       0KB
Trace table              9'648KB
Pageable                 2'031GB
Storage utilization       43%
Tasks waiting for a frame  0
Tasks waiting for a page  0/s

Dynamic memory management:
Permanent available      1'024GB
Reconfigurable available  2'048GB
Standby real stor. size   0KB
Reservd real stor. size   0KB
Increment size            1'024MB

Paging / spooling activity:
Page moves <2GB for trans. ....s
Fast path page-in rate    ....s
Long path page-in rate    ....s
Long path page-out rate   ....s
Page read rate            153k/s
Page write rate           135k/s
Change page rewrites       9/s
Ref only page write bypass 17096/s
Changed page writes MLOAD  135k/s
Ref only page writes MLOAD 0/s
Page read blocking factor  25
Page write blocking factor  21
Migrate-out blocking factor ...
Paging SSCH rate          12770/s
Paging SSCH rate, reads   11449/s
Paging SSCH rate, writes  1322/s
Pages read per SSCH       13
Pages written per SSCH     103
Read Loc-Rec CCWs per SSCH 1
Write Loc-Rec CCWs per SSCH 1
SPOOL read rate           0/s
SPOOL write rate          1/s

Reorder Settings:
Reorder for System        Off

Memory Constraint Relief:
Pageable memory <2G       On
Pageable memory >2G       On
Demand scans <2G          On
Demand scans >2G          On
Allocate pageable freeze   ---
Multiplier before Min/Max  ---
Turnover Rate Mult <2G    ---
Turnover Rate Mult >2G    ---
Requests waiting anywhere   0
Requests waiting <2G       0
Requests waiting >2G       0

Command ===>
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

XSTORE utilization:
Total available          0KB
Att. to virt. machines   0KB
Size of CP partition     0KB
CP XSTORE utilization     ...%
Low threshold for migr.   ....kB
XSTORE allocation rate    .../s
Average age of XSTORE blks ...s
Average age at migration  ...s

MDCACHE utilization:
Min. size in XSTORE       0KB
Max. size in XSTORE       0KB
Ideal size in XSTORE      0KB
Act. size in XSTORE       0KB
Bias for XSTORE           .00
Min. size in main stor.   0KB
Max. size in main stor.   4'096MB
Ideal size in main stor.   49'152 B
Act. size in main stor.   40'960 B
Bias for main stor.       1.00
MDCACHE limit / user      1'956KB
Users with MDCACHE inserts 1
MDISK cache read rate     39/s
MDISK cache write rate    ....s
MDISK cache read hit rate  10/s
MDISK cache read hit ratio 25%

VDISKS:
System limit (blocks)     2147m
User limit (blocks)       0
Main store page frames    0
Expanded stor. pages      0
Pages on DASD             0

Paging / spooling settings:
PAGING ALIAS              Off
PAGING HPF                On
PAGING warning threshold   90%
KEEPSLOT option           Yes
EARLYWRITES option        Yes

Agelist:
Target size                41'597MB
Actual size                41'597MB
Revalidation rate          609/s
Reval post-write rate      0/s
Steal rate                 152k/s
Pages Evaluated            10%
Writes rate                135k/s

```

Figure 9. Layout of Storage Utilization Report (FCX103)



## Field Descriptions:

### *Main Storage Utilization:*

This section provides information about the size and usage of central processor storage.

#### **Total real storage**

The total main storage size, including all dynamic memory reconfigurations and any offline frames.

#### **Total available**

The total main storage size. This is the *real* storage size, including hardware system area (HSA) frames. Your system might actually use less storage, if less storage was defined during system generation.

#### **Offline storage frames**

The number of total main storage frames that are offline, including all dynamic memory reconfigurations.

#### **SYSGEN storage size**

The storage size defined for the z/VM system.

#### **Shared storage**

The amount of main storage that is currently used by shared segments.

#### **FREE stor. subpools**

The storage that is taken up by subpools (total doublewords available, expressed in kB).

#### **Subpool stor. utilization**

The percentage of subpool free storage that is currently in use.

#### **Locked pages**

The number of dynamic paging area (DPA) storage pages, consisting of pages that have been locked, including CP free storage space.

#### **Reserved user storage**

The total amount of storage reserved, in bytes, using the SET RESERVED command.

#### **Set reserved SYSMAX**

The maximum amount of storage that can be reserved, in bytes, using the SET RESERVED command.

#### **Trace table**

The storage that is occupied by CP trace tables.

#### **Pageable**

The size of the pageable part of the DPA (that is, the total DPA size minus the locked pages and minus the trace table).

#### **Storage utilization**

The percentage of pageable storage that is used up by the working sets of currently-active users.

#### **Tasks waiting for a frame**

The number of CP tasks that are waiting for a page frame.

#### **Tasks waiting for a page**

The number of CP tasks that were found to be waiting for a page, per second.

### *Dynamic memory management:*

This section provides information about the system memory configuration. If support for CP dynamic memory downgrade is not available, you will see dashes in some of these fields.

#### **Permanent available**

The amount of permanent memory that is configured for use on the system.

#### **Reconfigurable available**

The amount of reconfigurable memory that is configured for use on the system.

#### **Standby real stor. size**

The standby memory size for the partition, including all dynamic memory reconfigurations.

#### **Reservd real stor. size**

The reserved memory size for the partition, including all dynamic memory reconfigurations.

**Increment size**

The effective storage increment size. Memory can be added or removed from the configuration in multiples of this size.

*Paging / spooling activity:*

This section provides information about I/O operations performed by the paging subsystem, which include I/O operations to PAGE, SPOL, and DRCT extents, and to mapped minidisks.

**Page moves <2GB for trans.**

The rate at which pages were moved below the 2GB line for translation. The field is available only for z/VM systems running in 64-bit mode.

**Fast path page-in rate**

The rate at which pages were paged in using fast-path logic.

**Long path page-in rate**

The rate at which pages were paged in using long-path logic.

**Long path page-out rate**

The rate at which pages were paged out using long-path logic.

**Page read rate**

The average number of pages read per second (includes any reads from a paging slot).

**Page write rate**

The average number of pages written per second (includes any writes to a paging slot).

**Change page rewrites**

The rate of changed pages written to an old auxiliary slot during frame reclaim from the global age list. This counter applies only to changed pages that have fixed paging slots.

**Ref only page write bypass**

The rate of referenced-only pages plus PNRs that were not written to auxiliary slots during frame reclaim from the global age list because the data in auxiliary storage is still valid.

A PNR is a page read-in on a block that was read but was never referenced.

**Changed page writes MLOAD**

The rate of changed pages written to a new auxiliary slot during frame reclaim from the global age list due to load balancing. New slots are selected to optimize contiguous slot allocation and to spread writes across all paging volumes.

This counter includes initial writes to DASD, writes of changed pages, and, when KEEPSLOT=NO is in effect, it includes referenced-only pages plus PNRs (because KEEPSLOT=NO caused the release of the original slot when the page was read).

A PNR is a page read-in on a block that was read but was never referenced.

**Ref only page writes MLOAD**

The rate of referenced-only pages written to a new auxiliary slot during frame reclaim from the global age list because the volume where it currently resides is draining.

**Page read blocking factor**

The average number of pages read per page fault. In general, this value is equal to the page write blocking factor.

**Page write blocking factor**

The average size of a page block written to auxiliary storage. A page block could span multiple volumes.

**Migrate-out blocking factor**

The average blocking factor for migrating pages from expanded storage to DASD.

**Paging SSCH rate**

The number of I/O operations per second performed by the paging subsystem to CP-owned volumes.

**Paging SSCH rate, reads**

The number of read operations per second performed by the paging subsystem from CP-owned volumes.

**Paging SSCH rate, writes**

The number of write operations per second performed by the paging subsystem to CP-owned volumes.

**Pages read per SSCH**

The average number of pages read by the paging subsystem from CP-owned volumes per Start Subchannel (SSCH) instruction performing a read operation.

**Pages written per SSCH**

The average number of pages written by the paging subsystem to CP-owned volumes per SSCH performing a write operation.

**Read Loc-Rec CCWs per SSCH**

The number of read Locate-Record channel command words (CCWs) done per read SSCH from CP-owned volumes.

**Write Loc-Rec CCWs per SSCH**

The number of write Locate-Record CCWs done per write SSCH from CP-owned volumes.

**SPOOL read rate**

The rate at which pages were read from SPOOL devices.

**SPOOL write rate**

The rate at which pages were written to SPOOL devices.

*Reorder Settings:*

This section provides information about reorder settings.

**Reorder for System**

Indicates the system default reorder setting: `On` or `Off`. The system setting applies to the shared frame owned list and the user frame owned list of every virtual machine in the system, with the exception of those virtual machines for which a `SET REORDER` command was issued to override the system default.

*Memory Constraint Relief:*

This section provides information about memory constraint relief.

**Pageable memory <2G**

Pageable pages <2G. When `On`, pageable pages can be placed <2G. When `Off`, pageable pages cannot be placed <2G. Depending on storage conditions, this constraint can go on and off.

**Pageable memory >2G**

Pageable pages >2G. When `On`, pageable pages can be placed >2G. When `Off`, pageable pages cannot be placed >2G. Depending on storage conditions, this constraint can go on and off.

**Demand scans <2G**

Demand scans can be done <2G. When `On`, demand scans can be done for pageable pages <2G. When `Off`, demand scans will not be done for pageable pages <2G. This does not affect frame table scans. After this constraint goes on, it is not turned off.

**Demand scans >2G**

Demand scans can be done >2G. When `On`, demand scans can be done for pageable pages >2G. When `Off`, demand scans will not be done for pageable pages >2G. This does not affect frame table scans. After this constraint goes on, it is not turned off.

**Allocate pageable freeze**

When `On`, the decision on how to allocate pageable pages, as well as the turnover rate multiplier values, will no longer change.

**Multiplier before Min/Max**

The calculated <2G turnover rate multiplier before any minimum or maximum is applied.

**Turnover Rate Mult <2G**

The <2G turnover rate multiplier. This is used to adjust the available list high thresholds to influence steal rates. It is used during available frame comparisons to cause appropriate consumption of available frames above and below 2G.

**Turnover Rate Mult >2G**

The >2G turnover rate multiplier. This is used to adjust the available list high thresholds to influence steal rates. It is used during available frame comparisons to cause appropriate consumption of available frames above and below 2G.

**Requests waiting anywhere**

The count of deferred frame requests that are waiting for an available frame that can be located anywhere. This count plus Requests waiting <2G plus Requests waiting >2G is the total wait count.

**Requests waiting <2G**

The count of deferred frame requests that are waiting for an available frame that can be located only <2G. This count plus Requests waiting anywhere plus Requests waiting >2G is the total wait count.

**Requests waiting >2G**

The count of deferred frame requests that are waiting for an available frame that can be located only >2G. This count plus Requests waiting anywhere plus Requests waiting <2G is the total wait count.

*XSTORE utilization:*

This section provides information about the size and utilization of expanded storage (XSTORE).

**Total available**

The total size of expanded storage.

**Att. to virt. machines**

The size dedicated to specific virtual machines.

**Size of CP partition**

The XSTORE size that is available for CP use.

**CP XSTORE utilization**

The percentage of the CP partition that is currently in use.

**Low threshold for migr.**

The minimum XSTORE size that is to be kept free for allocation of new blocks. Used XSTORE blocks will be migrated to DASD if less free XSTORE space is available.

**XSTORE allocation rate**

The rate at which XSTORE pages were allocated.

**Average age of XSTORE blks**

The average age of XSTORE pages in the CP partition, in seconds.

**Average age at migration**

The average age of XSTORE pages at migration. If the average age at migration is low, and is close to the average age of all XSTORE blocks in storage, the size of the CP partition is probably too small.

*MDCACHE utilization:*

This section provides information about the size and utilization of the minidisk cache.

**Min. size in XSTORE**

The minimum size of the full-track cache area in expanded storage.

**Max. size in XSTORE**

The authorized ceiling for the minidisk cache.

**Ideal size in XSTORE**

The ideal size of the full-track cache area in expanded storage.

**Act. size in XSTORE**

The current XSTORE size used for the minidisk cache.

**Bias for XSTORE**

The bias for XSTORE usage of the full-track cache.

**Min. size in main stor.**

The minimum size of the full-track cache in main storage.

**Max. size in main stor.**

The maximum size of the full-track cache in main storage.

**Ideal size in main stor.**

The ideal size of the full-track cache in main storage.

**Act. size in main stor.**

The main storage actually occupied by the full-track cache.

**Bias for main stor.**

The bias for main storage usage of the full-track cache.

**MDCACHE limit / user**

The maximum amount of MDCACHE storage into which any one user can insert minidisk data in an interval. This is a fair-share limit that will prevent a user from monopolizing the minidisk cache, unless the machine is exempted by means of the NOMDCFS directory option.

**Users with MDCACHE inserts**

The number of users that had blocks inserted into the minidisk cache during the fair-share interval.

**MDISK cache read rate**

The number of read I/O operations per second to the minidisk cache.

**MDISK cache write rate**

The number of write operations per second to the minidisk cache.

**MDISK cache read hit rate**

The number of read operations per second to cached minidisks, where the block to be read was already in the cache.

**MDISK cache read hit ratio**

The percentage of read operations to cache minidisks, where the block to be read was already in the cache. The minidisk cache read hit rate and the read hit ratio are indications of the efficiency of caching.

*VDISKS:*

This section provides information about virtual disk system settings and overall storage utilization by virtual disks.

**System limit (blocks)**

The system limit set for maximum storage utilization by virtual disks.

**User limit (blocks)**

The maximum number of VDISK blocks that can be defined per user.

**Main store page frames**

The number of main storage page frames used by virtual disks.

**Expanded stor. pages**

The number of expanded storage blocks used by virtual disks.

**Pages on DASD**

The number of virtual disk pages that have been migrated to DASD.

*Paging / spooling settings:*

This section provides information about settings of the z/VM paging and spooling subsystems.

**PAGING ALIAS**

Indicates the setting of the PAGING\_ALIAS feature:

**On**

SYSTEM-attached HyperPAV aliases will be used for paging.

**Off**

SYSTEM-attached HyperPAV aliases will not be used for paging.

**Unavailable**

This feature is not available.

**PAGING HPF**

Indicates the setting of the PAGING\_HPF feature:

**On**

High Performance FICON (zHPF) channel programs will be used for paging, if the hardware is capable.

**Off**

zHPF channel programs will not be used for paging.

**Unavailable**

This feature is not available.

**PAGING warning threshold**

The allocation warning level for paging volumes, as specified on a PAGING system configuration statement or a previous SET PAGING command.

**KEEPSLOT option**

Indicates the setting of the KEEPSLOT option:

**Yes**

This option is enabled.

**No**

This option is disabled.

**EARLYWRITES option**

Indicates the setting of the EARLYWRITES option:

**Yes**

This option is enabled.

**No**

This option is disabled.

*Agelist:*

This section provides information about the state of the age list that is used by the frame replenishment algorithm.

**Target size**

The target amount of storage in the age list, in bytes.

**Actual size**

The current amount of storage in the age list, in bytes.

**Revalidation rate**

The rate of age list page revalidations.

**Reval post-write rate**

The rate of age list page revalidations after the pages were written to auxiliary storage.

**Steal rate**

The rate of page steals (that is, frame reclaims).

**Pages Evaluated**

The percentage of the age list that has been evaluated and written to auxiliary storage, if necessary.

**Writes rate**

The rate of the amount of storage written to auxiliary storage, in bytes per second.

**Usage Notes:**

1. Most of the fields in this report (including all fields under Paging / spooling settings) come from STORAGE SAMPLE records and will not be updated if those records were not enabled when the data was collected. See [“FCX149, Monitor Settings – MONSET” on page 515](#) to determine whether STORAGE SAMPLE record collection is active.
2. This report displays storage usage data for z/VM systems with expanded storage (XSTORE). If expanded storage is unavailable or unsupported, the XSTORE-related fields contain dots or 0s.

**Effect of AVERAGE and INTERIM subcommands:**

Overall period averages are shown for all fields where a calculated average has meaning. All other fields show the current value or state.

## FCX104, Privileged Operations Screen – PRIVOP

**Performance Toolkit Subcommand:** “PRIVOP” on page 257.

The Privileged Operations Screen (FCX104):

FCX104		CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor	
Operation		Count	Rate/s	Operation		Count	Rate/s
Diagnose X'00'		0	.0	Diagnose X'04'		347	5.8
Diagnose X'08'		2	.0	Diagnose X'0C'		18	.3
Diagnose X'10'		0	.0	Diagnose X'14'		0	.0
Diagnose X'18'		0	.0	Diagnose X'20'		0	.0
Diagnose X'24'		11	.2	Diagnose X'28'		0	.0
Diagnose X'34'		0	.0	Diagnose X'3C'		0	.0
Diagnose X'40'		0	.0	Diagnose X'44'		0	.0
Diagnose X'48'		0	.0	Diagnose X'4C'		0	.0
Diagnose X'54'		0	.0	Diagnose X'58'		2	.0
Diagnose X'5C'		0	.0	Diagnose X'60'		0	.0
Diagnose X'64'		0	.0	Diagnose X'68'		22	.4
Diagnose X'70'		0	.0	Diagnose X'74'		0	.0
Diagnose X'7C'		94	1.6	Diagnose X'80'		0	.0
Diagnose X'84'		0	.0	Diagnose X'88'		0	.0
Diagnose X'8C'		0	.0	Diagnose X'90'		0	.0
Diagnose X'94'		0	.0	Diagnose X'98'		211	3.5
Diagnose X'9C'		0	.0	Diagnose X'A0'		0	.0
.							
.							
.							
EPSW Extract PSW		0	.0	ESEA Extr/Set Sec Auth		0	.0
IPTE Inv. Pg Tab Entry		0	.0	ISK Insert Stor. Key		0	.0
ISKE Ins. Stor Key Ext		0	.0	IUCV total instr. cnt		1156	19.3
Reserved				RSCH Resume Subchannel		0	.0
SSCH Start Subchannel		104	1.7	SIO Start I/O		0	.0
SIOF Start I/O Fast		0	.0	TCH Test Channel		0	.0
LCTL Load Control (32)		0	.0	LCTLG Load Control (64)		0	.0
LPSW Load PSW		0	.0	LPSWE Load PSW Extended		0	.0
PGIN Page-In		0	.0	PGOUT Page-Out		0	.0
PTLB Purge TLB		0	.0	RRB Reset Ref. Bit		0	.0
RRBE Reset Ref. Bit Ex		0	.0	SCK Set Clock		0	.0
SCKPF Set Clock Progr.		0	.0	SERV Service Call		0	.0
SIGP Signal Processor		24	.4	SPX Set Prefix		0	.0
SSK Set Storage Key		0	.0	SSKE Set Stor. Key Ext		0	.0
SSM Set System Mask		0	.0	STAP Store CPU Address		0	.0
STCTG Store Contrl (64)		0	.0	STCTL Store Contrl (32)		0	.0
STFL Store Facil. List		0	.0	STIDP Store CPU ID		0	.0
STNSM Store Then AND SM		0	.0	STOSM Store Then OR SM		0	.0
STPX Store Prefix		0	.0	TB Test Block		0	.0
TPROT Test Protection		0	.0	SIE Start Int. Exec.		15	.3
BSA Branch/Set Auth		0	.0	ESSA Extr/Set Stg Atrs		0	.0
IDTE Inv. DAT Tab. Ent		0	.0	PTFF Perf Tim Fac Func		0	.0
STFLE Store Fac Lst Ext		0	.0	STSI Store System Info		0	.0
Command ==> _							
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return							

Figure 10. Layout of Privileged Operations Screen (FCX104)

The privileged operations screen shows counts and frequency of usage of many DIAGNOSE and other privileged instructions, issued by virtual machines, that had to be simulated by CP (that is, were not emulated by SIE microcode). Such instructions are one source of CP overhead; their counts and frequencies could help you when analyzing virtual machine workloads.

### Notes:

1. When running Performance Toolkit against z/VM 5.3 or earlier data, you will see dashes in all fields on this screen.
2. When it is running against z/VM 5.4 data, Performance Toolkit can use trend files created using monitor data from any release; in which case, older data might cause dashes to appear in some fields. Note, however, that Performance Toolkit *cannot* use trend files that were actually created using older data for this screen.



## **Description of Privileged Operation Fields:**

### **Operation**

This column describes the privileged operation.

### **Count**

The number of simulations in the selected interval, that is, during the last measurement interval or since the last RESET, for 'current' or 'average' data selection respectively.

### **Rate/s**

The number of simulations per second.

All the values shown are based on the sums of the respective counters accross all online processors of the system.

### **Diagnose Codes**

Refer to your operating system's reference manuals for a full description of the functions performed by the different diagnose codes.

#### **Diagnose X'00'**

Store extended identification code.

#### **Diagnose X'04'**

Examine real storage. Performance Toolkit uses this diagnose code to retrieve performance-relevant data from CP control blocks. Performance Toolkit (or other real-time monitors) will usually cause most of the load.

#### **Diagnose X'08'**

Virtual console function. Used to run CP commands from programs running in a virtual machine.

#### **Diagnose X'0C'**

Pseudo timer.

#### **Diagnose X'10'**

Release pages.

#### **Diagnose X'14'**

Input spool file manipulation.

#### **Diagnose X'18'**

Standard DASD I/O.

#### **Diagnose X'20'**

General 370 synchronous I/O.

#### **Diagnose X'24'**

Query device type and features.

#### **Diagnose X'28'**

Dynamic channel program modification.

#### **Diagnose X'34'**

Read system dump spool file.

#### **Diagnose X'3C'**

Activate CP directory.

#### **Diagnose X'40'**

Clean up after virtual IPL by device.

#### **Diagnose X'44'**

Voluntary time slice end.

#### **Diagnose X'48'**

Second level SVC 76.

#### **Diagnose X'4C'**

Generate accounting records.

#### **Diagnose X'54'**

Control the function of the PA2 key.

**Diagnose X'58'**

3270 virtual console full-screen interface.

**Diagnose X'5C'**

Error message editing by EMSG setting.

**Diagnose X'60'**

Determine virtual machine storage size.

**Diagnose X'64'**

Finding, loading, and purging a named saved segment.

**Diagnose X'68'**

Virtual machine communication facility (VMCF).

**Diagnose X'70'**

TOD clock accounting interface.

**Diagnose X'74'**

Saving and loading a 3800 image library file.

**Diagnose X'7C'**

Logical device support facility.

**Diagnose X'80'**

MSSF hardware call.

**Diagnose X'84'**

Directory update in place.

**Diagnose X'88'**

Validate user authorization/link minidisk.

**Diagnose X'8C'**

Access 3270 display device information.

**Diagnose X'90'**

Read symbol table.

**Diagnose X'94'**

Set up virtual machine dump block (VMDUMP).

**Diagnose X'98'**

Real I/O (for example, VTAM).

**Diagnose X'9C'**

Voluntary time slice yield.

**Diagnose X'A0'**

Obtain ACIGROUP name.

**Diagnose X'A4'**

Synchronous DASD I/O for standard CMS blocks.

**Diagnose X'A8'**

Synchronous general I/O to fully supported devices.

**Diagnose X'B0'**

Get protected application re-IPL data.

**Diagnose X'B4'**

Reading, writing, and erasing virtual printer extended attribute buffer.

**Diagnose X'B8'**

Spool file external attribute buffer manipulation.

**Diagnose X'BC'**

Open and query spool file characteristics.

**Diagnose X'C4'**

Handle class override data for UCR file.

When running Performance Toolkit against z/VM 6.3 or later data, the value in this field should always be 0.

**Diagnose X'C8'**

Set language.

**Diagnose X'CC'**

Save message repository.

**Diagnose X'D0'**

Tape volume serial support.

**Diagnose X'D4'**

Alternate user ID support.

**Diagnose X'D8'**

Read spool file descriptor blocks.

**Diagnose X'DC'**

Declare/delete virtual buffer for application data.

**Diagnose X'E0'**

System trace file interface.

**Diagnose X'E4'**

Return minidisk real device information (get information about another user's minidisk).

**Diagnose X'EC'**

Query GUEST trace status.

**Diagnose X'F0'**

CSE communication (information transmission between communication virtual machines).

When running Performance Toolkit against z/VM 6.3 or later data, the value in this field should always be 0.

**Diagnose X'F8'**

Spool file origin information.

**Diagnose X'FC'**

Channel path reconfiguration interface.

**Diagnose X'210'**

Retrieve device information.

**Diagnose X'214'**

CMS page management.

**Diagnose X'218'**

Access to real processor CPU ID.

**Diagnose X'238'**

Time based unique identifiers.

**Diagnose X'23C'**

Address space services.

**Diagnose X'240'**

Access list services.

**Diagnose X'244'**

Minidisk mapping services.

**Diagnose X'248'**

Copy to primary address space.

**Diagnose X'250'**

Synchronous/asynchronous block I/O.

**Diagnose X'254'**

Access real subsystem (issue a limited set of channel programs to select tape control unit subsystems).

**Diagnose X'258'**

Page reference services.

**Diagnose X'25C'**

Directory query (VMUDQ macro).

**Diagnose X'260'**

Access certain virtual machine information.

**Diagnose X'264'**

CMS/CP communication via communication area.

**Diagnose X'268'**

370 accommodation services.

**Diagnose X'26C'**

Access certain system information.

**Diagnose X'270'**

Pseudo timer extended.

**Diagnose X'274'**

Set timezone interrupt flag.

**Diagnose X'278'**

Extract XLINK control blocks.

**Diagnose X'27C'**

Product enablement verification.

**Diagnose X'280'**

Set POSIX IDs — security values.

**Diagnose X'288'**

Control virtual machine time bomb.

**Diagnose X'290'**

Perform privileged spool functions.

**Diagnose X'29C'**

Set-POSIX-IDs services.

**Diagnose X'2A0'**

Query POSIX IDs.

**Diagnose X'2A4'**

POSIX Process ID (PID) services.

**Diagnose X'2E0'**

SYSEVENT Query virtual server (QVS).

**Diagnose X'2FC'**

Obtain certain guest performance data.

**Simulated Instructions and Events**

See [\*z/Architecture® Principles of Operation\*](#) for a description of the functions performed by the different privileged instructions.

**EPSW**

Count of 'Extract PSW' instructions simulated

**ESEA**

Count of 'Extract and Set Extended Authority' instructions simulated

**IPTE**

Count of 'Invalidate Page Table Entry' instructions simulated

**ISK**

Count of 'Insert Storage Key' instructions simulated (370 mode virtual machines only)

**ISKE**

Count of 'Insert Storage Key Extended' instructions simulated

**IUCV total instr. cnt**

Total number of IUCV calls

**RSCH**

Count of 'Resume Subchannel' instructions simulated

**SSCH**

Count of 'Start Subchannel' instructions simulated

**SIO**

Count of 'Start I/O' instructions simulated (370 mode virtual machines only)

**SIOF**

Count of 'START I/O Fast' instructions simulated (370 mode virtual machines only)

**TCH**

Count of 'Test Channel' instructions simulated (370 mode virtual machines only)

**LCTL**

Count of 'Load Control' instructions simulated

**LCTLG**

Count of 'Load Control' instructions simulated for 64-bit mode

**LPSW**

Count of 'Load PSW' instructions simulated

**LPSWE**

Count of 'Load PSW Extended' instructions simulated

**PGIN**

Count of 'Page In' instructions simulated

**PGOUT**

Count of 'Page Out' instructions simulated

**PTLB**

Count of 'Purge TLB' instructions simulated (purge translation look-aside buffer)

**RRB**

Count of 'Reset Reference Bit' instructions simulated (370 mode virtual machines only)

**RRBE**

Count of 'Reset Reference Bit Extended' instructions simulated

**SCK**

Count of 'Set Clock' instructions simulated

**SCKPF**

Count of 'Set Clock Programmable Field' instructions simulated

**SERV**

Count of 'Service Call' instructions simulated

**SIGP**

Count of 'Signal Processor' instructions simulated

**SPX**

Count of 'Set Prefix' instructions simulated

**SSK**

Count of 'Set Storage Key' instructions simulated (370 mode virtual machines only)

**SSKE**

Count of 'Set Storage Key Extended' instructions simulated

**SSM**

Count of 'Set System Mask' instructions simulated

**STAP**

Count of 'Store CPU Address' instructions simulated

**STCTG**

Count of 'Store Control' instructions simulated for 64-bit mode

**STCTL**

Count of 'Store Control' instructions simulated

**STFL**

Count of 'Store Facility List' instructions simulated

**STIDP**

Count of Store CPU ID instructions simulated

**STNSM**

Count of 'Store Then AND System Mask' instructions simulated

**STOSM**

Count of 'Store Then OR System Mask' instructions simulated

**STPX**

Count of 'Store Prefix' instructions simulated

**TB**

Count of 'Test Block' instructions simulated

**TPROT**

Count of 'Test Protection' instructions simulated

**SIE**

Count of 'Start Interpretive Execution' instructions simulated

**BSA**

Count of 'Branch and Set Authority' instructions simulated

**ESSA**

Count of 'Extract and Set Storage Attributes' instructions simulated

**IDTE**

Count of 'Invalidate DAT Table Entry' instructions simulated

**PTFF**

Count of 'Perform Timing Facility Function' instructions simulated

**STFLE**

Count of 'Store Facility List Extended' instructions simulated

**STSI**

Count of 'Store System Information' instructions simulated

**Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages will be shown for all load fields.

## FCX105, Performance History Data Screens - GRAPHDET/PLOTDET, GRAPHSUM/PLOTSUM, GRAPHVAR/PLOTVAR

---

**Performance Toolkit Subcommands:** [“GRAPHDET / PLOTDET” on page 182](#), [“GRAPHSUM / PLOTSUM” on page 185](#), [“GRAPHVAR / PLOTVAR” on page 189](#).

See [“Graphics Selection Menu” on page 958](#) for more information about this specific screen and [Appendix A, “Graphical Displays of Performance History Data,” on page 955](#) for more general information.

## FCX106, CP IUCV Services Screen – CPSEVR

**Performance Toolkit Subcommand:** “CPSEVR” on page 158.

The CP IUCV Services Screen (FCX106) shows the activity for all CP IUCV services and also for VMCF communication.

FCX106	CPU 2817 SER 1EE75 Interval 18:13:38 - 19:38:38				Perf. Monitor			
IUCV Service	Transfers TO		Transfers BY		Unsuccessful		Total	X-fers
	Count	Rate/s	Count	Rate/s	Count	Rate/s	Count	Rate/s
*MSGALL	0	.0	1	.0	0	.0	1	.0
*MSG	0	.0	6085	1.2	0	.0	6085	1.2
*BLOCKIO	125M	24587	125M	24587	5	.0	251M	49174
*RPI	11259	2.2	11259	2.2	0	.0	22518	4.4
*MONITOR	176	.0	176	.0	0	.0	352	.1
*SIGNAL	44320	8.7	44320	8.7	0	.0	88640	17.4
*CCS	0	.0	0	.0	0	.0	0	.0
*SPL	0	.0	0	.0	0	.0	0	.0
*SYMPTOM	0	.0	0	.0	0	.0	0	.0
*ACCOUNT	0	.0	249	.0	0	.0	249	.0
*LOGREC	0	.0	33	.0	0	.0	33	.0
*CRM	0	.0	0	.0	0	.0	0	.0
*IDENT	0	.0	0	.0	0	.0	0	.0
*CONFIG	0	.0	0	.0	0	.0	0	.0
*VSWITCH	3	.0	3	.0	0	.0	6	.0
*ASYNCMD	0	.0	0	.0	0	.0	0	.0
*SCLP	0	.0	0	.0	0	.0	0	.0
*VMEVENT	0	.0	0	.0	0	.0	0	.0
Total	125M	24598	125M	24599	5	.0	251M	49197
Tot VM IUCV	126M	24791	126M	24792	1071	.2	253M	49583
Tot VMCF	8167	1.6	8167	1.6	0	.0	16334	3.2
Total X-fer	126M	24792	126M	24794	1071	.2	253M	49586
Command ==>								
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return								

Figure 11. Layout of CP IUCV Services Screen (FCX106)

The information shown is based on CP monitor SYSTEM domain SAMPLE data. Fields which indicate some activity will appear highlighted on the screen.

### Field Description:

#### IUCV Service

The name of the IUCV system service

- **\*MSGALL** - Message All System Service
- **\*MSG** - Message System Service
- **\*BLOCKIO** - DASD Block I/O System Service
- **\*RPI** - Access Verification System Service
- **\*MONITOR** - Monitor System Service
- **\*SIGNAL** - Signal System Service
- **\*CCS** - Console Communication Service
- **\*SPL** - Spool System Service
- **\*SYMPTOM** - Symptom System Service
- **\*ACCOUNT** - Account System Service
- **\*LOGREC** - Error Logging System Service
- **\*CRM** - Collection Resource Management System Service
- **\*IDENT** - Identify System Service
- **\*CONFIG** - Enterprise Systems Connection Manager System Service
- **\*VSWITCH** - Virtual Switch System Service



- **\*ASYNCMD** - Asynchronous CP Command Response System Service
- **\*SCLP** - SCLP System Service
- **\*VMEVENT** - VM Event System Service
- **Total** - Totals for all CP IUCV services
- **TOT VM IUCV** - Totals for Virtual Machine IUCV Service
- **TOT VMCF** - Totals for VMCF
- **Total X-fer** - Overall totals for both VM IUCV and VMCF activities

#### **Transfers TO**

IUCV transfer activity from virtual machines **to** the CP service

#### **Transfers BY**

IUCV transfer activity **by** the CP service to a virtual machine

#### **Unsuccessful**

Unsuccessful attempts for IUCV data transfers

#### **Total X-fers**

Total number of IUCV transfers for the respective service

The actual activity is shown in the form of two values under each of the above headings:

#### **Count**

the number of transfers in the relevant interval

#### **Rate/s**

the transfer rate, in transfers per second, during the interval

Note that each successful data transfer is associated with two counters. One for the target (transfer to) and one for the source (transfer by).

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all load fields.

## FCX107, Channel Load Screen – CHANNEL

**Performance Toolkit Subcommand:** “CHANNEL” on page 147.

In the Channel Load Screen (FCX107), fields which indicate that the channel has been busy will be highlighted on the screen.

FCX107	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor
<b>CHPID</b>	<b>Chan-Group</b>	<b>&lt;%Busy&gt;</b>	<b>&lt;----- Channel %Busy Distribution</b>	<b>04:24:40-04:</b>	
<b>(Hex)</b>	<b>Descr</b>	<b>Qual</b>	<b>Cur Ave</b>	<b>0-10 11-20 21-30 31-40 41-50 51-60 61-70 71-80</b>	
31	ESCON	00	31 20	17 33 27 23 0 0 0 0	
32	ESCON	00	31 20	17 30 30 23 0 0 0 0	
76	FICON	00	27 27	0 0 100 0 0 0 0 0	
34	ESCON	00	18 15	0 90 10 0 0 0 0 0	
11	ESCON	00	17 15	3 93 3 0 0 0 0 0	
33	ESCON	00	17 14	3 93 3 0 0 0 0 0	
10	ESCON	00	16 14	17 83 0 0 0 0 0 0	
0E	ESCON	00	14 8	63 30 7 0 0 0 0 0	
0F	ESCON	00	14 8	63 30 7 0 0 0 0 0	
66	FICON	00	12 12	0 100 0 0 0 0 0 0	
80	FICON	00	12 12	0 100 0 0 0 0 0 0	
77	FICON	00	11 11	0 100 0 0 0 0 0 0	
14	FICON	00	4 4	100 0 0 0 0 0 0 0	
15	FICON	00	4 4	100 0 0 0 0 0 0 0	
4A	FICON	00	4 4	100 0 0 0 0 0 0 0	
13	FICON	00	3 3	100 0 0 0 0 0 0 0	
04	FICON	00	1 1	100 0 0 0 0 0 0 0	

See also LCHANNEL and FCHANNEL displays  
 Command ==>  
 F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return

Figure 12. Layout of Channel Load Screen (FCX107)

Overall channel load analysis is based on data from the Extended Channel-Path Measurement Facility, where available; otherwise high-frequency sampling data from the CP monitor SYSTEM domain is used.

Note that the information extends beyond the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window so as to view the remaining data.

The screen shows a list of the active channels. Only channels for which some 'busy' conditions have been recorded will be included in the display, and the message:

No channel busy conditions logged so far

will be shown if no busy conditions were recorded at all.

Channels will be ordered in %Busy sequence, either for the *current*, i.e. the last interval's %Busy values if 'current' data are being displayed, or for *average* %Busy values when 'average' performance data have been selected. Channels with equal %Busy figures will be sorted in ascending CHPID sequence.

### Field Description for Channel Load:

#### CHPID (Hex)

The hexadecimal 'channel path identifier'.

#### Chan-Group Descr

Description of the channel model group, as found in the Store-Channel-Path description.

#### Chan-Group Qual

The hexadecimal channel model group qualifier. Note that a valid qualifier is available only on systems with the extended channel measurement facility.

#### %Busy

The percentage of 'busy' conditions found, where

#### Cur

Shows the 'current' figures, calculated for the last interval

**Ave**

Shows the 'average' load, calculated for the whole period since the last RESET, or since Performance Toolkit was last invoked.

For FICON channels the 'busy' value is calculated as the percentage of actual work unit utilization compared to the maximum possible value. Channel 'busy' percentages cannot be calculated for Hipersocket channels; the string n/a is inserted instead.

Under the heading:

```
Channel %Busy Distribution hh:mm:ss-hh:mm:ss
```

you will find the channel load distribution for each of the active channels. It shows the relative frequency with which the channel busy percentage of separate measurement intervals was found in either of the 10% ranges 0-10, 11-20, ..., 91-100.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Interim averages are displayed for the INTERIM setting, and overall averages are shown for both the AVERAGE and CURRENT settings since average data is shown together with the current load. Sorting of the channels is based on the average %Busy percentage instead of the last interval's %Busy values.

## FCX108, General I/O Device Screen – DEVICE or DEVICE DASD

**Performance Toolkit Subcommand:** “DEVICE” on page 169.

The General I/O Device Screen (FCX108) will be shown if you enter the DEVICE subcommand with no parameters at all or with the DASD parameter, in which case only disks (and no other devices) will be included in the display. Devices that have been active during the last measuring interval are highlighted on the screen.

FCX108		CPU nnnn		SER nnnnn		Interval HH:MM:SS - HH:MM:SS		Perf. Monitor												
<-- Device Descr. -->		Mdisk Pa-		<-Rate/s->		<----- Time (msec) ----->		Req. <Percent>		SEEK Recov		<-Throttle->		Devs/						
Addr	Type	Label/ID	Links	ths	I/O	Avoid	Pend	Disc	Conn	Serv	Resp	CUMt	Qued	Busy	READ	Cyls	SSCH	Set/s	Dly/s	SSID
>> All DASD <<			....	..	1.3	.0	.252	11.6	.744	12.6	22.0	.047	.01	2	84	1315	0	...	.0	3729
C50D	3390-9	ATP111 CP	0	4	22.4	.0	.252	11.6	.685	12.5	21.4	.000	.21	28	85	0	0	...	...	0607
CD05	3390-9	ATP215 CP	0	4	22.3	.0	.267	13.4	.694	14.4	28.3	.000	.30	32	85	0	0	...	...	060F
C707	3390-9	ATP133 CP	0	4	22.3	.0	.250	11.4	.708	12.4	23.1	.000	.25	28	85	0	0	...	...	0609
CD0A	3390-9	ATP220 CP	0	4	22.3	.0	.257	12.2	.700	13.1	25.3	.000	.27	29	85	0	0	...	...	060F
C701	3390-9	ATP127 CP	0	4	22.3	.0	.251	11.6	.656	12.5	24.1	.000	.27	28	85	0	0	...	...	0609
A677	3390-9	BWP102 ..	1	4	11.1	.0	.113	.001	.235	.349	.349	.000	.00	0	5	1410	0	...	...	0347
BE14	HPAV	.....	...	4	10.5	.0	.256	9.55	3.49	13.3	13.3	2.26	.00	14	60	0	0	...	...	0600
BE11	HPAV	.....	...	4	10.4	.0	.257	9.96	3.36	13.6	13.6	2.17	.00	14	62	0	0	...	...	0600
BE13	HPAV	.....	...	4	10.4	.0	.259	9.97	3.53	13.8	13.8	2.30	.00	14	61	0	0	...	...	0600
BE17	HPAV	.....	...	4	10.4	.0	.259	9.86	3.50	13.6	13.6	2.27	.00	14	61	0	0	...	...	0600
BE10	HPAV	.....	...	4	10.4	.0	.258	10.1	3.42	13.7	13.7	2.18	.00	14	61	0	0	...	...	0600
BE15	HPAV	.....	...	4	10.4	.0	.257	9.66	3.78	13.7	13.7	2.50	.00	14	61	0	0	...	...	0600
BE16	HPAV	.....	...	4	10.3	.0	.257	9.81	3.69	13.8	13.8	2.44	.00	14	61	0	0	...	...	0600
BE18	HPAV	.....	...	4	10.3	.0	.261	10.1	3.64	14.0	14.0	2.40	.00	14	61	0	0	...	...	0600
BE12	HPAV	.....	...	4	10.3	.0	.260	10.1	3.70	14.1	14.1	2.49	.00	14	61	0	0	...	...	0600
E936	3390-3	PFE936 ..	4	4	1.2	.0	.081	.000	.119	.199	.199	.000	.00	0	0	1230	0	...	...	033A
Command ==> _																				
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return																				

Figure 13. Layout of General I/O Device Screen (FCX108)

The information shown is based on CP monitor MONITOR domain and I/O domain SAMPLE data.

Figure 13 on page 412 shows a screen with I/O devices sorted in ascending order of the real device numbers, as indicated by the underscore characters above the Addr field. Other possible sort sequences are indicated by dots above the corresponding columns. These sequences can be set using the SORT subcommand with the string shown in the lowest line of the column heading as a sort argument (see “SORT” on page 295) or by positioning the cursor on the heading of the field that is to be sorted and then pressing ENTER.

Some additional lines with device group averages will be shown at the top when device groups have been defined. (See “FCONTROL DEVGROU” on page 26.) The device group name will be displayed in place of the device number and device type. The Devs/SSID field in a device group line contains the number of devices over which the row's statistics were calculated.

For more information about one specific I/O device (especially disks), perform one of the following actions:

1. Place the cursor on the Addr field of the device you are interested in and press ENTER
2. Enter the following subcommand on the command line:

```
DEVICE devno
```

Information about the selected device is then displayed on “FCX110, I/O Device Details Screen – DEVICE” on page 420.

### Field Descriptions:

*Device Descr.:*

#### Addr

The real device number.

The first line, which contains the prefix >> A11 DASD << instead of a device number, provides the average load figures for all disks for which performance data were collected.

### Type

The device type and model. Special cases are as follows:

- For PAV alias devices of an IBM DASD subsystem, the device type is replaced by the following string:

```
-> xxxx
```

where xxxx is the device number of the base RDEV.

- For HyperPAV aliases, the device type is replaced by the following string:

```
HPAV
```

### Label/ID

The label or ID, as follows (note that items 3, 4, and 5 apply to DASDs that are not attached to guests):

1. For a device that is attached to a guest, this field contains:

```
>userid
```

where *userid* is the user ID of the guest to which the device is attached.

2. For a device that is not a DASD, this field contains a label that consists of dots.
3. For a HyperPAV alias device, this field contains a label that consists of dots.
4. For a PAV alias device, this field contains the disk label of the base RDEV.
5. For all other DASDs, this field contains the disk label, and, if the device is in the CP-owned list, the label will have a suffix of CP.

### Mdisk Links

For DASDs that are:

- Attached to SYSTEM, this field contains the number of minidisk links.
- Not attached to SYSTEM, this field always contains a 0.

When the device is other than a non-PAV DASD RDEV, a PAV base RDEV, or a HyperPAV base RDEV, dots are displayed in this field.

### Paths

The number of available channel paths to the device. If the number of channel paths is low, and their usage is high, this could cause excessive reconnect misses, which would result in abnormally high disconnect times.

For emulated FBA devices (EDEV), this number is always 1, even when multiple channel connections to the device are available.

*Rate/s:*

### I/O

The number of SSCH and RSCH operations executed per second from this system during the selected interval.

A suffix of T will be appended to the I/O rate numbers of all devices with an I/O rate that is currently being "throttled".

### Avoid

The number of real I/O operations per second that were avoided because the blocks were found in the minidisk cache.

*Time (msec):*

**Pend**

The "function pending" time. This is the average amount of time that an I/O request remained pending in the I/O subsystem due to path-busy conditions.

This value and the Disc, Conn, Serv, Busy, and Resp values are calculated based on information obtained from the subchannel measurement block. This information is important especially for disks. Dots are inserted for devices where the information is not available or is not calculated.

**Disc**

The "device disconnected" time. This is the average amount of time that the device remained disconnected from the channel while executing I/O requests. For DASDs, it includes the time spent on executing the SEEK and SET SECTOR orders and also any re-connect misses. High values could indicate overloaded paths or many long SEEKS to data at opposite ends of the device.

The value also includes the "device-active-only" time where that information is available.

**Conn**

The "device connected" time. This is the average time that the device was connected to the channel while executing I/O requests, that is, receiving commands or actually transferring data. For DASDs, it also includes the SEARCH time.

**Serv**

The service time for the device. This is the sum of the function pending, connected, and disconnected times.

**Resp**

The device response time. This is the device service time plus the time during which an I/O request was waiting to be started.

**CUWt**

The control unit queuing time, that is, the additional wait time caused when the target device was already busy with an I/O operation that was started from another system.

**Req. Qued**

The average I/O queue length in the selected interval.

The value is taken from CP monitor high-frequency sample data; it is an indicator for I/O contention on the device.

The I/O queue length (Req. Qued) is not meaningful for devices that are dedicated (attached) to guest systems. The guest system will not try to start another I/O operation before an active one has ended, but rather will do its own queuing. Guest system queuing cannot be shown by host performance monitors, so they will have to be analyzed using a performance monitor on the guest system.

The queue length reported here is the number of real I/O operations (channel programs) that CP's real I/O layer is waiting to start to the device or volume. This includes I/O operations performed for all reasons: for minidisk I/O, for paging I/O, and for spooling I/O, for example. For devices in the CP-owned list, additional queueing information can be found on one of these screens:

- [“FCX109, CP-Owned Device Screen – DEVICE CPOWNED” on page 416](#)
- [“FCX170, CPOWNED Device Screen – CPOWNLOG” on page 553.](#)

*Percent:*

**Busy**

The percentage of time during which the device was busy.

**READ**

The percentage of I/O requests that were for read operations only.

**SEEK Cyls**

The average number of cylinders skipped per seek order.

**Recov SSCH**

The number of SSCHs (*not* the rate) that were executed for error recovery purposes.

*Throttle:*

**Set/s**

The throttle I/O rate set for the device.

Because an average or total of this value for all DASDs is not meaningful, this field contains dots on the >> All DASD << line.

**Dly/s**

The rate at which arriving I/O operations were delayed due to throttling.

**Devs/SSID**

The number of devices or the SSID, as follows:

1. Devs - On the >> All DASD << line or on a device-group line, this field contains the number of devices that were used in the calculation.
2. SSID - For DASDs that have an SSID, this field contains the SSID. Otherwise, this field contains dots.

**Usage notes:**

1. Statistics depicted on the rows of this screen are calculated using only the measurement data for the listed device itself. The effects of PAV or HyperPAV are not included. For a depiction of volume-level statistics that account for the effect of PAV or HyperPAV, use the FCX329 screen (VOLUME) or the FCX330 screen (VOLLOG).
2. When a column is sorted, only the Device Data lines are sorted. The >> All DASD << and Device Group Data lines remain unsorted at the top of the screen.
3. Lines for any non-DASD devices (such as OSA or 3270, for example) are not included in the >> All DASD << line calculation.

**Missing interrupt detection:**

When missing interrupt detection has been activated, the performance monitor checks the status of all non-CP-owned disks at the end of each monitor sample interval. An interrupt is considered to be missing if a disk was found to be busy during two consecutive measurements and no I/O interrupt was received in between.

If a missing interrupt is detected on a real device, the following message will be shown in the error message line at the bottom of the screen while in performance monitor mode:

```
FCXPEI354A Device 'xxxx' waiting for interrupt
```

If monitor data are collected permanently (using the FC MONCOLL ON command), message FCXPEI354A will also be generated and inserted into the basic mode output screen.

**Effect of AVERAGE and INTERIM subcommands:** Overall or interim period averages will be shown for all load fields. Real averages will be shown for all fields that display device load figures. The Label/ID, Mdisk Links, and Devs/SSID fields will always contain current information.

## FCX109, CP-Owned Device Screen – DEVICE CPOWNED

**Performance Toolkit Subcommand:** “DEVICE” on page 169.

The CP-Owned Device Screen (FCX109) is shown if you specify the CPOWNED parameter on the DEVICE subcommand:

FCX109	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor
Page / SPOOL Allocation Summary					
PAGE slots available		2642m	SPOOL slots available		8411760
PAGE slot utilization		0%	SPOOL slot utilization		0%
T-Disk space avail. (MB)		14152	DUMP slots available		23587k
T-Disk space utilization		0%	DUMP slot utilization		0%
<---Device Descr.--->					

Figure 14. Layout of CP-Owned Device Screen (FCX109)

The screen shown in the above figure is a sample for the type of display you can expect when analyzing CP-owned disks. Because the performance of paging and spooling I/O is a key factor of general system performance, this display is a good starting point when problems in that area are suspected.

Two blocks of data are shown:

- A summary of the total page and spool space and its utilization
- A detailed list with all system area definitions and their utilization and activity.

This information is based on CP monitor STORAGE domain sample data and on MONITOR domain configuration records. It can be displayed only when data collection for these domains is active.

The sample shows a screen with the system areas sorted in ascending device number sequence, which is the initial setting. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the SORT subcommand with the appropriate column heading (see “SORT” on page 295), or by positioning the cursor on the heading of the field which is to be sorted and then pressing ENTER.

For more detailed information on one specific CP-owned disk, do one of the following:

- Place the cursor on the Addr field of the device you are interested in and press ENTER
- Enter the subcommand DEVICE *devno* on the command line.

### Field Descriptions for CP-Owned Device Performance Data:

*Page / SPOOL Allocation Summary:*

#### PAGE slots available

The total number of slots allocated as PAGE areas on the disks that are currently online.

#### PAGE slot utilization

The utilization, as a percentage, of the available PAGE area.

#### SPOOL slots available

The total number of slots allocated as SPOOL areas on the disks that are currently online.



**SPOOL slot utilization**

The utilization, as a percentage, of the available SPOOL area.

**T-Disk space avail. (MB)**

The total number of megabytes allocated in TDSK areas on the system disks that are currently online.

**T-Disk space utilization**

The utilization, as a percentage, of the available TDSK area.

**DUMP slots available**

The total number of slots allocated as SPOOL areas on disks which have been dedicated to DUMPs

**DUMP slot utilization**

The utilization, as a percentage, of the available SPOOL area on DUMP disks.

*Device Descr.:*

**Addr**

The real device number of the disk.

**Devtyp**

Device type and model of the disk

**Volume Serial**

Disk label

**Area Type**

Describes what the system area is used for. The area type indicated can be:

**DIRECT**

For directory cylinders

**DUMP**

For a spooling area on DUMP disks

**PAGE**

For a paging area

**SPOOL**

For a general spooling area

**T-DISK**

For temporary minidisk space

**Area Extent**

Start and end cylinders for the system area, or start and end slot number for FBA disks. Where the cylinder or block number range would require more space than is available for this field, the number of slots in the area is inserted instead.

**Used %**

Percentage of slots currently in use in this area. In the case of TDSK and DIRECTORY space, the "used" percentage is based on the number of used cylinders, not blocks.

Valid information for the remainder of the performance fields can be shown only for PAGE, SPOOL, and DUMP areas.

*Rate/s:*

**P-Rds**

Average page read rate during the last interval. This field and the "Page Write" (P-Wrt) field are inserted only for the first page area on a specific disk, but show the *total* activity on *all* PAGE areas on the same disk pack.

**P-Wrt**

Average page write rate during the last interval.

**S-Rds**

Average spool read rate during the last interval. This field and the "Spool Write" (S-Wrt) field are inserted only for the first spool area on a specific disk, but show the *total* activity on *all* SPOOL areas on the same disk pack.

**S-Wrt**

The average spool write rate during the last interval.

**Total**

The total page and spool read and write activity during the selected interval.

**SSCH+RSCH**

The number of SSCH and RSCH operations executed per second from this system during the selected interval.

**I/O Interferes**

The number of I/O requests that had interference in the last interval (for current data) or the average number of I/O requests which had interference per monitor sample interval if average data are shown. Some examples of I/O interference are:

- CP I/O done to the volume on behalf of privileged CP commands, for example: QUERY PATHS, QUERY CACHE, HYPERSWAP, and SNAPDUMP.
- CP I/O done on behalf of MONITOR data collection, for example: Cache statistics in the Device Activity record.
- CP I/O done on behalf of CP error recovery, for example: State change processing.
- User I/O to non-paging/non-spool extents on the volume.

**Queue Lngth**

The number of page and spool reads and writes that are currently active or queued in CP's paging layer. This is a single sample if "current" data are being shown; it is the average of all samples since the last reset if "average" data are shown.

Note that this value reflects queue lengths for CP paging and spooling activity only; it does *not* indicate anything about the queueing of other I/O to the same device (that is, by users).

For more information about I/O queueing at the device, see the Req. Qued field of [“FCX108, General I/O Device Screen – DEVICE or DEVICE DASD”](#) on page 412.

**Serv Time/Page**

The smoothed service time for processing a page, in milliseconds (msec).

**MLOAD Resp Time**

The smoothed response time for CP I/O operations to this area, in milliseconds (msec).

The value is normalized to reflect the response time for servicing a single page. It also takes into account the PIOBK queue length for the device.

**Block Page Size**

The average number of pages in the blocks available on the device.

**%Used for Alloc**

The percentage of times the volume was chosen for allocation when looked at while choosing a volume.

**IOM**

The I/O protocol that is used to communicate with the device:

**C**

For command mode

**T**

For transport mode

**E**

For SCSI I/O

**Usage note:**

Volumes with multiple extents will result in multiple detail lines, one detail line per extent. The data for the behavior of the volume is displayed in the first detail line. Each subsequent detail line indicates only the size of the subsequent extent, with the remainder of the fields containing dashes (---).

**Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages will be shown for all load fields. Current values will be inserted for the Used % values of T-DISK and DIRECT areas, which are only extracted, for performance reasons, when the screen is actually selected for viewing or printing.

## FCX110, I/O Device Details Screen – DEVICE

**Performance Toolkit Subcommand:** “DEVICE” on page 169.

The I/O Device Details Screen (FCX110) will be shown if you specify a *devno* on the DEVICE subcommand, or by pressing ENTER after placing the cursor on the device number field on the General I/O Device Screen (FCX108), CP-Owned Device Screen (FCX109), or CU-Cached Disks Screen (FCX111).

This screen gives detailed information on the load of a single I/O device. Although intended primarily for analyzing the I/O load to different minidisks on SYSTEM-attached disks (online SEEKS analysis), it can also give some additional information for other devices.

FCX110	CPU nnnn	SER nnnnn	Interval	HH:MM:SS	- HH:MM:SS	Perf. Monitor
<b>Detailed Analysis for Device 0145 ( CP OWNED )</b>						
Device type :	3380-A	Function pend.:	4.8ms	Device busy :	23%	
VOLSER :	VMPG1C	Disconnected :	38.5ms	I/O contention:	7%	
Nr. of LINKs:	41	Connected :	7.3ms	Reserved :	0%	
Last SEEK :	524	Service time :	50.6ms	SENSE SSCH :	0	
SSCH rate/s :	4.6	Response time :	54.1ms	Recovery SSCH :	0	
Avoided/s :	.0	CU queue time :	.0ms	Throttle del/s:	...	
Status:	SHARABLE					
<b>System Page/Spool I/O Details</b>						
Page reads/s :	6.2	Total pages/s :	10.9	PG serv. time:	4.8ms	
Page writes/s :	4.7			PG resp. time:	4.8ms	
Spool reads/s :	.0	System I/Os /s :	3.4	PG queue len.:	.00	
Spool writes/s:	.0	User interfer./s:	.0	Avail. bsize :	14	
<b>Path(s) to device 0145:</b>						
Channel path status :	01	41				
	ON	ON				
<b>Device Overall CU-Cache Performance Split</b>						
DIR ADDR VOLSER	IO/S	%READ	%RDHIT	%WRHIT	ICL/S	BYP/S
8A 0145 VMPG1C	5.5	54	0	0	.0	.0
						'NORMAL' I/O only
<b>MDISK Extent Userid Addr IO/s VSEEK Status LINK MDIO/s</b>						
+-----+-----+-----+-----+-----+-----+-----+						
0071 - 0092	SNAMAIN	03CF	.0	0	WR	1 .0
0189 - 0204	SNAMAIN	06C3	.0	0	WR	1 .0
0205 - 0228	SNAMAIN	06C4	.0	0	WR	2 .0
0289 - 0298	SNAMAIN	07C2	.0	0	WR	1 .0
0299 - 0299	CHGCS	0191	.0	0	RR	6 .0
+-----+-----+-----+-----+-----+-----+-----+						
0300 - 0599	SYSTEM PAGE		RD/s	WR/s	MLOAD	Used IO/s
	LOAD	====>	6.2	4.7	6.1	47% ....
+-----+-----+-----+-----+-----+-----+-----+						
0600 - 0619	SNAMAIN	07C1	.0	0	WR	1 .0
0620 - 0664	CHVSCS	0193	.0	6	RR	4 .0
0821 - 0884	MAINT	0A95	.0		owner	
	XAACNT	0195	1.2	0	WR	1 1.2
+-----+-----+-----+-----+-----+-----+-----+						
<b>Command ==&gt; _</b>						
<b>F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return</b>						

Figure 15. Layout of I/O Device Details Screen (FCX110)

The screen in the above figure is a sample for the type of display you can expect when analyzing a disk containing active minidisks. It shows performance data for a 3380 disk which is connected to 3880-23 controllers, but for which caching has been disabled. Basically, the information shown is grouped into up to five blocks:

- A first block containing general device performance data. This block will always be shown, but its layout depends on the device type being analyzed.
- A second block with details on system I/O (to paging and spooling areas on the device, for example). Shown only when some system I/O activity has ever taken place.
- A third block containing information on the data path(s) available to the device. Always shown, always same layout regardless of device type.

- A fourth block containing information on the use made of the control unit cache. This information is actually extracted from FCX111, CU-Cached Disks Screen - CACHDBSE, and it can be shown only when the necessary prerequisites are met: I/O domain of CP MONITOR active, and data extraction from CP MONITOR enabled for the Performance Toolkit. See “CACHDBSE” on page 143 for a more detailed description of these requirements.
- A fifth block containing data for minidisk load analysis. This block will be shown only when SYSTEM-attached or -owned disks are analyzed which either contain minidisks that have been linked by a logged-on user, or which contain system areas (paging, spooling, or dump areas, for example). The line No active MDISKS on this volume - no SEEKS analysis possible will be displayed instead for all other disks.

Note that SEEKS analysis *is* possible, of course, but only based on CP monitor data: you will have to select one of the specific DASD seeks displays to obtain information on the distribution of I/O activity for such a disk (see “SEEKDET” on page 284 for more information).

The minidisk load analysis part of the data may extend beyond the number of lines that can be shown simultaneously on the screen, in which case the FORWARD and BACKWARD subcommands (or the corresponding PF-keys) can be used to view the remaining data.

The example shown is for a disk which is shared with another VM system; the minidisk cache cannot be used for I/O to minidisks on this pack.

Where use of the minidisk cache is possible, the minidisk load analysis part of the screen has a slightly different layout to allow displaying information on minidisk cache efficiency:

MDISK Extent	Userid	Addr	IO/s	VSEEK	Status	LINK	VIO/s	%MDC
0051 - 0060	BGF	0191	.0	0	WR		1	.0 ...
0121 - 0170	CAR	0319	6.2	0	RR			
	KSC	0319	3.7	0	RR			
	BERS	0319	1.4	0	RR			
	ALO	0319	.8	0	RR			
	FBL	0319	.5	0	RR		814	24.6 99
.... - ....	....							

Note that in this case, the information extends beyond the 80 columns of an ordinary screen. Use the LEFT/RIGHT subcommands, or the corresponding PF-keys, to shift the display window to view the remaining data.

## Field Description for General Device Performance Data:

### Device type

Real device type and model of the device found on the address being analyzed

### SSCH rate/s

Average number of SSCH and RSCH operations per second during the last measuring interval. Please note that the total number of SSCHs executed per second on a disk will usually not match exactly the total of the I/O rates displayed for all minidisks on the pack. This inconsistency is to be expected because:

- The values come from different sources and cannot be extracted at exactly the same time
- Some of the minidisks which contributed to the total may have been detached, or their users logged off, in the mean time, and
- I/O activity on CP areas (paging, spooling, etc.) cannot be shown.

### Status

General device status information. The status indicated can be

#### RESERVE HELD

A RESERVE to the real DASD is held from this system

#### RESERVE PEND

A RESERVE to the real DASD is pending from this system

**RELEASE PEND**

A RELEASE to the real DASD is pending

**MDCACHE USED**

Minidisk caching is used for at least one minidisk on this device

**SHARABLE**

The disk has been defined as 'SHARED' in the HCPRIO or configuration file, or by command

**HOT I/O DEV.**

Hot I/O device

**INTERV. REQD.**

Operator intervention is required on the real device

**RECORDING**

Intensive recording mode is active for the device

**DEV. OFFLINE**

The *device* is OFFLINE

**PATH OFFLINE**

All *path(s)* to the device are OFFLINE

**ONLINE**

Device ONLINE - no special status to be indicated. This status will be inserted only when no other status applies.

***for DASD only:*****VOLSER**

Disk label

**Nr. of LINKs**

The total number of LINKs to minidisks found on the disk pack. As for the total I/O rate above, this number may differ slightly from the LINK total found when adding up all the LINKs to each minidisk.

**Last SEEK**

The address of the last real SEEK executed for the disk. For FBA or SCSI disks, this is the block or address on which that last emulated I/O operation was performed.

**Avoided/s**

The rate at which I/O requests could be satisfied via the minidisk cache, without execution of real I/O operations.

***for GRAF devices only:*****Features**

Indicates extended data stream features available on the device. Features can be

**C**

Device has extended color support

**H**

Device has extended highlighting support

**PS**

Device has programmed symbols support

**Screen stat.**

Special GRAF status information. The status indicated can be

**DISA**

Disabled

**DISASQ**

Disabling sequence

**DROP**

DROP processing

**LOGON**

User logged on

**ENA**

Enabled

***for high speed devices only:*****Function pend.**

The average time, in milliseconds, during which an I/O operation remained pending in the I/O subsystem during the last measuring interval due to path busy conditions.

This value, and also the next five values, is calculated based on data from the 'host measurement facility'. Dots will be inserted for devices where this information is not available.

**Disconnected**

The average time, in milliseconds, during which the device remained disconnected while handling an I/O operation. Abnormally high 'disconnected' times (as in the example shown) could be caused by an overloaded channel path which did lead to many re-connect misses, or, for DASD, by long SEEKs due to I/O to data at opposite ends of the device.

**Connected**

The average time, in milliseconds, during which the device was connected to the channel path while handling an I/O operation (receiving commands or transferring data). This time also includes SEARCH time for DASD.

**Service time**

The average service time of the device, i.e. the average time required by the device for executing an I/O operation (in milliseconds). The service time is an indication of device performance only, it does not include queueing delays.

**Response time**

The average I/O response time for the device. The I/O response time is the service time plus any queueing delays, it is calculated from the service time and the average I/O request queue length.

**CU queue time**

The time when the device was logically disconnected from the channel subsystem while it was busy with an I/O operation initiated from another system.

The field will contain valid data only if the control unit queuing measurement facility is installed.

**Device busy**

Percent of elapsed time during which the device was busy with an I/O operation. This value is calculated based on the I/O rate of the device and its service time, using the data provided by the 'host measurement facility'.

**I/O contention**

The smoothed relative I/O queue length, in percent. The value is based on the number of IORBKs found queued during the last few measurements. (see [“DEVICE” on page 169](#)).

This value should be very low, it indicates an I/O performance problem because of an overloaded device (too many users having high-activity minidisks allocated on the same real device) if it is often high.

Please note that the contention is calculated based on one sample per measurement interval. The resulting contention figures should be taken as rough indicators only.

**Reserved**

Percentage of measurements during which the device was found to have been 'reserved' by real reserve/release logic. This value is a smoothed average, based on the last few samples, and it is inserted only for DASD.

**SENSE SSCH**

The number of 'SENSE' SSCH operations during the last measuring interval. This should normally be zero: SENSE operations are usually performed when a previous I/O operation has terminated with an error condition (unit check).

**Recovery SSCH**

The number of recovery SSCH operations during the last measuring interval.

**Throttle del/s**

The rate at which I/Os to the device were delayed due to throttling (applicable only when throttling set on).

**Field Descriptions for System Page/Spool I/O Details:**

These fields describe the CP I/O activity to system areas (e.g. paging and spooling areas), and to minidisks which have been mapped for paging I/O to VM data spaces. The fields shown and their meaning are

**Page reads/s**

The number of pages read per second

**Page writes/s**

The number of pages written per second

**Spool reads/s**

The number of spool pages read per second

**Spool writes/s**

The number of spool pages written per second

**Total pages/s**

The total number of pages moved per second

**System I/Os /s**

The number of actual SSCH or RSCH operations per second for system areas

**User interf./s**

The rate at which user-generated I/Os caused interference with system-generated I/O activity (system I/O program could not be resumed)

**PG serv. time**

The smoothed service time per PIOBK (page I/O CCW package), in milliseconds. Continuously updated by CP after each I/O.

**PG resp. time**

The smoothed response time per PIOBK. Based on the smoothed service time, but taking into account also the PIOBK queue length. CP uses this value when selecting suitable disks for the next paging or spooling write operation.

**PG queue len.**

The smoothed PIOBK queue length. This value is determined by sampling the actual current queue length, and then smoothing it according to the formula

$$(15 * \text{old value} + \text{current value}) / 16$$

**Avail. bsize**

The average number of consecutive free slots found during the last measuring interval. Dots will be inserted when no new slot sizes have been determined in an interval.

**Field Description for Path Information:**

The path information section shows the channel paths by which the I/O device is connected to the system, based on the specifications in the IOCP.

For the emulated FBA devices (EDEV), only the first path ID is shown, even when multiple channel connections to the device are available.

The format is:

<b>Path(s) to device nnnn:</b>	<b>p1</b>	<b>p2</b>	<b>p3</b>	<b>p4</b>
Channel path status :	xxx	xxx	xxx	xxx

where:

**nnnn**

Is the real device number



**p1 p2 ..**

Are the path numbers of all generated paths to the device

**xxx**

Is the channel path status. It can be either 'ON' for 'ONLINE', or 'OFF' for 'OFFLINE'.

### **Field Descriptions for Cache Usage:**

The output shown is an extract of the data displayed on FCX111, CU-Cached Disks Screen - CACHDBSE. Refer to [“CACHDBSE” on page 143](#) for the corresponding field descriptions.

Please be aware that cache performance data are asynchronously collected by the CP MONITOR I/O domain. Since the data collect interval of the CP MONITOR is different from the one used for collecting detailed device data, the I/O rates shown will usually be different too. See also the **Usage notes** for FCX111, CU-Cached Disks Screen - CACHDBSE for further guidelines.

### **Field Descriptions for Minidisk Load Analysis:**

Minidisks found on the device which belong to, or are linked by currently logged on users will be displayed in ascending sequence of start cylinder or block. Minidisks for which any I/O activity was found during the last measuring interval will be highlighted.

Since both the total I/O count to each minidisk and also the individual contribution to this total by individual users are shown, this display can be used as a very powerful tool for detecting the reason for high I/O contention on a specific device.

### **MDISK Extent**

Start and end cylinders for minidisks on CKD disks. For minidisks on FBA devices the numbers of the first and last block of the minidisk are inserted instead.

### **Userid**

User identifications of users linked to the minidisk. The first entry shown for any minidisk is always the minidisk owner, regardless of I/O activity. Where other users are linked to the same minidisk, up to four of them will be shown if they have contributed to the minidisk's I/O activity in the preceding measuring interval, in descending order of I/O activity.

### **Addr**

Virtual minidisk address

### **IO/s**

Number of I/O requests per second to this particular minidisk (to this virtual address of this user)

### **VSEEK**

Last virtual SEEK address. The seek address has only little meaning for CMS minidisks where the user usually does not know where specific files have been allocated, and where the address may not even be updated for normal CMS file I/O. It provides an indication on the dataset being processed if the minidisk is used by a VSE or MVS system.

### **Status**

Minidisk status information. The status information consists of two parts. The first part indicates how the minidisk has been linked by the user: the characters 'WR' are shown for minidisks linked in write mode and the characters 'RR' for minidisks in read-only mode (always shown). The string 'owner' will be inserted instead for the first entry of minidisks which are in use, but whose owner does not have them linked.

The second part of the status information is inserted only if the minidisk was found to be in a special status at the end of the measuring interval. The status indicated can be:

### **BUSY**

Indicates that the minidisk was busy with an I/O operation when the last sample was taken

### **SUSP**

Indicates that an I/O operation has been suspended

### **RPEND**

Indicates that a 'resume SSCH' operation was pending

**QUED**

Indicates that an I/O operation was queued for the minidisk

**RS/RL**

Indicates that the minidisk has been defined in the directory with virtual reserve/release support

**TDISK**

Indicates that the minidisk has been defined as a temporary minidisk

Only the first status found (from top to bottom of the above list) will be inserted, even if more than one status applies.

**LINK**

Number of LINKs to the minidisk

**VIO/s**

Total number of *virtual* I/O requests per second to the minidisk. This field is shown only for disks which can use the minidisk cache.

**%MDC**

Minidisk cache efficiency, shown as the percentage of virtual I/O requests which did not require a real I/O operation. This field is shown only for disks which can use the minidisk cache; dots will be inserted when the value cannot be calculated because the virtual I/O rate is 0.

Although not absolutely accurate (there is no fixed 1:1 relation between virtual and real I/O operations to a disk, especially to non-CMS minidisks, even if no minidisk cache is available), this value should be a good indicator of minidisk cache efficiency for a specific CMS minidisk.

**MDIO/s**

Total number of real I/O requests per second to the minidisk

**Cache-eligible minidisks:**

If a minidisk has been defined in the directory as eligible for caching in the expanded storage minidisk cache (the default, valid unless a 'CACHOPT' statement specifying 'NOMDC' follows the minidisk entry), the character 'C' will be used as left and right delimiter instead of the standard '|'.

If the minidisk cache is used for a minidisk (less real I/Os measured than there were virtual I/Os, i.e. %MDC value > 0), these delimiters and the extent information for the minidisk will be highlighted.

**Additional fields for system areas on the disk:**

System area fields are separated from each other and from general user minidisk entries by a dashed line.

The layout for system areas is as follows:

```
+-----+
| 0412 - 0471   SYSTEM  type   RD/s  WR/s  MLOAD  Used  IO/s |
|              LOAD    =====>  4.1   2.3   16.5   74%   6.4 |
+-----+
```

**SYSTEM type**

The area type can be:

**DIRECT**

For CP directory space (allocation DRCT)

**PAGE**

For a paging area (allocation PAGE)

**SPOOL**

For a spooling area (allocation SPOL)

**TDSK**

For a temporary minidisk area (allocation TDSK)

The following three fields contain meaningful data only if the system area is of the type PAGE or SPOOL:

**RD/s**

Page or SPOOL read rate during the last interval (number of PIOBKs read per second)

**WR/s**

Page or SPOOL write rate during the last interval (number of PIOBKs written per second)

**MLOAD**

The smoothed response time for CP I/O operations to this area, in milliseconds (msec).

The value is normalized to reflect the response time for servicing a single page. It also takes into account the PIOBK queue length for the device.

**Used**

Percentage of blocks (PAGE and SPOOL areas) or of cylinders (DIRECT and TDSK) currently in use in this area.

**IO/s**

Information for calculating I/O rates to system areas directly is not available. The I/O rate will be inserted for SPOOL areas as the sum of the read and write request rate, and for PAGE areas as the difference between the sum of all read and write requests and the number of requests which could be appended to an existing channel program without a separate SSCH/RSCH. Where more than one PAGE or SPOOL area exists on the same disk pack, the load information is inserted for the first area only (no information is available to tell how the load is split). Dots are inserted for the I/O rate in further PAGE or SPOOL areas, and for all other types of system areas.

See also the system I/O details section of this screen for information on I/O handling to system areas on the disk.

For temporary minidisk areas you will see both a system area entry, specifying the total size of the TDSK space available, and ordinary minidisk entries for each of the temporary minidisks which are currently defined.

**Effect of AVERAGE and INTERIM subcommands:** None. Averages cannot be shown for the detailed I/O device display.

## FCX111, CU-Cached Disks Screen – CACHDBSE

**Performance Toolkit Subcommand:** “CACHDBSE” on page 143.

The CU-Cached Disks Screen (FCX111) shows performance data for disks connected to a cached 3880-13 or 3880-23 control unit or an IBM DASD subsystem, and it gives information on the use made and the effectiveness of the cache. Devices which have been active during the last measuring interval will be highlighted on the screen.

FCX111			CPU nnnn		SER nnnnn	Interval		HH:MM:SS - HH:MM:SS		Perf. Monitor		
Device			Overall CU-Cache Performance					Split				
DIR	ADDR	VOLSER	IO/S	%READ	%RDHIT	%WRHIT	ICL/S	BYP/S	IO/S	%READ	%RDHIT	
02	2000	530RES	.6	4	94	81	.0	.0	.2	5	83	(N)
									.4	4	100	(S)
									.0	...	...	(F)
02	2001	530SPL	.0	0	...	100	.0	.0	No SEQ./	CACHE	FW	
02	2002	530PAG	.0	100	100	...	.0	.0	'NORMAL'	I/O only		
02	2003	530W01	.0	100	66	...	.0	.0	'NORMAL'	I/O only		
02	2004	530W02	.0	...	...	...	.0	.0	'NORMAL'	I/O only		
02	2005	530PG1	.0	...	...	...	.0	.0	'NORMAL'	I/O only		
02	2006	530PG2	.0	100	0	...	.0	.0	'NORMAL'	I/O only		
02	2007	530PG3	.0	33	0	100	.0	.0	No SEQ./	CACHE	FW	
02	2008	530PG4	.0	...	...	...	.0	.0	'NORMAL'	I/O only		
02	2009	530PG5	.0	100	0	...	.0	.0	.0	100	0	(N)
									.0	100	0	(S)
									.0	...	...	(F)
02	200A	LNXC07	.4	0	...	91	.0	.0	.2	0	...	(N)
									.2	0	...	(S)
									.0	...	...	(F)
02	200B	LNW02	.0	0	...	60	.0	.0	.0	0	...	(N)
									.0	0	...	(S)
									.0	...	...	(F)
02	200C	LNW03	.0	0	...	50	.0	.0	No SEQ./	CACHE	FW	
02	200D	LNW04	.0	0	...	44	.0	.0	No SEQ./	CACHE	FW	
Select a device for I/O device details												
Command ==>												
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return												

Figure 16. Layout of CU-Cached Disks Screen (FCX111)

The information shown is based on CP monitor MONITOR and I/O domain data.

Figure 16 on page 428 shows a screen with performance information for 3380 disks that are connected over two paths to storage directors of two different 3880-23 control units. Only one line per device is shown, containing the sum of the relevant data from both 3880-23 control units. Output lines are sorted in ascending order of the DASD device numbers, as indicated by the underscore characters above the ADDR heading. Other possible sort sequences are indicated by dots above the corresponding columns. These sequences can be set using the SORT subcommand with the column heading (see “SORT” on page 295) or by positioning the cursor on the heading of the field that is to be sorted and then pressing ENTER. The cache has not been activated for all disks, hence the ratio of 0 for some of them.

For more information about the load of one specific disk, perform one of the following actions:

1. Place the cursor on the ADDR field of the device you are interested in and press ENTER
2. Enter the following subcommand on the command line:

```
DEVICE devno
```

A cached 3880-13 or 3880-23 control unit accepts both NORMAL and SEQUENTIAL I/O requests, and an IBM DASD subsystem can additionally handle FAST WRITE requests. Separate counts are kept by the control units and collected by the CP MONITOR for these modes, so separate performance information can be shown. The first few output line columns (with the heading Overall CU-Cache Performance) are summaries which include all request types. The last three columns with the heading 'Split' are included to indicate the contribution of each of the different request types to the total load. Where only 'NORMAL' I/O requests have been executed, as in the example above, no load split is shown. The string

'NORMAL' I/O only (for 3880) No SEQ./ CACHE FW (for IBM DASD subsystems)

will be inserted for all devices where no SEQUENTIAL or CACHE FAST WRITE I/O requests were made during the interval shown. The string:

initial sample

and dots will be inserted for devices where no valid delta values could be found (initial display, or not valid counts found).

If there was any SEQUENTIAL or CACHE FAST WRITE I/O activity on the device, the performance data for NORMAL, SEQUENTIAL and, for IBM DASD subsystems, CACHE FAST WRITE I/O will be inserted separately as shown in the following example:

FCX111	CPU	nnnn	SER	nnnnn	Interval	hh:mm:ss	-	hh:mm:ss	PERF.	MONITOR
Device		Overall CU-Cache Performance			Split					
DIR	ADDR	VOLSER	IO/S	%READ	%RDHIT	%WRHIT	ICL/S	BYP/S	IO/S	%READ %RDHIT
xx	xxxx	xxxxxx	n.n	nnn	nnn	nnn	n.n	n.n	n.n	nnn (N)
									n.n	nnn (S)
									n.n	nnn (F)
xx	xxxx	xxxxxx	n.n	nnn	nnn	nnn	n.n	n.n	n.n	nnn (N)
									n.n	nnn (S)
									n.n	nnn (F)

The strings (N) or (S) on the right indicate that the data inserted under the 'Split' heading is for NORMAL or SEQUENTIAL I/O respectively. For an IBM DASD subsystem a third line for CACHE FAST WRITE will be added, with the designator (F) on the right.

#### Field Description for Overall Performance:

The following field descriptions apply to both the CACHDBSE and CACHDLOG *devno* displays, with the exception of the DIR and ADDR fields which will be replaced by the field Interval End Time on the CACHDLOG display.

##### DIR

3880-13/23 storage director ID. For IBM DASD subsystems the last two hexadecimal digits of the SSID will be inserted.

##### ADDR

Disk device number

##### VOLSER

Disk label

##### IO/S

Sum of NORMAL, SEQUENTIAL and FAST WRITE (for IBM DASD subsystems only) I/O requests per second (READ + WRITE)

##### %READ

Overall percentage of READ requests (NORMAL, SEQUENTIAL and FAST WRITE) to the total I/O activity of the device.

Note that dots will be inserted when this value is not relevant.

##### %RDHIT

Overall percentage of READ hits, i.e. the percentage of READ requests which did not need DASD access (NORMAL, SEQUENTIAL and FAST WRITE).

Note that dots will be inserted when this value is not relevant.

##### %WRHIT

Ratio of WRITE hits to WRITE requests. On 3880 control units this refers to 'write normal' hits, and on IBM DASD subsystems to DASD fast write hits.

Note that dots will be inserted when this value is not relevant.

**ICL/S**

Number of INHIBIT CACHE LOAD requests per second

**BYP/S**

Number of BYPASS CACHE requests per second

**Field Description for Normal / Sequential / Fast Write Split:****IO/S**

Number of I/O requests per second over this storage director.

**%READ**

Percentage of READ requests.

Note that dots will be inserted when this value is not relevant.

**%RDHIT**

READ percentage, i.e. the percentage of READ requests which did not need DASD access.

Note that dots will be inserted when this value is not relevant.

Where separate 'normal', 'sequential' and 'cache fast write' load information is available the strings:

**(N)**

For normal I/O

**(S)**

For sequential I/O

**(F)**

For cache fast write I/O

will be appended to the end of split lines.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all load fields.

## FCX112, User Resource Usage – USER

**Performance Toolkit Subcommand: “USER” on page 346.**

The User Resource Usage report (FCX112), which shows general information about resource usage by individual users, is displayed if you issue the USER subcommand with no *userid* parameter. Output lines with data of users that have been active during the last interval are highlighted on the screen.

FCX112	CPU nnnn				SER nnnnn	Interval HH:MM:SS - HH:MM:SS				Perf. Monitor										
-----	<----- CPU Load ----->				<----- Virtual I/O/s ----->				<----- User Time ----->				<----- Spool----->				<----- VConf ----->			
-----	<----- Seconds----->				<----- Diag ----->				<----- Minutes----->				<----- Pages ----->				<----- Scat. ----->			
-----	UR	Pg/s	User	Status	Logd	Act.	Pages	SPg/s	MDC/s	Share	Scat.	Nr of								
Userid	%CPU	TCPU	VCPU	Ratio	Total	DASD	Avoid	98	UR	Pg/s	User	Status	Logd	Act.	Pages	SPg/s	MDC/s	Share	Scat.	Nr of
>>Mean>>	.09	.051	.031	1.63	.0	.0	.1	.0	.0	.0	---	---	1.0	.4	.0	.0	.0	---	...	39
User Class Data:																				
FTP	.00	.000	.000	....	.0	.0	.0	.0	.0	.0	---	---	1.0	1.0	.0	.0	.0	---	...	2
User Data:																				
CFT2ND	2.88	1.730	1.124	1.54	.0	.0	.0	.0	.0	.0	EME,CL3,DISP	1	1	.0	.0	.0	100	.478		
CFT2NDA	.00	.001	.001	1.00	.0	.0	.0	.0	.0	.0	EME,---,DORM	1	1	.0	.0	.0	100	.000		
CHG00007	.04	.026	.019	1.37	.5	.3	1.7	.0	.0	.0	ESA,---,DORM	1	1	.0	.0	.0	100	.000		
DGA41	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
DGA42	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
DGA43	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
DGA44	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
DTCVSW1	.00	.000	.000	.00	.0	.0	.0	.0	.0	.0	ESA,---,DORM	1	1	.0	.0	.0	100	.000		
DTCVSW2	.00	.000	.000	.00	.0	.0	.0	.0	.0	.0	ESA,---,DORM	1	1	.0	.0	.0	100	.000		
DWPLEXA1	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	STOP	...		
FTPSEERVE	.00	.000	.000	.00	.0	.0	.0	.0	.0	.0	XC,---,DORM	1	1	.0	.0	.0	100	.000		
FTPSERV50	.00	.000	.000	.00	.0	.0	.0	.0	.0	.0	ESA,---,DORM	1	1	.0	.0	.0	100	.000		
GCS	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
MISCSERV	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
MPROUTE	.07	.043	.031	1.39	.0	.0	.0	.0	.0	.0	EME,CL0,DISP	1	1	.0	.0	.0	3000	.000		
OPERATNS	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
OPERATOR	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
OPERSYMP	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
PERFSVMA	.15	.088	.021	4.19	.4	.2	.0	.0	.0	.0	ESA,---,DORM	1	1	.0	.0	.0	3.0%A	.000		
PORTMAP	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
PVM	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	3000	...		
REXECD	.00	.000	.000	.00	.0	.0	.0	.0	.0	.0	ESA,---,DORM	1	1	.0	.0	.0	100	.000		
RSCS	.00	.001	.000	.00	.1	.0	.0	.0	.0	.0	ESA,---,DORM	1	1	.0	.0	.0	100	.000		
RSCSDNS	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
RSTL3	.15	.089	.021	4.24	.7	.7	3.5	.0	.0	.0	ESA,---,DORM	1	1	.0	.0	.0	3.0%A	.000		
RXAGENT1	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
SCSICPIO	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
SFSFCFT	0	0	0	....	0	0	0	0	0	0	XC,---,DORM	1	0	0	0	0	1500	...		
SMTP	.00	.001	.001	1.00	.0	.0	.0	.0	.0	.0	ESA,---,DORM	1	1	.0	.0	.0	100	.000		
TCPIP	.02	.010	.004	2.50	.0	.0	.0	.0	.0	.0	ESA,---,DORM	1	1	.0	.0	.0	3000	.000		
TCPIP50	.00	.001	.001	1.00	.1	.0	.0	.0	.0	.0	ESA,---,DORM	1	1	.0	.0	.0	3000	.007		
TEST50	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
UFTD	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
VMNFS	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
VMSEVRP	0	0	0	....	0	0	0	0	0	0	XC,---,DORM	1	0	0	0	0	1500	...		
VMSEVR	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	1500	...		
VMSEVS	0	0	0	....	0	0	0	0	0	0	XC,---,DORM	1	0	0	0	0	1500	...		
VMSEVU	0	0	0	....	0	0	0	0	0	0	XC,---,DORM	1	0	0	0	0	1500	...		
YVETTE	0	0	0	....	0	0	0	0	0	0	ESA,---,DORM	1	0	0	0	0	100	...		
Command ==>>																				
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return															

Figure 17. Layout of User Resource Usage report (FCX112)

Figure 17 on page 431 shows a report sorted in ascending order of the user IDs, as indicated by the underscore characters above the Userid heading. Other possible sort sequences are indicated by dots above the corresponding columns. These sequences can be set using the SORT subcommand with the column heading (see “SORT” on page 295) or by positioning the cursor on the heading of the field that is to be sorted and then pressing Enter.

The first entry is reserved for the overall system averages, regardless of the current sorting sequence. It is identified by: >System<-userid.

Some additional lines with user class averages will be shown at the top when user classes have been defined (see “FCONTROL UCLASS” on page 89). The *userid* shown will then be the user class name.

For more detailed information about a specific user, do one of the following:

1. Place the cursor on the Userid field of the virtual machine you are interested in and press Enter
2. On the command line, enter the subcommand: USER *userid*

## Field Descriptions:

The following field descriptions apply to the USER and USERLOG *userid* reports, with the exception of the Userid field, which is replaced by the Interval End Time field in the USERLOG report.

### Userid

The user identification. The >>Mean>> line shows the overall average values for all users.

*CPU Load:*

### %CPU

The percentage of total CPU used. This value is based on the usage of a single processor. Values that exceed 100% are possible for virtual MP users.

### TCPU

The total CPU time used during the interval, in seconds.

### VCPU

The virtual CPU time used during the interval (emulation time), in seconds.

### T/V Ratio

The ratio of total to virtual (emulation) CPU time used.

*Virtual IO/s:*

### Total

The total virtual I/O rate per second in the selected interval.

### DASD

The virtual I/O rate to disks.

### Avoid

The rate of I/O operations that were avoided because the requested blocks could be found in the minidisk cache or because the I/O operations were done to virtual disks.

### Diag 98

The Diagnose X'98' rate. Diagnose X'98' allows virtual machines to execute I/O operations very efficiently, bypassing normal CCW translation logic.

### UR

The rate of start requests to virtual UR devices.

### Pg/s

The page rate (total pages read from or written to DASD).

### User Status

Status information about the virtual machine, as collected at the end of the last monitor sample interval.

Possible values for the status indicators (xxx,yyy,zzzz) follow:

For xxx:

#### EME

The virtual machine architecture mode was previously ESA/390, but it is now z/Architecture.

#### ESA

The virtual machine architecture mode is ESA/390.

#### XC

The virtual machine architecture mode is ESA/XC.

#### Z

The virtual machine architecture mode is z/Architecture.

#### ZXC

The virtual machine architecture mode is z/XC.

#### ???

The virtual machine architecture mode is unknown.

For yyy:



**CLn**

Indicates that the virtual machine was in the queue as a class *n* user.

Dashes indicate that the virtual machine is not in the queue.

For zzzz, more than one state could apply at the same time, but only the first state found is inserted (the search order is from top to bottom as follows):

**IDLE**

The virtual machine was found to have been idle longer than the initial period set by the FC FORCEUSR subcommand. The status field is highlighted in red and reverse video.

**LOOP**

The virtual machine was found to have been looping longer than the initial time period set by the FC FORCEUSR subcommand. The status field is highlighted in red and reverse video.

**DORM**

The virtual machine is in DORMANT state.

**DISP**

The virtual machine is in the dispatch list.

**ELIG**

The virtual machine is in the eligible list (in resource wait).

A status of LOGOFF is shown for virtual machines that are no longer active, but these virtual machines' data are still included when displaying average data because their last session fell within the average period processed.

A status of LOGON is shown for virtual machines that are in the process of logging on.

*UsrTime Minutes:*

**Logd**

The time, in minutes, that the user was found to be logged on.

**Act.**

The time, in minutes, that the user was found to have been active. The whole length of a sample interval is counted as active time if the user consumed some virtual CPU time or executed any I/O operations during the interval.

*Spool:*

**Total Pages**

The total number of pages spooled for the user.

**Rate SPg/s**

The spool read and write rate per second.

**MDC Insert MDC/s**

The minidisk cache insertions per second. This value is an approximate indicator. Due to the collection mechanism, the numbers shown can be considerably lower than the users' actual MDC insertion activity, especially for long monitor intervals.

**Share**

The absolute share (in *nn%A* format) or relative share of the user. One of the following suffixes could be appended to the share value:

**m**

If a maximum share has been set for the user.

**h**

If the maximum share has been set with a hard limit.

This field does not display the value of the maximum share; it displays whether a maximum share, a hard limit, or both have been set.

For z/VM 5.4 or later, a virtual machine's share is distributed dynamically across its started virtual processors rather than its defined virtual processors.

For special cases, this field contains one of these values:

**DED**

All defined virtual processors work as dedicated on logical processors. Because only processors of the primary logical processor type can be dedicated, this string can appear only for users with virtual CPUs of type CP or IFL.

**MIX**

The types of defined virtual processors in running state are different, and the summary share value cannot be calculated.

**STOP**

All defined virtual processors are in stopped state.

For related information, see the **Share** field description in [“FCX115, User Resource Details – USER userid” on page 443.](#)

**Vconf Rebal Scat.**

The weighted average grade of how scattered the virtual configuration tended to be placed.

**Nr of Users**

The number of users that contributed to the average shown on this line (values are inserted only for lines with averages; this field is blank for user data lines).

**Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages or totals are shown for all fields, with the exception of the user status and share fields, which always display the virtual machine status found when the last sample was taken. All of a user's sessions during the selected period contribute to the averages shown. Note also that all of the activity fields (rates, for example) are based on each user's connection time within the overall period, *not* the total elapsed time for the entire period.

## FCX113, User Page Data Screen – UPAGE

**Performance Toolkit Subcommand:** “UPAGE” on page 335.

The User Page Data Screen (FCX113) contains detailed information on each user's paging activities during the last measuring interval (for current data), or since the last RESET (when average data has been selected). Users who have been active during the last interval are shown highlighted on the screen.

FCX113		CPU 3090	SER nnnnn	Interval HH:MM:SS - HH:MM:SS		Perf. Monitor													
-----		Data		Paging Activity/s				Number of Pages											
		Owned	Reads	Write	Page	--Page Migration--				WSS	Resrld	--Resident--				--Locked--			
Userid						>2GB>	X>MS	MS>X	X>DS			R<2GB	R>2GB	L<2GB	L>2GB	XSTOR	DASD	Stor	Nr of
>System<		10.6	51.1	34.8	.0	.0	.0	.0	.0	1446k	65609	8776	1372k	0	10	0	840841	6341M	56
User Class Data:																			
PERF		.0	12.8	12.4	.0	.0	.0	.0	.0	4234	0	0	4234	0	0	0	8824	3072M	1
Service		.0	.0	.0	.0	.0	.0	.0	.0	493	512	1	511	0	19	0	910	32768K	1
User Data:																			
APCSCIF		0	.0	.0	.0	.0	.0	.0	.0	238	0	0	1	0	0	0	242	64M	
BK200001		0	.0	.0	.0	.0	.0	.0	.0	849408	0	0	786462	0	0	0	540	3328M	
BK200007		0	.0	.0	.0	.0	.0	.0	.0	849408	0	3052	783410	0	0	0	786982	3328M	
BK200008		0	.0	.0	.0	.0	.0	.0	.0	849408	0	19	786443	0	0	0	786982	3328M	
BK300012		0	1321	353	.0	.0	.0	.0	.0	1074k	0	24	1525k	0	0	0	4195k	16G	
BK400015		0	231	700	.0	.0	.0	.0	.0	26277k	0	194642	22470k	0	0	0	15340k	100G	
CFT2NDA		0	.1	.2	.0	.0	.0	.0	.0	268	0	0	268	0	0	0	16255	64M	
CHG00001		0	.0	.2	.0	.0	.0	.0	.0	181	0	0	183	0	0	0	421	65M	
CHG00006		0	.0	1.1	.0	.0	.0	.0	.0	355	0	0	357	0	0	0	577	65M	
CHG00008		0	.0	.0	.0	.0	.0	.0	.0	49	0	0	49	0	0	0	396	65M	
CHG00010		0	.1	.2	.0	.0	.0	.0	.0	255	0	0	255	0	0	0	420	65M	
CHSCTSTN		0	.0	.0	.0	.0	.0	.0	.0	2	0	0	2	0	0	0	29	16M	
DG100037		0	.0	168	.0	.0	.0	.0	.0	270528	0	4	223832	0	0	0	516343	2048M	
DG100038		0	374	.0	.0	.0	.0	.0	.0	340183	0	29	123137	0	1	0	516343	2048M	
DTCVSW1		0	.0	.0	.0	.0	.0	.0	.0	46	0	0	47	0	1	0	2853	32M	
DTCVSW2		0	.0	.0	.0	.0	.0	.0	.0	46	0	0	47	0	1	0	2853	32M	
D250A		128	.0	.0	.0	.0	.0	.0	.0	190	0	0	2	0	1	0	202	32M	
D2501		464	.0	.3	.0	.0	.0	.0	.0	44	0	0	111	0	0	0	897	32M	
FTPSERVE		0	.0	.0	.0	.0	.0	.0	.0	19	0	0	20	0	1	0	1601	32M	
GCS		0	.0	.0	.0	.0	.0	.0	.0	52	0	0	1	0	1	0	52	16M	
GERTNERJ		0	.1	.1	.0	.0	.0	.0	.0	29	0	0	29	0	0	0	297	32M	

Select a user for user details  
 Command ==> \_  
 F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 18. Layout of User Page Data Screen (FCX113)

Figure 18 on page 435 shows a screen sorted in ascending order of the user IDs, as indicated by the underscore characters above the Userid heading. Other possible sort sequences are indicated by dots above the corresponding columns. These sequences can be set using the SORT subcommand with the column heading (see “SORT” on page 295) or by positioning the cursor on the heading of the field that is to be sorted and then pressing ENTER.

The first entry is reserved for the overall system average for all users, regardless of the current sorting sequence. It is identified by: >System<-userid.

Some additional lines with user class averages will be shown at the top when user classes have been defined (see “FCONTROL UCLASS” on page 89). The 'userid' shown will then be the user class name.

Use the scrolling PF-keys or subcommands to view data not visible on the initial screen, or use the LOCATE subcommand if you want to search for a specific user.

For more detailed information about a specific user, do one of the following:

1. Place the cursor on the Userid field of the virtual machine you are interested in and press Enter
2. On the command line, enter the subcommand: USER *userid*

### Field Descriptions:

The following field descriptions apply to the UPAGE and UPAGELOG *userid* screens, with the exception of the Userid field, which is replaced by the Interval End Time field on the UPAGELOG screen.

#### Userid

The user identification.

**Data Spaces Owned**

The number of non-base address spaces owned by the virtual machine.

*Paging Activity/s:*

*Page Rate:*

**Reads**

The rate at which pages were read in via paging operation

**Write**

The rate at which pages were written out via paging operation

**Page Steals**

The rate at which pages were stolen from the virtual machine's address spaces while they were private.

*Page Migration:*

**>2GB>**

The rate at which page frames were moved below the 2GB line for translation.

**X>MS**

The rate at which page frames were moved from XSTORE to main storage

**MS>X**

The rate at which page frames were migrated out from main storage to XSTORE

**X>DS**

The rate at which page frames were migrated out from XSTORE to DASD

*Number of Pages:*

**WSS**

The number of pages in the user's projected working set

**Resrvd**

The number of reserved pages for the user

**Resident R<2GB**

The number of pages that are resident below the 2GB line.

**Resident R>2GB**

The number of pages that are resident above the 2GB line. This field is available only when running in 64-bit mode.

**Locked L<2GB**

The number of pages locked in the user's private address spaces below the 2GB line.

**Locked L>2GB**

The number of private pages that are locked above the 2GB line.

**XSTOR**

The number of pages of this user which are currently resident in XSTORE (without dedicated XSTORE)

**DASD**

The number of page slots used on system paging areas on DASD

**Stor Size**

The machine's virtual storage size.

Where the user's storage consists of discontinuous extents, this is the number of bytes addressable by the user, *not* the highest address that can be referenced.

**Nr of Users**

The number of users which contributed to the average shown on this line (values are inserted only for lines with averages; the field remains empty for normal user data lines).

**Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages will be shown for all data fields. All of a user's sessions during the selected period contribute to the averages shown. Note also that all of the activity fields (rates) are based on each user's connection time within the overall period, *not* the total elapsed time for the entire period.

## FCX114, User Wait States Screen – USTAT/USTATG

**Performance Toolkit Subcommand:** “USTAT / USTATG” on page 351.

User wait state data can be shown in one of two formats:

1. The User Wait States Screen (FCX114), a standard display with user activity and wait state percentages shown numerically
2. GDDM graphics display where user activity and wait state are shown in bar chart format (requires access to GDDM code and a display terminal with graphics capability).

## Standard User Wait State Display

The User Wait States Screen (FCX114), obtained by issuing the USTAT subcommand, is based on CP monitor data. It shows user status information collected by the high-frequency sampling facility of the CP monitor.

Status information of users which have been active during the last interval is shown highlighted on the screen.

[illegible]

Figure 19. Layout of User Wait States Screen (FCX114)

The sample shows a screen with the users sorted in descending user activity sequence, the initial setting (indicated by the underscore characters above the %ACT heading). Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the SORT subcommand with the appropriate column heading as sort argument (see [“SORT” on page 295](#) and [“USTAT / USTATG” on page 351](#)), or by positioning the cursor on the heading of the field that is to be sorted and then pressing Enter.

The first entry is reserved for the overall system average for all users, regardless of the current sorting sequence. It is identified by the >>Mean>> line, which could indicate system bottlenecks that might be masked by a high CPU usage of a minority of users: paging or I/O bottlenecks will appear as a high percentage of in-queue samples showing page wait or I/O wait (or instruction simulation wait for CMS diagnose I/O).

Some additional lines with user class averages will be shown at the top when user classes have been defined (see “FCONTROL UCLASS” on page 89). The *userid* shown will then be the user class name.

For more detailed information about a specific user, do one of the following:

1. Place the cursor on the `Userid` field of the virtual machine you are interested in and press `Enter`
2. On the command line, enter the subcommand: `USER userid`

**Field Descriptions:**

The following field descriptions apply to the USTAT/USTATG and USTATLOG displays, with the exception of the Userid field, which is replaced by the Interval End Time field on the USTATLOG display.

Users' wait states are tested for in the following sequence:

- I/O wait
- Console function wait
- Instruction simulation wait
- Page wait
- CPU wait
- Running
- SVM wait and in the eligible list
- Loading
- Dormant
- Dormant and in SVM wait
- I/O active
- Test idle wait
- SVM wait and test idle wait
- Page active wait
- Other

Only the first matching wait state found is accounted for.

**Userid**

The user identification.

**%ACT**

The percentage of samples during which the user was found to be active, consuming CPU or waiting for some service.

**%RUN**

The percentage of active samples during which the user was found to be running on a logical processor.

**%CPU**

The percentage of active samples during which the user was found to be waiting to run on the logical processor.

**%LDG**

The percentage of active samples during which the user was found in loading state, that is, with a very high page read rate.

**%PGW**

The percentage of active samples during which the user was found in page wait state.

**%IOW**

The percentage of active samples during which the user was found in I/O wait state.

**%SIM**

The percentage of active samples during which the user was found in instruction simulation wait, that is, waiting for CP to complete simulation of a privileged instruction.

**%TIW**

The percentage of active samples during which the user was found in test idle wait, that is, waiting to be dropped from the dispatch list.

**%CFW**

The percentage of active samples during which the user was found in console function wait, that is, waiting for CP to run one of the console function commands that allow a user to alter the state and configuration of the virtual machine (ATTACH, DETACH, or LINK, for example).

**%TI**

The percentage of active samples during which the user was found in SVM wait and test idle wait. SVM wait state indicates that the user was waiting for the completion of a communication with a service virtual machine.

**%EL**

The percentage of active samples during which the user was found in SVM wait and in the eligible list (resource wait).

**%DM**

The percentage of active samples during which the user was found in SVM wait and in the dormant list. (This status is considered to be active because the machine is waiting for something, and is not idle.)

**%IOA**

The percentage of active samples during which the user was found waiting for something other than the above while an I/O operation started by this machine had not yet completed.

**%OTH**

The percentage of active samples during which the user was found in other states.

**%LIM**

The percentage of active samples during which the user was found in the limit list, that is, in the dispatch list waiting for CPU, but prevented from running by a maximum share setting.

**%PGA**

The percentage of samples showing page-active wait, that is, the virtual machine has loaded a wait PSW while waiting for the completion of a page read operation. This state is to be expected for guest systems running with PAGEX ON.

*%Samples showing:* The columns under this heading show the relative frequency with which a machine was found in the different lists and classes:

**Q0**

The percentage of active samples during which the user was found in Q0 (in the dispatch list as a class 0 user).

**Q1**

The percentage of active samples during which the user was found in Q1 (in the dispatch list as a class 1 user).

**Q2**

The percentage of active samples during which the user was found in Q2 (in the dispatch list as a class 2 user).

**Q3**

The percentage of active samples during which the user was found in Q3 (in the dispatch list as a class 3 user).

**E0-3**

The percentage of active samples during which the user was found in the eligible list (any class).

**vCPU Type**

Indicates the type of virtual processor:

**CP**

Central Processor

**ICF**

Internal Coupling Facility

**IFL**

Integrated Facility for Linux



**ZIIP**  
IBM Z Integrated Information Processor

**MIX**  
The MP user has different PU types.  
This field contains a value during measured interval.

**vCPU Cnt**  
The count of all defined virtual processors during measured interval.

**Nr of Users**  
The number of users that contributed to the average shown on this line (values are inserted only for lines with averages; this field is blank for user data lines).

**Effect of AVERAGE and INTERIM subcommands:** Overall or interim period averages will be shown for all data fields. All of a user's sessions during the selected period contribute to the averages shown.

**Graphical Display of User Status Data**

This type of display, obtained by issuing the USTATG subcommand, often gives an overview of user wait states which is easier to interpret than the large amount of figures shown in the user status display obtained with the USTAT subcommand, but it requires access to GDDM code and also a display terminal with graphics capability.




Figure 20. Layout of Graphic User Wait State Display

Wait state information for all users can be shown (use the scrolling PF-keys if necessary), and users can be sorted according to different criteria by supplying the name of the sorting variable with the USTATG subcommand.

**Field Description:**

Each line of the bar chart represents the wait state information for one virtual machine. The userid of the machine is inserted in front of the corresponding bar.

### User % in-queue (%xxx sort)

The bar chart shown under this heading gives the percentage of high-frequency samples where the user was found to be active (in queue). This is the equivalent of the '%ACT' value in the standard user status display.

'%xxx' indicates the current sorting sequence.

### Status while in queue

The bar chart shown under this heading indicates the relative frequency with which the machine was found in different states while in-queue.

No load bar can be shown if none of the accumulated samples found the user to be in a non-dormant state. The string:

```
No activity
```

will then be inserted instead of a load bar.

### Description

Explains which colors and shading patterns were used for the different states shown under the heading 'Status while in queue'.

Note that percent of samples showing wait time in the limit list (for users with max. share settings) and percent of samples showing 'page-active' wait time (for users with PAGEX ON) will not be displayed separately but will be included in the 'CPU wait' and 'Page wait' percentages respectively.

An additional line at the bottom of the display shows the average of all users' in-queue states; it corresponds to the >System< entry on the standard user status display. This is the main indicator of paging or I/O bottlenecks on the system.

### Saving / Printing Screen:

The graphical part of the screen can be saved on disk for further use. The field:

```
( FCONUST 1 )
```

at the right hand side of the bottom line can be overwritten. It allows you to:

- Define the file name of the disk file: overwrite the string FCONUST with the file name to be used for the disk file
- Set the number of copies to be created when printing the screen contents: overwrite the number 1 with a digit corresponding to the number of copies to be printed.

You can then use the allocated PF-keys to let the data be written to disk:

- Press PF6 (or the PF-key to which you assigned the 'SAVE' function) to save the graphics part of the display as file 'xxxxxxx ADMGDF A', where xxxxxxxx is the file name previously entered in the '( FCONUST 1 )' field in the bottom line of the screen. The resulting file can later be retrieved by GDDM for re-displaying, or it can be converted and included in a script file for processing by DCF (Document Composition Facility, 5748-XX9).
- Press PF9 (or the PF-key to which you assigned the 'PRINT' function) to save the graphics part of the display as file 'xxxxxxx ADMPRINT A', where xxxxxxxx is the file name previously entered in the '( FCONUST 1 )' field in the bottom line of the screen. It can then be printed using standard GDDM facilities. line.

### Effect of AVERAGE and INTERIM subcommands:

Overall or interim period averages will be shown for all data fields. The graphical representation of user wait states will be based on average values when an AVERAGE subcommand has been entered prior to invoking USTATG. You cannot switch between average and current data without first leaving the USTATG display.

## FCX115, User Resource Details – USER *userid*

**Performance Toolkit Subcommand:** “USER” on page 346.

The User Resource Details report (FCX115) is shown when you specify a user ID on the USER subcommand, or when you press Enter after selecting a user with the cursor on the general user, user paging, or user status report:

```

FCX115      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS      Perf. Monitor

Detailed data for user TCPIP50
Total CPU      : .0%      Storage def.    :      128MB  Page fault rate: .0/s
Superv. CPU    : .0%      Resident <2GB  :      1060  Page read rate : .0/s
Emulat. CPU    : .0%      Resident >2GB  :     12064  Page write rate: .0/s
I/O rate       : .1/s     Proj. WSET     :     12837  Pgs moved >2GB>: .0/s
DASD IO rate   : .0/s     Reserved pgs   :         0  Main > XSTORE  : .0/s
UR I/O rate    : .0/s     Locked pages   :        287  XSTORE > main  : .0/s
Diag. X'98'    : .0/s     XSTORE dedic.  :         0MB  XSTORE > DASD  : .0/s
*BLOCKIO      : .0/s     XSTORE pages   :         0  SP00L pg reads : .0/s
Last IPL       : CMS      DASD slots     :         0  SP00L pg writes: .0/s
Config mode    : ESA390   IUCV X-fer/s   :      .0/s   MDC insert rate: .0/s
Base CPU type  : CP       Share              :     3000   MDC I/O avoided: .0/s
Base CPU affin.: ON      Max. share        :      ...

#I/O active    :      0    Active          :      0%    PSW wait      : ..%   I/O act.      : ..%
Stacked blk    : ..      Page wait       : ..%    CF wait       : ..%   Eligible      : ..%
Stat.: ESA,P06,DORM  I/O wait       : ..%    Sim. wait     : ..%   Runnable      : ..%

Proc. %CPU %CP %EM IO/S Status Type Aff SRel. SAbs. Limit MRel. MAbs.
00 BI .0 .0 .0 .1 ESA,P06,DORM CP ON 500 .....
01 .0 .0 .0 .0 ESA,P06,DORM CP ON 500 .....
02 .0 .0 .0 .0 ESA,P06,DORM CP ON 500 .....
03 .0 .0 .0 .0 ESA,P06,DORM CP ON 500 .....
04 .0 .0 .0 .0 ESA,P06,DORM CP ON 500 .....
05 .0 .0 .0 .0 ESA,P06,DORM CP ON 500 .....

Data Space Name Size Mode PgRd/s PgWr/s XRd/s XWr/s Migr/s Steal/s
BASE 128MB Priv .0 .0 .0 .0 .0 .0

Device activity and status:
0009 3215 .0 000C 254R CL *, EOF NOH NCNT
000D 254P CL A, CO 01, NOH NCNT 000E 1403 CL A, CO 01, NOH NCNT
0190 3390 .0 B8A0,RR, 214Cyl,----->0 0191 3390 .0 B877,WR, 2Cyl,----->0
0198 3390 .0 B877,RR, 20Cyl,----->0 019D 3390 .0 B8A0,RR, 292Cyl,----->0
019E 3390 .0 B8A0,RR, 500Cyl,----->0 0500 CTCA .0 NOT READY
0501 CTCA .1 NOT READY 0551 3390 .0 B8A2,RR, 40Cyl,----->0
0591 3390 .0 B8A1,RR, 122Cyl,----->0 0592 3390 .0 B8A1,RR, 140Cyl,----->0
5050 OSA .0 QDIO->SUBNT240 SYSTEM 5051 OSA .0 QDIO->SUBNT240 SYSTEM

Command ===>
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

```

Figure 21. Layout of User Resource Details report (FCX115)

The information shown is grouped into the following sections:

1. The first section, Detailed data for user *userid*, which has the same layout for all users, shows general performance information for a user. Fields that indicate any user activity are highlighted.
2. For virtual MP users only, an additional section displays the CPU and I/O load of each of the virtual processors, and the corresponding status information. For example:

```

Proc. %CPU %CP %EM IO/S Status Type Aff SRel. SAbs. Limit MRel. MAbs.
00 BI .0 .0 .0 .1 ESA,P06,DORM CP ON 500 .....
01 .0 .0 .0 .0 ESA,P06,DORM CP ON 500 .....
02 .0 .0 .0 .0 ESA,P06,DORM CP ON 500 .....
03 .0 .0 .0 .0 ESA,P06,DORM CP ON 500 .....
04 .0 .0 .0 .0 ESA,P06,DORM CP ON 500 .....
05 .0 .0 .0 .0 ESA,P06,DORM CP ON 500 .....

```

- The next section displays detailed information about the paging activity for each of the data spaces the user owns. For example:

Data Space Name	Size	Mode	PgRd/s	PgWr/s	XRd/s	XWr/s	Migr/s	Steal/s
BASE	128MB	Priv	.0	.0	.0	.0	.0	.0

- If you selected an SFS server, an additional section displays summary information about shared file system load and performance. For example:

SFS Load	REQ. Servr	<- Load / Req.->	<- Elapsed Time per -->
USERID	/sec %BUSY	SEC. I/O SAC	BLKIO LKWT CKPT ESMC ..
EDLBACK	21.0 9	.085 2.78 5.0	.051 .011 1.296 ... ..

This is an extract from the Shared File System Server report. The information shown is available only when the necessary prerequisites are met (CP MONITOR data collection is activated); otherwise, it is omitted. See [“SFS” on page 289](#) for more information.

- The final section, `Device activity and status`, lists all virtual devices that are currently defined for the user. Where possible, the activity of individual devices is also calculated and displayed. Devices that have been active during the last measuring interval are highlighted.

The following information might be appended to the header line with the user ID of the virtual machine that is being analyzed:

**sec. user:**

The ID of a virtual machine that is receiving console output from this user using the Single Console Image Facility (SCIF), if the SCIF has been activated.

**alt. user:**

The ID of an alternate virtual machine for which work is being done (set using Diagnose X'D4').

**Field Descriptions:**

*Detailed data for user userid:*

**Total CPU**

The percentage of total CPU used.

All of the CPU load percentages are based on the usage of a single processor. Values that exceed 100% are possible for virtual MP users.

**Note:** The CPU load value for the SYSTEM user ID is taken from the %SYS column of the general CPU report. It is updated only when an automatic update causes recalculation of all performance values, but not when an update for the detailed user report alone is forced by pressing the Enter key.

**Superv. CPU**

The percentage of supervisor CPU used.

**Emulat. CPU**

The percentage of emulation mode CPU used (virtual CPU).

**I/O rate**

The average number of non-spooled I/O operations per second during the last interval.

**DASD IO rate**

The average number of I/O operations per second to disk during the last interval.

**UR I/O rate**

The average I/O rate to virtual unit record devices (spooled).

**Diag. X'98'**

The average rate of Diagnose X'98' instructions per second.

**\*BLOCKIO**

The average rate of block I/O operations run per second, using the \*BLOCKIO service.

**Last IPL**

Indicates the device number (DEV *vdev*) of the last IPLed device for the virtual configuration, or the name (*sysname*) of the last named saved system to be IPLed. (The named saved system was previously saved using the SAVESYS command.) DEV NONE indicates that no device or named saved system was IPLed.

**Config mode**

Displays the virtual configuration mode for the virtual machine: CF, ESA390, LINUX, or VM.

**Base CPU type**

Indicates the base virtual CPU type. Possible values are:

**CP**

Central Processor

**ICF**

Internal Coupling Facility

**IFL**

Integrated Facility for Linux

**ZIIP**

IBM Z Integrated Information Processor

**Base CPU affin.**

Indicates the CPU affinity for the base virtual CPU.

**Storage def.**

The virtual storage size of the virtual machine, as defined in the directory or by the DEFINE STORAGE command.

Where the user's storage consists of discontinuous extents, this is the number of bytes that are addressable by the user, *not* the highest address that can be referenced.

**Resident <2GB**

The number of pages that were resident below 2GB at the end of the last measuring interval.

**Resident >2GB**

The number of pages that were resident in main storage above the 2GB line at the end of the last measuring interval.

**Proj. WSET**

The number of pages in the projected working set at the end of the interval.

**Reserved pgs**

The number of main storage pages that are reserved for this user.

**Locked pages**

The number of main storage pages that are locked for this user.

**XSTORE dedic.**

The amount of expanded storage that is dedicated to this user.

**XSTORE pages**

The number of expanded storage blocks that are occupied by this user's pages.

**DASD slots**

The number of page slots used on system page areas on DASD.

**IUCV X-fer/s**

The number of IUCV data transfers per second (from and to this virtual machine).

**Share**

The user's share, in the format *nnnnn* for relative share, or as *nnnnn%* for an absolute share percentage.

For special cases, this field contains:

**DED**

All running virtual processors are dedicated to logical (host) processors.

**MIX**

All running virtual processors have mixed types. Actual values of share per CPU type are shown in the Activity of Virtual Processors for Virtual MP Users section.

**STOP**

All virtual processors are in stopped state.

Note that this field does not display the value of the maximum share; it might contain only an indicator of maximum share settings (no limit, soft limit, or hard limit). For related information, see the Share field description in [“FCX112, User Resource Usage – USER” on page 431](#).

**Max. share**

The user's maximum share, if available, in the format *nnnnn* for a relative share, or as *nnnnn%* for an absolute maximum share percentage.

The character H is appended if the share setting is "hard".

**Page fault rate**

The number of page faults per second during the last measuring interval.

**Page read rate**

The number of pages read per second during the last measuring interval.

**Page write rate**

The number of pages written per second during the last measuring interval.

**Pgs moved >2GB>**

The number of pages moved below the 2GB line per second for translation during the last measuring interval.

**Main > XSTORE**

The number of pages migrated per second from main storage to expanded storage during the last measuring interval

**XSTORE > main**

The number of pages migrated per second from expanded storage to main storage during the last measuring interval.

**XSTORE > DASD**

The number of pages migrated per second from expanded storage to external paging devices (paging disks).

**SPOOL pg reads**

The number of spool pages read per second during the last measuring interval.

**SPOOL pg writes**

The number of spool pages written per second during the last measuring interval.

**MDC insert rate**

The number of minidisk cache insertions per second during the last measuring interval

**MDC I/O avoided**

The average rate of I/O operations avoided because the requested data were found in storage, in the minidisk cache.

**MDC setting**

This field is inserted only if a special minidisk cache setting is active for the user. The settings tested for and inserted are:

**NOMDCFS**

If the NOMDCFS option has been set for the virtual machine, indicating that its use of the minidisk cache is not to be throttled by "fair share" considerations.

**Set OFF**

If minidisk caching has been disabled for the virtual machine.

**#I/O active**

The number of active high-speed I/O devices at the end of the measuring interval. This number has little meaning for ordinary CMS virtual machines working with CMS files where only one disk I/O

operation can be active at a time (executed via a Diagnose call to CP), but it gives an indication about I/O overlap for other operating systems, such as z/OS or z/VSE.

### **Stacked blk**

The type of block stacked by CP for this user. Block types can be:

- CPEBK - A CP task execution block has been stacked.
- IORBK - An I/O request block or a timer request block (TRQBK) has been stacked.

### **Stat.**

Status information about the virtual machine, as collected at the end of the last monitor sample interval.

Possible values for the status indicators (xxx,yyy,zzzz) follow:

For xxx:

#### **EME**

The virtual machine architecture mode was previously ESA/390, but it is now z/Architecture.

#### **ESA**

The virtual machine architecture mode is ESA/390.

#### **XC**

The virtual machine architecture mode is ESA/XC.

#### **Z**

The virtual machine architecture mode is z/Architecture.

#### **ZXC**

The virtual machine architecture mode is z/XC.

#### **???**

The virtual machine architecture mode is unknown.

For yyy:

#### **DCS**

The user is disconnected.

#### **DEV**

The virtual machine was IPLed from a device.

#### **Pnn**

nn is the number of defined virtual processors. This value is displayed for virtual MP users.

#### **QDS**

The user is designated as a quick dispatcher (QUICKDISP).

#### **SHR**

The virtual machine was IPLed from a shared system.

For zzzz:

#### **CFWT**

The virtual machine is in console function wait.

#### **DORM**

The virtual machine is in DORMANT state.

#### **ELIG**

The virtual machine is in the eligible list.

#### **IDLE**

The virtual machine is idle.

#### **IOWT**

The virtual machine is in I/O wait.

#### **IUCV**

The virtual machine is in IUCV wait.

**PEND**

The virtual machine is in disconnected wait state.

**PGWT**

The virtual machine is in page wait.

**PSWT**

The virtual machine is in enabled PSW wait.

**RNBL**

The virtual machine is runnable.

**SIMW**

The virtual machine is in instruction simulation wait.

**SLP**

The virtual machine is sleeping.

**STOP**

The virtual machine is stopped.

**TIO**

The virtual machine is in a TIO loop.

**TSCH**

The virtual machine is in a TSCH loop.

A merged status is shown for MP users.

**Active**

The percentage of samples where the virtual machine was found in queue, that is, active at the moment when data was collected. This value, and also the next six user status values, are identical to the status information included in [“FCX114, User Wait States Screen – USTAT/USTATG” on page 438](#). These values are calculated based on all samples collected since the user logged on or since the status counters were last reset.

You can find the time of the last RESET by displaying the USTAT report. The time during which user status sample counts have been accumulated is shown in the top line of the USTAT report.

**Page wait**

The percentage of in-queue samples where the virtual machine was found in page wait.

**I/O wait**

The percentage of in-queue samples where the virtual machine was found in I/O wait. A virtual machine is considered to be in I/O wait when it is in one of these states:

- Flagged as being in I/O wait (CCW translation, for example)
- In instruction simulation wait with a high-speed I/O operation outstanding
- In PSW wait and not idle, with a high-speed I/O operation outstanding.

**PSW wait**

The percentage of in-queue samples where the virtual machine was found in enabled PSW wait, but is not idle, and is not waiting for a high-speed I/O operation to end.

**CF wait**

The percentage of in-queue samples where the virtual machine was found in console function wait.

**Sim. wait**

The percentage of in-queue samples where the virtual machine was found in instruction simulation wait, that is, waiting for CP to simulate a privileged instruction with no high-speed I/O operation outstanding.

**I/O act.**

The percentage of in-queue samples where a high-speed I/O operation was still active (VMDIOACT field in VMDBK > 0) and where the virtual machine was *not* in an I/O wait state.



**Eligible**

The percentage of in-queue samples where the virtual machine was found to be in the eligible list, that is, waiting for a constrained resource.

**Runnable**

The percentage of in-queue samples where the virtual machine was found to be runnable, that is, not waiting for pages, I/O, completion of a console function, instruction simulation, or storage, and not idle. This value also includes any other wait states not mentioned in the preceding list (IUCV soft wait, for example).

*Activity of virtual processors for virtual MP users:*

One line with CPU performance data, the I/O activity generated, and with the processor status is shown for each of the virtual processors. None of these data are shown for ordinary UP users, where the load generated by the single processor would equal the totals inserted above.

**Proc.**

The processor ID in hexadecimal format.

A **B** indicates the base processor and an **I** indicates the IPL processor.

**%CPU**

The percentage of total CPU used.

**%CP**

The percentage of supervisor mode CPU used.

**%EM**

The percentage of emulation CPU used.

**IO/S**

The virtual I/O rate for the processor. The I/O rates shown for each processor do not include I/O counts of devices running under I/O assist.

**Status**

The status information for the VMDBK that describes the virtual processor. See the Stat field description for the meaning of the different status indicators.

**Type**

Indicates the type of virtual processor:

**CP**

Central Processor

**ICF**

Internal Coupling Facility

**IFL**

Integrated Facility for Linux

**ZIIP**

IBM Z Integrated Information Processor

**Aff**

Indicates the CPU affinity for the virtual processor:

**OFF**

CPU affinity has not been set on. When OFF is in effect, the virtual processor is dispatched on only logical processors of primary type.

**ON**

CPU affinity has been set on with the SET CPUAFFINITY command and this virtual processor is dispatched on a logical processor of the same type.

**SUP**

CPU affinity is suppressed. CPU affinity has been set on with the SET CPUAFFINITY command, but there is currently no matching logical processor on which to run.

**SRel.**

The virtual machine's relative share summary value for this CPU, if set. Dots are inserted if an absolute share has been set for this CPU type for the virtual machine, or if the processor is in stopped state.

**SAbs.**

The virtual machine's absolute share summary value (as a percentage) for this CPU, if set. Dots are inserted if a relative share has been set for this CPU type for the virtual machine, or if the processor is in stopped state.

**Limit**

The type of limit set for the maximum share: `Soft` or `Hard`. Dots are inserted if no limit has been set.

**MRel.**

The virtual machine's maximum share relative value, if set. Dots are inserted if no maximum relative share has been set.

**MAbs.**

The virtual machine's maximum share absolute value (as a percentage), if set. Dots are inserted if no maximum absolute share has been set.

When the virtual processor is dedicated to a logical processor, the string (`dedicated on logical processor ###`) is shown instead of all share parameters.

*Data space information:*

One line with data space information is shown for each of the data spaces owned by the user.

**Data Space Name**

The name given to the data space.

**Size**

The size of the data space.

**Mode**

The mode in which the data space is used. It can be:

**Priv**

A private data space that can be accessed only by its owner.

**Publ**

A public data space that can be accessed by all users.

**Shrd**

A shared data space that can be accessed by users who have been specifically authorized to work with the data space.

**PgRd/s**

The page read rate for the data space.

**PgWr/s**

The page write rate for the data space.

**XRd/s**

The rate at which pages for the data space have been moved from expanded storage to main storage.

**XWr/s**

The rate at which pages for the data space have been moved from main storage to expanded storage.

**Migr/s**

The rate at which pages have been migrated from expanded storage to auxiliary storage (DASD).

**Steal/s**

The page steal rate for the data space.

*Shared file system load and performance:*

The information is extracted from the Shared File System Server report and is only displayed if data is available; it is the summary information for the selected user. See [“FCX114, User Wait States Screen – USTAT/USTATG”](#) on page 438 for descriptions of the different fields. Be aware that the load information

shown is typically based on different periods: the SFS data are based on the CP MONITOR data collection interval, not on the normal data collection interval of Performance Toolkit.

#### *Device activity and status:*

Two virtual device entries fit on each line. The general format of these entries follows:

```
xxxx tttt nn.n sssssssssssssssssssssss a
```

where:

#### **xxxx**

Is the virtual device number.

#### **tttt**

Is the virtual device type. The string VDSK is inserted for virtual disks in storage.

#### **nn.n**

Is the average I/O rate for the device during the last measuring interval. This field is calculated only for devices where virtual I/O counts are available in the VDEVBLOK of the device (MDISKS, VCTCA, and all dedicated devices, including dialed screens, for example). Where it cannot be calculated (for virtual UR devices, for example), the I/O rate field is left blank.

#### **sss...sss**

Contains device-specific status information. The format of the information shown in this field varies with the device type. The following data can be expected:

#### *Virtual UR equipment:*

The format is always CL x, eeee hhhh cccc, where:

#### **x**

Is the spooling class of the device.

#### **eeee**

Is EOF or NOEOF for UR input devices, depending on whether the device has been set to EOF or NOEOF. For UR output devices, this field defines the current copy count for the device in the format CO nn.

#### **hhhh**

Is HLD or NOH, depending on whether the device has been set to HOLD or NOHOLD.

#### **cccc**

Is CONT or NCNT, depending on whether the device has been set to CONT or NOCONT.

The hhhh and cccc fields are overlaid with the string NOT READY if the device was found to be not ready at the end of the measuring interval.

#### *Virtual CTCA:*

The status field contains the string NOT READY if a virtual CTCA has been defined, but is not yet coupled to its equivalent in an other user's virtual machine, or, if connected, this field has the format:

```
uuuuuu>--->yyyy sssssss
```

where:

#### **uuuuuu**

Is the user ID on the other side of the virtual CTCA.

#### **yyyy**

Is the virtual device number of the CTCA in the other user's virtual machine.

#### **ss..ss**

Is status information for the virtual CTCA. The field might be blank if there was no special status at the end of the measuring interval, or it might contain one of the following:

#### **CNTL WT**

If a CONTROL command was waiting in the adapter.

**HALT (Y)**

If the CTC was halted by the Y-side.

**READ WT**

If a READ command was waiting in the adapter.

**WRT WT**

If a WRITE command was waiting in the adapter.

**xxxx ACT**

Indicates that the specified command indicated was active in the adapter.

**xxxx PND**

Indicates that the specified command was pending for the adapter.

The first status found is inserted.

**Virtual OSA:**

The general format of the status field is:

```
tttt->llllllll uuuuuuuu
```

where:

**tttt**

Is the network device type. Possible values are:

**BRDG**

For an IQD adapter for a bridge port

**HiSo**

For a HiperSockets adapter

**IEDN**

For an IEDN adapter

**INMN**

For an INMN adapter

**QDIO**

For a QDIO adapter

**VSwi**

For a QDIO VSWITCH connection

**????**

Unrecognized format.

**lllllll**

Is the name of the guest LAN to which the virtual OSA adapter is coupled.

**uuuuuuuu**

Is the user ID that owns the guest LAN.

The network device type is always shown, but the remaining information can be displayed only for adapters that are coupled to a guest LAN.

*Virtual screens (GRAF):*

The status field contains the string NOT READY if a virtual GRAF device has been defined, but no DIAL command has connected it yet to a real display terminal directly or via VM/Pass-Through Facility or VM SNA Console Support (VSCS), or, if connected, it contains a string of the format:

```
yyyy,DIALED ,h c ps
```

where:

**yyyy**

Can be the real device number of a dialed display terminal or the string PVM or SNA, depending on whether the screen from the DIAL command was issued is connected locally or via VM/Pass-Through Facility or VSCS. In the case of a DIAL command from an SNA terminal, the LU name of the terminal might be inserted instead of the string SNA.

**DIALED**

Indicates that the real display terminal has been connected via a DIAL command as opposed to a direct ATTACH via DEDICATE directory statement or ATTACH command. If an LU name was inserted for yyyy, the string is shortened to DIAL.

**h**

Can be blank or H, in which case it indicates that the display terminal has extended highlighting capability.

**c**

Can be blank or C, in which case it indicates that the display terminal has extended color capability.

**ps**

Can be blank or PS, in which case it indicates that the the display terminal has programmed symbols capability.

*Minidisks:*

The general format is:

```
yyyy,rr,ssssCYL,nnnnnn
```

where:

**yyyy**

Is the real device number of the disk pack on which the minidisk is located. Dashes are displayed for virtual disks in storage.

**rr**

Is the read/write status of the minidisks. It can be RR for disks with read-only status or WR for disks that are linked in write mode.

This field indicates the read/write status as defined from CP. The actual CMS access might cause the minidisk to be in read-only mode, even though CP would allow write access.

**ssssCYL**

Is the size of the minidisk in cylinders. For minidisks allocated on FBA disks, this is shown in the format ssssssB, to indicate the number of blocks in the minidisk. The format ssssskB is used if the number of blocks exceeds 999999, where kB is equal to 1000 blocks.

**nnnnnn**

Indicates the last virtual seek address on the device (the address is shown in the form - ->nnn with a preceding arrow), or it provides information about the actual minidisk status if a special status is found. The seek address has little meaning for CMS minidisks, where the user usually does not know where specific files have been allocated, but it provides an indication on the dataset being processed if the minidisk is in OS or DOS format.

If the minidisk status is shown, this status information can be one of the following:

**BUSY**

Indicates that the minidisk was busy at the end of the measuring interval

**QUED**

Indicates that an SSCH request has been queued for the minidisk, and it was waiting for an unsolicited DE at the end of the measuring interval.

**RPEND**

Indicates a "resume pending" condition for the minidisk at the end of the measuring interval.

**RS/RL**

Indicates that the minidisk has been defined with virtual reserve/release.

**SUSP**

Indicates a "subchannel suspended" condition at the end of the measuring interval.

**TDISK**

Indicates that the minidisk has been defined as a temporary minidisk

The *ssssCYL* and the *nnnnnn* parts of the minidisk information can be replaced by the string *nnn% MDC eff.* if fewer real I/O operations are made to the minidisk than virtual I/O operations, where:

***nnn***

Is the efficiency expressed as the percentage of the total number of virtual I/O operations shown to the minidisk that did not result in a real I/O operation.

Although not absolutely correct (there is no fixed 1:1 relationship between the number of virtual and real I/O operations even without caching, and especially for non-CMS minidisks), the information can serve as an approximate indicator of minidisk cache efficiency for CMS minidisks.

***Dedicated devices:***

The format is:

```
yyyy, ATTACHED
```

where:

***yyyy***

Is the real device number of the device that has been attached to the user. In the case of dedicated display terminals, the extended highlighting, color, and programmed symbols capability of the terminal is appended as described for dialed virtual screens above.

The character *a* indicates whether I/O assist is available and active for a specific virtual device.

The character *A* is inserted for I/O devices that are eligible for I/O assist, and this character is highlighted if I/O assist is active for the device. An I/O rate of 0 is shown for XA-mode guests.

**Missing interrupt detection:**

The performance monitor checks the status of all virtual devices at the end of each measuring interval. An interrupt is considered to be missing if a device was found to be busy during two consecutive measurements, and if no I/O interrupt was received in between.

If a missing interrupt is detected on a virtual device in the user's machine configuration, the string *INT .MISS* is inserted at the end of the virtual device's status information field, and the message:

```
FCXPEU355M Interrupt(s) missing for device(s) of user uuuuuuuu
```

is shown in the error message area at the bottom of the report.

**Effect of AVERAGE and INTERIM subcommands:**

None. Averages cannot be shown for the detailed user report.

## FCX116, Shared File System Server Screen – SFS

**Performance Toolkit Subcommand:** “SFS” on page 289.

The Shared File System Server Screen (FCX116), which shows summary performance data for all Shared File System server machines, will be shown if you enter the SFS subcommand with no *userid*. SFS servers which have been active during the last measuring interval will be highlighted on the screen.

FCX116	CPU	nnnn	SER	nnnnn	Interval	HH:MM:SS	-	HH:MM:SS	Perf.	Monitor	
-----	REQ.	Servr	<-	Load /	Req.->	<-	Elapsed Time	per	-->	Actv	Nr.
USERID	/sec	%BUSY	SEC.	I/O	SAC	BLKIO	LKWT	CKPT	ESMC	AGNT	RBK
EDLBACK	21.0	8	.085	2.78	5.0	.051	.011	1.296	...	1.7	0
EDLSERV	.0	...	...	...	...	...	...	...	...	.0	0
EDLSRV	.0	...	...	...	...	...	...	...	...	.0	0
EPLREF	.0	...	...	...	...	...	...	...	...	.0	0
FLDTEST	.0	...	...	...	...	...	...	...	...	.0	0
SERVERW	.0	...	...	...	...	...	...	...	...	.0	0
SERVERX	.0	...	...	...	...	...	...	...	...	.0	0
SERVERY	.0	...	...	...	...	...	...	...	...	.0	0
SERVER3	2.2	1	.014	1.18	4.3	.011	...	...	...	.0	0
SERVER7	.1	0	.005	.05	6.7	.019	...	...	...	.0	0
SERVER8	9.8	5	.025	1.66	5.8	.015	.042	...	...	.2	0
SERV7	.1	0	.018	1.09	.0	...	...	...	...	.0	0
SQLMACHZ	.0	...	...	...	...	...	...	...	...	.0	0
TEMPSEV	...	...	...	...	...	...	...	...	...	...	...
TSTMACH1	.0	...	...	...	...	...	...	...	...	.0	0
VMSYS	.0	...	...	...	...	...	...	...	...	.0	0

Select an SFS server for server details  
Command ==> -  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 22. Layout of Shared File System Server Screen (FCX116)

Performance analysis of SFS servers is based on CP monitor USER and APPLDATA domain SAMPLE data.

The sample shows a screen with a number of SFS server machines sorted in ascending userid sequence, as indicated by the underscore characters above the USERID heading. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the 'SORT' subcommand with the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the heading of the field that is to be sorted and then pressing ENTER.

For more detailed information on one of the shared file system servers either

- Place the cursor on the USERID field of the server you are interested in and press ENTER, or
- Enter the subcommand 'SFS *userid*' on the command line

### Field Description:

#### USERID

The userid of the SFS server machine

#### REQ./sec

Number of file pool requests handled per second

#### Servr %BUSY

Percent of time during which the server machine was active handling requests

#### <- Load / Req.->

The columns under this heading give some information on the average load caused by one file pool request during the selected interval.

#### SEC.

Total service time per average file pool request, in seconds

**I/O**

Total number of I/Os per average file pool request

**SAC**

Number of Storage Access Component calls per average file pool request

*<- Elapsed Time per -->:*

The columns under this heading give information on the elapsed time required for some key activities. They can give pointers to possible bottlenecks if the performance of the SFS server is not adequate.

**BLKIO**

The average service time per block I/O

**LKWT**

The average time spent in lock wait

**CKPT**

The average time spent in checkpoint wait

**ESMC**

The average wait time per External Security Manager call

**Actv AGNT**

The average number of active agents in the interval

**Nr. RBK**

The total number of roll-backs due to deadlocks

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all load fields. All of a server's sessions during the selected period contribute to the averages shown. Note also that all of the activity fields (rates) are based on each server's connect time within the overall period, *not* the total elapsed time for the whole period.



## FCX117, Shared File System Server Details Screen – SFS

**Performance Toolkit Subcommand:** “SFS” on page 289.

The Shared File System Server Details Screen (FCX117), which shows detailed performance data for Shared File System server machines, will be shown if you specify a *userid* on the SFS subcommand, or via cursor-selection of the server on the general Shared File System screen:

```
FCX117      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor

Detailed Data for SFS Server EDLBACK
File Pool Req.: 6163      Blocks per BIO : 7.64      Avg. block I/O: .059s
FP Request rate: 20.6/s   Blocks per I/O : 4.06      Avg. lock wait: .005s
Agents held : 2.3        SAC calls / FPR: 5.9      Avg. check pt.: 2.125s
Active agents : 1.7      Deadlocks w RBK: 0         Avg. ESM call : .....s

Server States:
Total busy : 8.3%      Page wait : .0%      QSAM I/O wait : .0%
CPU (running) : 7.1%   Checkpoint wait: 1.2%

Time per File Pool Request:
Total time : .085s     Lock wait time : .000s     ESM call time : .000s
CPU time : .003s      Block I/O time : .083s     Other : .000s

I/Os per File Pool Request:
Total I/Os : 2.63      Catalog reads : .01      Control Writes : .00
File reads : 2.14      Catalog writes : .02     LOG reads : .00
File writes : .23      Control reads : .00      LOG writes : .20

File Pool Request Percentages:
Open read : 10.4%      Get dir. entry : .0%      Grant/rev auth.: .0%
Open update : 5.0%     Rename : .0%             Lock & unlock : 2.0%
Read : 51.0%          Create alias : .0%       Query : .0%
Write : 8.5%          Cr/del direct. : .0%     Refresh direct.: .0%
Close : 15.4%         O/G/C directory: 4.7%     Other requests : 1.8%
Delete : .7%

Command ==> _
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return
```

Figure 23. Layout of Shared File System Server Details Screen (FCX117)

Load fields for which some activity was found during the selected interval will be highlighted.

Performance analysis of SFS servers is based on CP monitor USER and APPLDATA domain SAMPLE data.

### Field Description:

#### Detailed Data for SFS Server *userid*

The first block of data provides some general information on the activity and utilization of the shared file system server machine *userid*. More detailed information for specific areas is available in the other blocks.

#### File Pool Req.

The number of file pool requests handled by this server during the selected interval

#### FP Request rate

The number of file pool requests handled per second

#### Agents held

The average number of server agents in use

#### Active agents

The average number of server agents which were active, processing SFS requests

#### Blocks per BIO

The average number of blocks in a block I/O request issued by the server

#### Blocks per I/O

The average number of blocks per CP I/O operation

**SAC calls / FPR**

The number of calls to the Storage Access Component per file pool request

**Deadlocks w RBK**

The number of roll-backs due to deadlock situations

**Avg. block I/O**

The average service time per block I/O request

**Avg. lock wait**

The average time spent in lock wait

**Avg. check pt.**

The average time required per checkpoint

**Avg. ESM call**

The average wait time per call of an External Security Manager exit

**Server States:**

The fields under this heading show the total server utilization, and the contribution of different components to this utilization.

**Total busy**

Total server utilization, i.e. the percentage of elapsed time during which the server was either active or waiting for some service

**CPU (running)**

Percentage of elapsed time during which the server was running on a processor

**Page wait**

Percentage of time during which the server was found in page wait state (waiting for a page to be read by CP)

**Checkpoint wait**

Percentage of time needed for handling checkpoints

**QSAM I/O wait**

Percentage of time during which the server was waiting for QSAM I/O to complete

**Time per File Pool Request:**

The fields under this heading show the average time it took to service a file pool request, and the contribution of different components to this service time.

**Total time**

Total service time, in seconds, for an average file pool request during the selected interval

**CPU time**

The average CPU time per file pool request

**Lock wait time**

The average lock wait time per file pool request

**Block I/O time**

The average time needed for block I/O execution per file pool request

**ESM call time**

The average time spent in an External Security Manager exit per file pool request

**Other**

The difference between the total time per file pool request, and the sum of the CPU, lock wait, block I/O and ESM call times

**I/Os per File Pool Request:**

The fields under this heading show the total number of I/Os required per file pool request, and how different activities contributed to this total.

**Total I/Os**

Total number of I/Os per file pool request during the selected interval

**File reads**

The number of I/Os needed for 'file reads' per file pool request

**File writes**

The number of I/Os needed for 'file writes' per file pool request

**Catalog reads**

The number of I/Os needed for catalog reads per file pool request

**Catalog writes**

The number of I/Os needed for catalog writes per file pool request

**Control reads**

The number of I/Os needed for control reads per file pool request

**Control writes**

The number of I/Os needed for control writes per file pool request

**Log reads**

The number of I/Os needed for log reads per file pool request

**Log writes**

The number of I/Os needed for log writes per file pool request

**File Pool Request Percentages:**

The fields under this heading show the percentages by which different request types contributed to the total number of file pool requests.

**Open read**

Percentage of open file read requests

**Open update**

Percentage of open update requests

**Read**

Percentage of read file requests

**Write**

Percentage of write file requests

**Close**

Percentage of close file requests

**Delete**

Percentage of delete file requests

**Get dir. entry**

Percentage of get directory entry requests

**Rename**

Percentage of rename requests

**Create alias**

Percentage of create alias requests

**Cr/del direct.**

Percentage of create and delete directory requests

**O/G/C directory**

Percentage of open, get, and close directory requests

**Grant/rev auth.**

Percentage of grant and revoke authorization requests

**Lock & unlock**

Percentage of lock and unlock requests

**Query**

Percentage of query requests

**Refresh direct.**

Percentage of refresh directory requests

**Other requests**

Percentage of other requests for which no separate numbers are supplied above

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all load fields. All of a server's sessions during the selected period contribute to the averages shown. Note also that all of the activity fields (rates) are based on the server's connect time within the overall period, *not* the total elapsed time for the whole period.

## FCX119, SPOOL File Summary Screen – SPOOL or SPOOL NOSYSTEM

**Performance Toolkit Subcommand:** “SPOOL” on page 296.

The SPOOL File Summary Screen (FCX119) will be shown if you enter the SPOOL subcommand with no parameters at all, or with the NOSYSTEM parameter, in which case system files will be excluded from the display. It is available only in virtual machines with privilege class D (to run Diagnose code X'D8' and the CP 'PURGE' command).

FCX119	CPU nnnn	SER nnnnn	Interval	HH:MM:SS	- HH:MM:SS	Perf. Monitor
<b>All Files</b>						
<b>SPOOL File Overview</b>						
Number of files		Reader	Printer	Punch	System	Total
		9028	1607	2	173	10810
Number of blocks		204883	3486	301	34251	242921
Number of records		8512610	128497	12558	33983	8687648
<b>S</b>	<b>Owner-ID</b>	<b>File</b>	<b>Class</b>	<b>Que</b>	<b>Ac</b>	<b>Blocks</b>
.	PKUN	0028	A	PUN	RDR	6651
.	DMI	0450	A	PUN	RDR	4552
.	DMI	0448	A	PUN	RDR	3317
.	MAINT21	0012	A	PRT	RDR	2338
.	CHRS	8072	T	CON	PRT	OU
.	BMO	8407	A	PUN	RDR	1944
.	WIL	0006	A	PUN	RDR	1592
.	*NSS	1548	A	SYS	-	AC
.	DDIRLIB	5579	X	PUN	RDR	901
.	WHU	0037	B	PUN	RDR	892
.	TOP	0391	A	PUN	RDR	855
.	*NSS	1607	A	SYS	-	836
Select 'P' to purge files, or enter 'SPool User' for user summary						
Command ==>						
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return						

Figure 24. Layout of SPOOL File Summary Screen (FCX119)

The screen shown in the above figure is an example for a summary SPOOL file screen as you may get on a large system with many SPOOL files. The SPOOL file displays will be useful especially when the SPOOL space utilization on your system reaches a dangerous level (as indicated by the percentages shown on the DEVICE SYSOWN display or the utilization percentage shown on the REDISP screen) and when you have to decide which files to purge in order to keep the system running properly.

Two blocks of data are shown:

- A summary of the total SPOOL space utilization.
- A detailed list with the up to the 100 largest SPOOL files found, ordered in descending file size sequence (based on the SPOOL blocks used, not the number of records which is not relevant in this context). Active files (that is, files in use) will be highlighted.

To purge a file, you can enter a 'P' selection character in front of the file to be deleted. Enter the command 'SPool User' to select the SPOOL File User Summary Screen (FCX120), which will show you the SPOOL space utilization per user (see “FCX120, SPOOL File User Summary Screen – SPOOL USER” on page 464).

### Field Description for SPOOL File Summary Data:

#### SPOOL File Overview

This first part of the screen shows the following:

- The total number of SPOOL files
- The total number of SPOOL blocks in use
- The total number of records

in the columns:

**Reader**

For files in a reader queue

**Printer**

For files in a print queue

**Punch**

For files in a punch queue

**System**

For system files

**Total**

For the overall totals of the system

*Detailed SPOOL File Display***S**

This is the column where you can enter the selection character 'P' if you want to purge a file.

**Caution:** Make sure a file is no longer needed before purging it or save it first to tape (using the CP command SPXTAPE LOAD).

**Owner-ID**

The user identification of the file owner. The string \*NSS will be shown for all kinds of saved segments.

**File**

The SPOOL file identification number

**Class**

The SPOOL file class, shown as a string

*c ttt*

where:

**c**

Is the actual SPOOL file class

**ttt**

Is the SPOOL file type. It can be

**CON**

For console log files

**PUN**

For punch files

**PRT**

For print files

**SYS**

For system files (shared segments)

**Que**

The current queue where the file is found. It can be:

**PUN**

For files on the punch queue

**PRT**

For files on the print queue

**RDR**

For files queued to a virtual machine's reader

-

For system files (shared segments)

**Ac**

The current status of the file, when active. It can be:

**OU**

For files being written (open for output)

**IN**

For files being read

**AC**

For system files (shared segments) which are in use

**Blocks**

The number of SPOOL file blocks occupied by the file

**Records**

The number of records in the file

**Date**

The date, in format *ccyy/mm/dd*, when the file was first opened (created)

**Time**

The time, in format *hh/mm*, when the file was first opened (created)

**Name**

The file name given to the file

**Type**

The file type given to the file

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None. The *current* SPOOL utilization is always shown.

## FCX120, SPOOL File User Summary Screen – SPOOL USER

**Performance Toolkit Subcommand:** “SPOOL” on page 296.

The SPOOL File User Summary Screen (FCX120) will be shown if you enter the SPOOL subcommand with the USER parameter. It is available only in virtual machines with privilege class D (to run Diagnose code X'D8' and the CP PURGE command).

FCX120	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	PERF. MONITOR
<b>All Files</b>				
<b>SPOOL File Overview</b>		<b>Reader</b>	<b>Printer</b>	<b>Punch</b>
Number of files		8942	1606	2
Number of blocks		205850	4874	17
Number of records		8546103	206123	679
<b>System</b>				
Total		173	34251	33983
10723		244992	8786888	
<b>Owner-ID Files Blocks</b>				
<b>&gt;System&lt;</b>				
<b>BEGL</b>				
<b>USWA</b>				
<b>WRO</b>				
<b>PSPTOOLS</b>				
<b>DDIRLIB</b>				
<b>NBRU</b>				
<b>SALI</b>				
<b>TBEL</b>				
<b>HHUB</b>				
<b>ECFDMGRO</b>				
<b>CHBITFFM</b>				
<b>Owner-ID Files Blocks</b>				
<b>MCOG</b>				
<b>BMO</b>				
<b>PKUN</b>				
<b>MBN</b>				
<b>MAINT21</b>				
<b>TOP</b>				
<b>WIL</b>				
<b>BRUG</b>				
<b>SRUO</b>				
<b>ZK</b>				
<b>FAM</b>				
<b>HSH</b>				
<b>Owner-ID Files Blocks</b>				
<b>WID</b>				
<b>DMI</b>				
<b>BKAN</b>				
<b>BARSMGR</b>				
<b>RNA</b>				
<b>DPE</b>				
<b>ATR</b>				
<b>JA</b>				
<b>MAKO</b>				
<b>ECFDMGR</b>				
<b>UHAL</b>				
<b>CHRSCS</b>				
Select a user for user details				
Command ===>				
F1=Help F2=Redis F3=Quit F7=Back F8=Forward F9=Top F12=Bottom				

Figure 25. Layout of SPOOL File User Summary Screen (FCX120)

The screen shown in the above figure is an example for a SPOOL file user summary screen you may get on a large system with many SPOOL files. The SPOOL file displays will be useful especially when the SPOOL space utilization on your system reaches a dangerous level (as indicated by the percentages shown on the DEVICE SYSOWN display or the utilization percentage shown on the REDISP screen) and when you have to decide which files to purge in order to keep the system running properly.

Two blocks of data are shown:

- A summary of the total SPOOL space utilization.
- A list with summaries of user SPOOL file utilization, ordered in descending sequence of SPOOL blocks used (the number of records is not relevant in this context). Up to 300 of the largest SPOOL space users will be shown, and users with active files (that is, files in use) will be highlighted.

To get more detailed information on the SPOOL file utilization of one specific user, just move the cursor to the selected userid and press ENTER, or enter the command 'SPOOL *userid*' (see “FCX121, SPOOL File User Details Screen – SPOOL” on page 466 for the screen layout).

Enter the command SPool to select the general SPOOL file summary, with the largest 100 SPOOL files on the system (see “FCX119, SPOOL File Summary Screen – SPOOL or SPOOL NOSYSTEM” on page 461).

### Field Description for SPOOL File Summary Data:

#### SPOOL File Overview

This first part of the screen shows the following:

- The total number of SPOOL files
- The total number of SPOOL blocks in use
- The total number of records

in the columns:



**Reader**

For files in a reader queue

**Printer**

For files in a print queue

**Punch**

For files in a punch queue

**System**

For system files

**Total**

For the overall totals of the system

This part is identical to the first part of the general SPOOL file summary (see [“FCX119, SPOOL File Summary Screen – SPOOL or SPOOL NOSYSTEM”](#) on page 461).

*User Summary*

Three similar columns are shown with user summary information:

**Owner-ID**

The user identification of the file owner. The string >System< is shown as owner for all kinds of saved segments.

**Files**

The number of SPOOL files owned by the user

**Blocks**

The total number of SPOOL file blocks occupied by a user's files

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None. The *current* SPOOL utilization will always be shown.

## FCX121, SPOOL File User Details Screen – SPOOL

**Performance Toolkit Subcommand:** “SPOOL” on page 296.

The SPOOL File User Details Screen (FCX121) will be shown if you specify a *userid* on the SPOOL subcommand. It is available only in virtual machines with privilege class D (to run Diagnose code X'D8' and the CP PURGE command).

FCX121	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS				PERF. MONITOR				
Owner BARSMGR											
SPOOL File Overview				Reader	Printer	Punch	System	Total			
Number of files				149	1	0	0	150			
Number of blocks				5513	1	0	0	5514			
Number of records				252590	39	0	0	252629			
S	OriginID	File	Class	Que	Ac	Blocks	Records	Date	Time	Name	Type
.	BARSSUBA	7949	T CON	RDR		1420	67675	2003/10/20	18:31		
.	BARSSUB7	7947	T CON	RDR		72	3203	2003/10/20	18:32		
.	BARSSUB2	7955	T CON	RDR		70	3176	2003/10/20	18:31		
.	BARSSUB1	7954	T CON	RDR		69	3088	2003/10/20	18:31		
.	BARSSUB5	7946	T CON	RDR		69	3103	2003/10/20	18:31		
.	BARSSUB6	7956	T CON	RDR		67	3044	2003/10/20	18:31		
.	BARSSUB9	7953	T CON	RDR		66	2956	2003/10/20	18:32		
.	...	...	.	...	...	..	....	.....	....		
.	...	...	.	...	...	..	....	.....	....		
.	BARSSUB3	7932	T CON	RDR		1	15	2003/10/20	06:36		
.	BARSSUBC	8015	A PUN	RDR		1	13	2003/10/21	06:39	QUERY	ALL
.	BARSMGR	7995	T CON	PRT OU		1	39	2003/10/21	06:32		
Select 'P' to purge files											
Command ==>											
F1:HELP F2:REDISP F3:QUIT F7:BACK F8:FORM F9:TOP F12:BOT											

Figure 26. Layout of SPOOL File User Details Screen (FCX121)

The screen shown is an example for the 'SPOOL File User Details' display, showing all SPOOL files owned by the selected user.

Two blocks of data are shown:

- A summary of the total SPOOL space utilization by this user
- A detailed list with all SPOOL files owned by this user, in descending file size sequence (based on the SPOOL blocks used, not the number of records which is not relevant in this context). Active files (that is, files in use) will be highlighted.

To purge a file, enter a 'P' selection character in front of the file to be deleted.

### Field Description for SPOOL File Summary Data:

#### SPOOL File Overview

This first part of the screen shows the following:

- The total number of SPOOL files
- The total number of SPOOL blocks in use
- The total number of records

for the selected user in the columns:

#### Reader

For files in the reader queue

#### Printer

For files in the print queue

#### Punch

For files in the punch queue

**System**

For system files

**Total**

For the overall totals of the user

*Detailed SPOOL File Display***S**

This is the column where you can enter the selection character 'P' if you want to purge a file.

**Caution:** Make sure a file is no longer needed before purging it or save it first to tape (using the CP command SPXTAPE LOAD).

**OriginID**

The userid of the SPOOL file originator

**File**

The SPOOL file identification number

**Class**

The SPOOL file class, shown as a string

*c ttt*

where:

**c**

Is the actual SPOOL file class

**ttt**

Is the SPOOL file type. It can be:

**CON**

For console log files

**PUN**

For punch files

**PRT**

For print files

**SYS**

For system files (shared segments)

**Que**

The current queue where the file is found. It can be:

**PUN**

For files on the punch queue

**PRT**

For files on the print queue

**RDR**

For files queued to a virtual machine's reader

-

For system files (shared segments)

**Ac**

The current status of the file, when active. It can be:

**OU**

For files being written (open for output)

**IN**

For files being read

**AC**

For system files (shared segments) which are in use

**Blocks**

The number of SPOOL file blocks occupied by the file

**Records**

The number of records in the file

**Date**

The date, in format *yyyy/mm/dd*, when the file was first opened (created)

**Time**

The time, in format *hh/mm*, when the file was first opened (created)

**Name**

The file name given to the file

**Type**

The file type given to the file

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None. The *current* SPOOL utilization is always shown.

## FCX122, Remote CP Command Execution Screen – CP

**Performance Toolkit Subcommand:** “CP” on page 103.

You will see this screen after entering an explicit CP command (with the CP prefix) while in remote monitoring mode, provided that you have been given 'command' authorization on the remote system (CMD specification in the requestor's entry of file FCONRMT AUTHORIZ on the remote system).

FCX122      CPU 3090   SER 11619   Command executed : 15:04:55   CHVM4

```
cp ind user perfsvm
USERID=PERFSVM    MACH=XA    STOR=0010M VIRT=V XSTORE=NONE
IPLSYS=CMS        DEVNUM=00009
PAGES: RES=000480 WS=000605 LOCK=000000 RESVD=000000
NPREF=000437 PREF=000000 READS=002389 WRITES=002436
XSTORE=000139 READS=650442 WRITES=652672 MIGRATES=002089
CPU 00: CTIME=13:55 VTIME=003:28 TTIME=004:50 IO=080046
RDR=000000 PRT=000000 PCH=000000
R;
```

Command ==> \_  
F1=Help   F4=Top   F5=Bot   F7=Bkwd   F8=Fwd   F12=Return

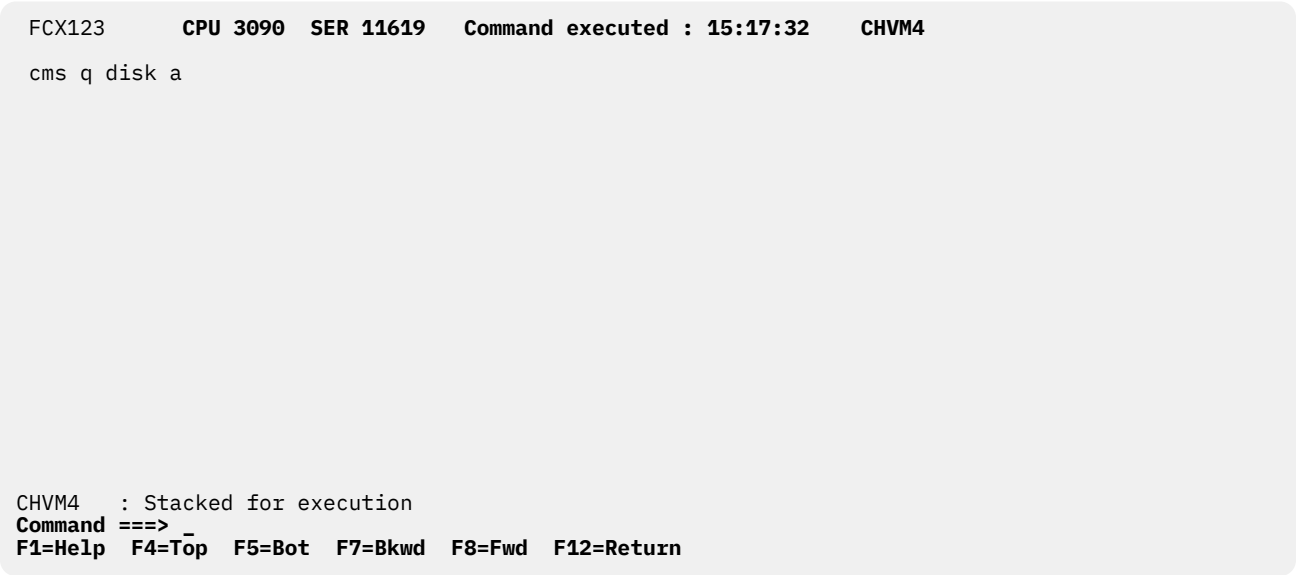
*Figure 27. Layout of Remote CP Command Execution Screen (FCX122)*

The information shown on this screen is the output of the CP command entered. Only up to 8kB of the output will be processed; any remainder will simply be discarded.

# FCX123, Remote CMS Command Execution Screen – CMS


**Performance Toolkit Subcommand:** “CMS” on page 102.

You will see the Remote CMS Command Execution Screen (FCX123) after entering an explicit CMS command (with the CMS prefix) while in remote monitoring mode with an APPC/VM session, provided that you have been given 'command' authorization on the remote system (CMD specification in the requestor's entry of file FCONRMT AUTHORIZ on the remote system).



*Figure 28. Layout of Remote CMS Command Execution Screen (FCX123)*

You will see only the confirmation message:



after entering a CMS command in a remote monitoring session via RSCS links.

## Description:

This screen only confirms that the CMS command entered has been stacked for execution. Any output of the CMS command will be written to the basic mode display of the remote system, but is not retrieved and included on this screen.

# FCX124, Performance Data Selection Menu Screen – MONITOR, MENU

**Performance Toolkit Subcommand:** “MONITOR” on page 110, “MENU” on page 235.

The Performance Data Selection Menu Screen (FCX124):

FCX124

Performance Screen Selection (FL720)

Perf. Monitor

General System Data

1. CPU data menu\*

2. Storage utilization

3. SSI data menu\*

4. Priv. operations

5. System counters

6. CP IUCV services

7. SPOOL file display\*

8. LPAR data menu\*

9. Shared segments

A. Shared data spaces

B. Virt. disks in stor.

C. Transact. statistics

D. Monitor data

E. Monitor settings

F. System settings

G. System configuration

H. VM Resource Manager

I. Exceptions

K. User defined data\*

I/O Data

11. Channel load

12. Control units

13. I/O device menu\*

14. PCI Function menu\*

15. Cache extend. func.\*

16. Reserved

17. DASD seek distance\*

18. I/O prior. queueing\*

19. I/O configuration

1A. I/O config. changes

User Data

21. User resource usage\*

22. User paging menu\*

23. User wait states\*

24. User response time\*

25. Resources/transact.\*

26. User communication\*

27. Multitasking users\*

28. User configuration\*

29. Linux systems\*

2A. CPU Pools menu\*

2B. User Processor Time

History Data (by Time)

31. Graphics selection

32. History data files\*

33. Benchmark displays\*

34. Correlation coeff.

35. System summary\*

36. Auxiliary storage

37. CP communications\*

38. DASD load

39. Minidisk cache\*

3A. Storage mgmt. data\*

3B. Proc. load & config\*

3C. LPAR logs menu\*

3D. Response time (all)\*

3E. RSK data menu\*

3F. Scheduler queues

3G. Scheduler data

3H. SFS/BFS logs menu\*

3I. System log

3K. TCP/IP data menu\*

3L. User communication

3M. User wait states

Pointers to related or more detailed performance data  
can be found on displays marked with an asterisk (\*).

Command ==> \_

F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 29. Layout of Performance Data Selection Menu Screen (FCX124)

See the "Operation in Performance Monitor Mode" chapter in the [z/VM: Performance Toolkit Guide](#) for more information.

## FCX126, LPAR Load Screen – LPAR

**Performance Toolkit Subcommand:** “LPAR” on page 220.

The LPAR Load Screen (FCX126) provides a list of the logical partitions defined for the system and the load for each of the defined logical cores.

```
FCX126      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor

LPAR Data, Collected in Partition FCFT

Processor type and model : 2827-HA1
Nr. of configured partitions: 21
Nr. of physical cores : 101
Dispatch interval (msec) : dynamic

Partition Nr.  Upid #Core Weight Wait-C  Cap %Load CID %Busy %Ovhd  --- --- --- Type  AbsCap MT GrpCapNm GrpCap
EVIC          18    41     2    10     NO   NO   ...  00    .0    .0 --- --- --- CP    ... ..
          10     NO   ...  01    .0    .0 --- --- --- CP    ... ..
FCFT          19    38    18    10     NO   NO   1.1  00    8.0  .1 --- --- --- CP    ... ..
          10     NO   ...  01    7.5  .1 --- --- --- CP    ... ..
          10     NO   ...  02    2.7  .1 --- --- --- CP    ... ..
          10     NO   ...  03    2.8  .1 --- --- --- CP    ... ..
          10     NO   ...  04    1.7  .0 --- --- --- CP    ... ..
          10     NO   ...  05    .3    .0 --- --- --- CP    ... ..
          10     NO   ...  06    1.0  .0 --- --- --- CP    ... ..
          10     NO   ...  07    .0    .0 --- --- --- CP    ... ..
          10     NO   ...  08    .0    .0 --- --- --- CP    ... ..
          10     NO   ...  09   39.1  .0 --- --- --- CP    ... ..
          10     NO   ...  0A   37.2  .1 --- --- --- CP    ... ..
          10     NO   ...  0B   12.4  .2 --- --- --- CP    ... ..
          10     NO   ...  0C    .0    .0 --- --- --- ZIIP  ... ..
          10     NO   ...  0D    .0    .0 --- --- --- ZIIP  ... ..
          10     NO   ...  0E    .1    .1 --- --- --- IFL   ... ..
          10     NO   ...  0F    .1    .0 --- --- --- IFL   ... ..
          10     NO   ...  10    .9    .1 --- --- --- IFL   ... ..
          10     NO   ...  11    .0    .0 --- --- --- IFL   ... ..
PHOS          20    39     5    10     NO   NO   .0  00    .0    .0 --- --- --- CP    ... ..
          10     NO   ...  01    .1    .1 --- --- --- CP    ... ..
          10     NO   ...  02    .1    .1 --- --- --- CP    ... ..
          10     NO   ...  03    .1    .1 --- --- --- CP    ... ..
          10     NO   ...  04    .1    .1 --- --- --- CP    ... ..
VMSYSONE     21    ...     0    10     NO   NO   ...  ...  ...  ... --- --- --- ..    ... ..

General LPAR mgmt overhead : ...
Overall physical load      : 4.7

Command ==>
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return
```

Figure 30. Layout of LPAR Load Screen (FCX126)

The information shown is based on CP monitor SYSTEM domain SAMPLE data. The name of active LPARs is highlighted.

### Field Descriptions:

LPAR Data, Collected in Partition *xxxx* - The partition name of the partition which collected the LPAR data.

Processor type and model: *nnnn-ccc* - The processor type *nnnn* and model designation *ccc*

Nr. of configured partitions: *n* - The number of logical partitions configured in this processor complex.

Nr. of physical cores: *n* - The number of cores in the processor complex.

Dispatch interval (msec): *xxxxxx* - The dispatch interval, as set when the logical partitions were defined, in milliseconds. The string *dynamic* is shown when the interval can be dynamically adapted by the PR/SM microcode.

The next two lines provide information about overall system load:



- **General LPAR mgmt. overhead:** - The general LPAR management time, expressed as a percentage of the available CPU time (elapsed time times the number of physical cores); it does not include LPAR management overhead for specific logical cores which is already shown in column %Ovhd.
- **Overall physical load:** - The sum of the %Load values for all active logical partitions plus the General LPAR mgmt. overhead (if available), that is, it gives you the overall load for normal cores in the processor complex and doesn't include the load of partitions with ICF or IFL processors.

**Note:** Correct overall physical load figures for normal cores will be calculated on systems with enhanced LPAR monitor support. Otherwise the numbers will be misleading on systems with ICF cores, IFL cores, or both, because ICF and IFL processors are not counted as part of the 'Number of physical cores' but they are treated like normal cores when providing logical core data. Because the overall physical load is calculated based on the assumption that there is no data from ICF or IFL cores, the resulting value will be too high.

### **Partition**

The name given to the logical partition. It is highlighted if the partition is active, that is, if at least one logical processor is assigned to the partition.

### **Nr.**

The number of the logical partition

### **Upid**

The user partition ID number for the logical partition.

### **#Cores**

The number of cores defined for the partition.

### **Weight**

The weight defined for this CPU type in the partition. Weight values range from 1 to 999, with a default of 10. The string DED is inserted when the partition is dedicated.

### **Wait-C**

The wait completion definition for the logical core. It can be:

#### **YES**

If the wait completion flag has been set on, indicating that the logical core will keep running on a real core until its time slice is complete, even though it is waiting

#### **NO**

If the wait completion flag is off

### **Cap**

The PR/SM initial capping setting for the partition.

#### **YES**

An initial cap is set. The partition is capped to its entitlement.

#### **NO**

An initial cap is not set.

### **%Load**

The relative load of the partition on the whole processor complex, expressed as the relation of the total time that logical cores were dispatched for this partition compared to the totally available processor time (number of cores times elapsed time), expressed in percent. If the enhanced LPAR monitor support for secondary cores such as zIIPs, IFLs, and ICFs is not available, then a load value will be inserted for only those partitions having CP cores, while dots will be shown for the relative load of partitions not having any CP cores.

### **CID**

A hexadecimal representation of the logical core ID.

### **%Busy**

Percentage of time the logical core was busy, defined as the percentage of elapsed time during which a real core was assigned to the logical core. Note that the value will also include wait time in the system that is running in the partition if wait completion has been enabled for the processor.

**%Ovhd**

Percentage of elapsed time the logical core spent for LPAR management.

**Type**

The core type of the logical processors defined for the partition. Possible values are:

**CP**

Central Processor

**ICF**

Internal Coupling Facility

**IFL**

Integrated Facility for Linux

**ZIIP**

IBM Z Integrated Information Processor

**Note:** For the ZIIP cores, the monitor data shows IIP. Performance Toolkit translates this value as ZIIP.

**AbsCap**

The absolute capping value PR/SM is enforcing for cores of this type. A value of 100 represents one whole physical core's worth of power.

If the partition does not have an absolute cap, you will see dots in this field.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

**MT**

The maximum multithreading level supported in the partition.

When running Performance Toolkit against z/VM 6.3 or earlier data, you will see dashes in this field.

**GrpCapNm**

The group capping name assigned to this partition.

If the partition does not have a group cap, you will see dots in this field.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

**GrpCap**

The group capping value assigned to this partition.

If the partition does not have a group cap, you will see dots in this field.

When running Performance Toolkit against z/VM 6.3 or earlier data, you will see dashes in this field.

**Usage note:**

The term *core* describes a unit that contains one or more processors. All processors of a core have the same processor type. For machine types that do not have the multithreading (MT) facility, the machine has physical processors and the partitions have logical processors. For machine types that have the MT facility, the machine has physical cores, each of which contains two physical processors. A partition that has not activated the MT facility continues to have logical processors. A partition that has activated the MT facility has logical cores, each of which contains one or more logical processors.

**Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages are displayed for all processor load fields.

## FCX127, User Response Time Screen – URESP

**Performance Toolkit Subcommand:** “URESP” on page 344.

The User Response Time Screen (FCX127) shows user response time information for each active user on the system who completed at least one transaction in a sample interval. Information for users which have been active in the interval is shown highlighted on the screen.

FCX127	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor					
-----	Think	<- Response Time ->	<- Transact/Hour ->							
-----	Time	Triv. N-Triv	Mean	Trv/h Ntrv/h Totl/h						
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----
>System<	9.2	.075 .756	.682	780 6402 7182	26	213	239			17
<b>User Class Data:</b>										
FCONXs	5.2	... .221	.221	0 690 690	0	23	23			1
General	27.0	.140 .119	.124	330 1032 1362	11	34	45			6
Service	28.3	.032 .000	.026	390 90 480	13	3	16			4
TCPs	3.6	.000 .992	.980	60 4590 4650	2	153	155			6
<b>User Data:</b>										
AUTOBUCK	...	.001 ...	.001	30 0 30	1	0	1			
CFT2NDA	95.6	... .884	.884	0 132 132	0	4	4			
CFT2NDA0	26.2	.154 .001	.129	300 60 360	10	2	12			
DTCVSW1	35.7	... .000	.000	0 180 180	0	6	6			
DTCVSW2	35.8	... .000	.000	0 180 180	0	6	6			
FTPSEVE	60.0	... .000	.000	0 90 90	0	3	3			
FTPSRV50	60.0	.000 ...	.000	60 0 60	2	0	2			
MONWRITE	8.0	... .032	.032	0 480 480	0	16	16			
MPROUTE	.4	... 1.729	1.729	0 1800 1800	0	60	60			
OPERATOR	16.6	.032 ...	.032	210 0 210	7	0	7			
PVM	...	.001 ...	.001	30 0 30	1	0	1			
REXECD	60.0	... .000	.000	0 90 90	0	3	3			
RSCS	29.9	.037 ...	.037	150 0 150	5	0	5			
RSTL3	5.2	... .221	.221	0 690 690	0	23	23			
Select a user for user details										
<b>Command ===&gt;</b>										
<b>F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return</b>										

Figure 31. Layout of User Response Time Screen (FCX127)

The information shown is based on CP monitor USER domain EVENT data.

Note that EVENT data collection for all users can cause considerable overhead on a large system. Activate this part of data collection only for periods where you really need the information for analysis.

The sample shows a screen with the users sorted in ascending userid sequence, the initial setting (indicated by the underscore characters above the Userid heading). Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the SORT subcommand with the appropriate column heading as sort argument (see “SORT” on page 295), or by positioning the cursor on the heading of the field that is to be sorted and then pressing Enter.

The first entry is reserved for the overall system average for all users, regardless of the current sorting sequence. It is identified by: >System<-userid.

Some additional lines with user class averages will be shown at the top when user classes have been defined (see “FCONTROL UCLASS” on page 89). The *userid* shown will then be the user class name.

For more detailed information about a specific user, do one of the following:

1. Place the cursor on the Userid field of the virtual machine you are interested in and press Enter
2. On the command line, enter the subcommand: USER *userid*

### Field Description:

The following field descriptions apply to the URESP and URESPLOG *userid* displays, with the exception of the Userid field, which is replaced by the Interval End Time field on the URESPLOG display.

#### Userid

The user identification.

**Think Time**

The think time in seconds, that is, the time between transactions. Think times of more than 15 minutes are ignored when calculating this value.

*Response Time* shows response times in seconds:

**Triv.**

The response time for trivial transactions.

Trivial transactions with a duration of more than 5 seconds are ignored.

**N-Triv**

The response time for non-trivial transactions.

Non-trivial transactions with a duration of more than 60 seconds are ignored.

**Mean**

The average response time for all transactions.

*Transact/Hour* shows transaction rates, in transactions per hour:

**Trv/h**

The trivial transaction rate.

**Ntrv/h**

The non-trivial transaction rate.

**Totl/h**

The total transaction rate.

*Total Trans:*

**Trv-T**

The total number of trivial transactions.

**Ntrv-T**

The total number of non-trivial transactions.

**Totl-T**

The total number of transactions.

**Nr of Users**

The number of users that contributed to the average shown on this line (values are inserted only for lines with averages, this field is blank for user data lines).

**Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages or totals will be shown. All of a user's sessions during the selected period contribute to the averages shown. Note also that all of the activity fields (rates) are based on each user's connection time within the overall period, *not* the total elapsed time for the entire period.

## FCX128, Graphics Selection Menu Screen – GRAPHICS

---

**Performance Toolkit Subcommand:** “GRAPHICS” on page 184.

See “Graphics Selection Menu” on page 958 for more information on this specific display, and Appendix A, “Graphical Displays of Performance History Data,” on page 955 for more general information.

## FCX129, User Resources per Transaction Screen – UTRANS

**Performance Toolkit Subcommand:** “UTRANS” on page 357.

The User Resources per Transaction Screen (FCX129) shows the average user resource consumption per transaction for each user. Information for users which have been active in the interval is shown highlighted on the screen.

FCX129	CPU nnnn	SER nnnnn	Interval	HH:MM:SS	- HH:MM:SS	Perf. Monitor					
-----	<-Transact->		<----- Resource Utilization per Transaction ----->								
	Total	Rate	CPU	Disk	Page	Xpage	<-Spooling->	<-IUCV/VMCF->			
-----	Count	/hour	msec	I/O	Flts	Reads	SP_RD	SP_WR	IU_BY	IU_TO	Nr of Users
>System<	660	6600	3.3	4.9	7.9	.0	.0	.0	.7	.7	20
<b>User Class Data:</b>											
FCONXs	54	540	22.7	.2	.5	.0	.0	.0	1.2	1.2	1
General	104	1040	6.4	30.8	49.9	.0	.0	.0	.5	.5	8
Service	41	410	.3	.0	.0	.0	.0	.0	.7	1.0	5
TCPs	461	4610	.7	.0	.0	.0	.0	.0	.7	.7	6
<b>User Data:</b>											
AUTOBUCK	0	0	...	...	...	...	...	...	...	...	...
CFT2NDA	7	75	55.3	62.9	741.4	.0	.0	.0	.0	.0	...
CFT2NDA0	28	280	6.9	30.6	.0	.0	.0	.1	.0	.0	...
DISKACNT	0	0	...	...	...	...	...	...	...	...	...
DTCVSW1	13	130	.0	.0	.0	.0	.0	.0	.0	.0	...
DTCVSW2	13	130	.0	.0	.0	.0	.0	.0	.0	.0	...
EREP	0	0	...	...	...	...	...	...	...	...	...
FTPSEVER	6	60	.0	.0	.0	.0	.0	.0	.0	1.0	...
FTPSRV50	5	50	.0	.0	.0	.0	.0	.0	.0	1.0	...
MONWRITE	43	430	1.9	44.4	.0	.0	.0	.0	1.1	1.1	...
MPROUTE	125	1250	2.0	.0	.0	.0	.0	.0	1.1	1.1	...
OPERATOR	21	210	.1	.0	.0	.0	.0	.1	.0	.0	...
PVM	0	0	...	...	...	...	...	...	...	...	...
REXECD	8	80	.6	.0	.0	.0	.0	.0	3.8	5.3	...
RSCS	12	120	.6	.0	.1	.0	.0	.0	.0	.0	...
RSCSAUTH	0	0	...	...	...	...	...	...	...	...	...
RSTL3	54	540	22.7	.2	.5	.0	.0	.0	1.2	1.2	...
SMTP	40	400	.1	.0	.0	.0	.0	.0	.0	.2	...
TCPIP	143	1430	.3	.0	.0	.0	.0	.0	1.3	1.2	...
TCPIP50	142	1420	.1	.0	.0	.0	.0	.0	.0	.0	...
Select a user for user details											
Command ==>											
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return											

Figure 32. Layout of User Resources per Transaction Screen (FCX129)

The information shown is based on CP monitor USER domain EVENT data.

Note that EVENT data collection for all users can cause considerable overhead on a large system. Activate this part of data collection only for periods where you really need the information for analysis.

The sample shows a screen with the users sorted in ascending userid sequence, the initial setting (indicated by the underscore characters above above the Userid heading). Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the SORT subcommand with the appropriate column heading as sort argument (see “SORT” on page 295), or by positioning the cursor on the heading of the field that is to be sorted and then pressing Enter.

The first entry is reserved for the overall system average for all users, regardless of the current sorting sequence. It is identified by: >System<-userid.

Some additional lines with user class averages will be shown at the top when user classes have been defined (see “FCONTROL UCLASS” on page 89). The userid shown will then be the user class name.

For more detailed information about a specific user, do one of the following:

1. Place the cursor on the Userid field of the virtual machine you are interested in and press Enter
2. On the command line, enter the subcommand: USER userid

## Field Descriptions:

The following field descriptions apply to the UTRANS and UTRANLOG *userid* screens, with the exception of the Userid field, which is replaced by the field Interval End Time on the UTRANLOG screen.

### Userid

The user identification.

### Transact

This section includes information about the number of transactions and their frequency.

### Total Count

The number of transactions.

### Rate/hour

The transaction rate, in transactions per hour.

### Resource Utilization per Transaction

This section includes information about usage figures for a variety of resources.

### CPU msec

The average CPU used per transaction, in msec.

### Disk I/O

The average number of disk I/O requests per transaction.

### Page Flts

The average number of page faults per transaction.

### Xpage Reads

The average number of pages moved from expanded storage to central storage per transaction.

### SP\_RD

The average number of UR read operations from spool per transaction.

### SP\_WR

The average number of UR write operations to spool per transaction.

### IU\_BY

The average number of IUCV data transfers from this user to another virtual machine, or to CP, per transaction.

### IU\_TO

The average number of IUCV data transfers to this user, per transaction.

### Nr of Users

The number of users that contributed to the average shown on this line (values are inserted only for lines with averages; the field remains empty for normal user data lines).

## Effect of AVERAGE and INTERIM subcommands:

Overall or interim period averages or totals will be shown. All of a user's sessions during the selected period contribute to the averages shown. Note also that all of the activity fields (rates) are based on each user's connection time within the overall period, *not* the total elapsed time for the entire period.

## FCX130, System Counters Screen – SYSTRANS

**Performance Toolkit Subcommand:** “SYSTRANS” on page 317.

The System Transaction Statistics Screen (FCX130):

FCX130	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS						Perf. Monitor			
System Transaction Statistics:												
Response Time Statistics (Percent of Samples Below Threshold)												
Resp. time (s)	<.01	<.02	<.04	<.06	<.10	<.20	<.40	<.60	<1.0	<2.0	<4.0	<6.0
Triv. % in range	69	4	4	2	3	7	8	3	0	0	0	0
Triv. accum. %	69	73	77	79	82	89	97	100	100	100	100	100
N-trv % in range	16	7	5	3	3	10	19	9	11	10	5	1
N-trv accum. %	16	23	28	31	34	44	62	71	83	93	97	99
Resource Usage / Transaction (Percent of Samples Below Threshold)												
CPU msec/trans.	<1	<2	<5	<10	<20	<50	<100	<200	<500	<1000	<2000	<5000
Triv. % in range	55	6	11	9	7	7	3	1	0	0	0	0
Triv. accum. %	55	61	72	82	89	96	98	100	100	100	100	100
N-trv % in range	8	2	13	13	9	16	12	10	9	6	2	1
N-trv accum. %	8	10	23	36	45	61	73	83	92	97	99	100
.												
.												
.												
IUCV xfers/trn.	<1	<2	<4	<6	<10	<20	<40	<60	<100	<200	<400	<600
Triv. % in range	90	3	6	0	1	0	0	0	0	0	0	0
Triv. accum. %	90	93	98	99	100	100	100	100	100	100	100	100
N-trv % in range	94	1	2	1	1	1	0	0	0	0	0	0
N-trv accum. %	94	95	97	98	99	99	100	100	100	100	100	100
VMCF xfers/trn.	<1	<2	<4	<6	<10	<20	<40	<60	<100	<200	<400	<600
Triv. % in range	92	0	5	1	2	1	0	0	0	0	0	0
Triv. accum. %	92	92	96	97	99	99	100	100	100	100	100	100
N-trv % in range	85	0	4	2	5	2	2	0	0	0	0	0
N-trv accum. %	85	85	89	91	96	98	100	100	100	100	100	100
Command ==>												
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return							

Figure 33. Layout of System Transaction Statistics Screen (FCX130)

The information shown is based on CP monitor USER domain EVENT data. It can be displayed only when data collection for that domain is active and when monitor data extraction by Performance Toolkit has been activated.

**Note that EVENT data collection for all users can cause considerable overhead on a large system. Activate this part of data collection only for periods where you really need the information for analysis.**

### Field Description:

This screen gives you some statistics on the frequency with which response times and the resources used per transaction fall within certain ranges of values. This should give you information on the distribution of these values which could not be obtained by giving just simple averages.

Numbers are given in each case for:

#### Triv. % in range

The percentage of all measurements for trivial transactions in the interval which did fall into the respective range of values.

#### Triv. accum. %

The percentage of all measurements for trivial transactions in the interval which was found to be lower than the upper boundary of the range, that is, this is the accumulated total percentage of all measurements below and within the range shown.



**N-trv % in range**

The percentage of all measurements for non-trivial transactions in the interval which did fall into the respective range of values.

**N-trv accum. %**

The percentage of all measurements for non-trivial transactions in the interval which was found to be *lower* than the upper boundary of the range, that is, this is the accumulated total percentage of all measurements below and within the range shown.

**Response Time Statistics (Percent of Samples Below Threshold)** shows the measured response time, in seconds.

**Resource Usage / Transaction (Percent of Samples Below Threshold)** shows the resources used per transaction for the following:

**CPU msec/trans.**

The amount of CPU used per transaction, in msec.

**Disk I/Os/trans.**

The number of disk I/O operations per transaction.

**SPOOL reads/trn.**

The number of UR read operations per transaction.

**SPOOL wrts/trns.**

The number of UR write operations per transaction.

**Page faults/trn.**

The number of page faults encountered per transaction.

**XSTORE reads/tr.**

The number of page-in operations from expanded storage per transaction.

**IUCV xfers/trn.**

The number of IUCV data transfer operations per transaction (both directions).

**VMCF xfers/trn.**

The number of VMCF data transfer operations per transaction (both directions).

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all load fields.

## FCX131, I/O Device Configuration Screen – DEVCONF

**Performance Toolkit Subcommand:** “DEVCONF” on page 168.

In the I/O Device Configuration Screen (FCX131), information for I/O devices which are online is shown highlighted on the screen:

FCX131	CPU nnnn	SER nnnnn	Status	HH:MM:SS	Perf. Monitor							
<---- Ranges ---->												
Device-No	Subch.-ID	Device Type	<- Channel Path Ids ->								Control Unit	Status
04C0-04C3	0093-0096	CTCA	1	2	3	4	5	6	7	8	....	Offline
04C4-04CF	0097-00A2	CTCA	16	.	.	.	.	.	.	.	3088	Online
04D0-04D7	00A3-00AA	CTCA	.	.	.	.	.	.	.	.	....	Offline
04D8-04DB	00AB-00AE	CTCA	16	.	.	.	.	.	.	.	3088	Online
04DC	00AF	CTCA	.	.	.	.	.	.	.	.	....	Offline
04DD-04DE	00B0-00B1	CTCA	16	.	.	.	.	.	.	.	3088	Online
04DF	00B2	CTCA	16	.	.	.	.	.	.	.	3088	Online
0580-0585	00E3-00E8	3480	0B	4B	.	.	.	.	.	.	3480-02	Offline
0586-0587	00E9-00EA	3480	.	.	.	.	.	.	.	.	....	Offline
0588-058D	00EB-00F0	3480	0B	4B	.	.	.	.	.	.	3480-02	Offline
058E-058F	00F1-00F2	3480	.	.	.	.	.	.	.	.	....	Offline
0600-061F	0103-0122	3380-K	06	4E	.	.	.	.	.	.	3880-23	Online
0720-073F	0163-0182	3380-K	07	4F	.	.	.	.	.	.	3880-23	Online
081A	0183	37x5	04	.	.	.	.	.	.	.	3745-D1	Online
0840-085F	0184-01A3	3390-3	14	20	54	60	.	.	.	.	3990-3	Online
0860-087F	01A4-01C3	3390	.	.	.	.	.	.	.	.	....	Offline
0B19	0206	37x5	0D	4D	.	.	.	.	.	.	3745-01	Online
Command ==> _												
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return												

Figure 34. Layout of I/O Device Configuration Screen (FCX131)

The information shown is based on CP monitor MONITOR domain SAMPLE data.

The I/O devices are sorted in ascending device number sequence; that sequence cannot be changed (no SORT subcommand for this screen).

### Field Description:

#### Ranges

Information for a single device or for a range of similar I/O devices, with similar channel paths and status, can be shown on one line. The following data will be shown:

##### Device-No

The device number(s) of the I/O device(s)

##### Subch.-ID

The subchannel ID(s) of the I/O device(s)

#### Device Type

The device type and model number.

The string '(E)' will be inserted for devices which are emulated on some other device type. The string '(X)' will be inserted for devices which support XRC. These flags may be combined, for example, '(EX)'.

#### Channel Path Ids

Shows which channels are active for data transfer to the device. Up to eight channel connections are supported by the architecture; dots will be inserted for the ones which are not active (not installed or offline).

For the emulated FBA devices (EDEV), only the first path ID is shown, even when multiple channel connections to the device are available.

#### Control Unit

The control unit type and model number. For different models of IBM DASD subsystems, they will be recognized and displayed as such only if the only if they are operating in enhanced mode; otherwise they will be displayed as 3990-3 control units.

**Status**

The device status (online or offline).

The information shown is initially provided by a MONITOR domain record when Performance Toolkit connects to the \*MONITOR service. If the program has been set up to collect the monitor data from another machine's disk file, the program will attempt to obtain the monitor domain data from the first monitor sample in the disk file; the data cannot be shown if the required information is not found in that first sample.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX132, User Communication Screen – UCOMM

**Performance Toolkit Subcommand:** [“UCOMM” on page 331.](#)

The User Communication Screen (FCX132) shows information on the use made of the IUCV and VMCF communication facilities by individual users. Information for users which have been active in the interval is shown highlighted on the screen.

FCX132	CPU	nnnn	SER	nnnnn	Interval	HH:MM:SS	-	HH:MM:SS	Perf.	Monitor																
-----	Msg	V	Total	<-	IUCV	Msg/s	-->	<-	VMCF	Msg/s	-->	<--	IUCV	in	Q	-->	VMCF	Connections	<--	Per	Message	----	Nr of			
-----	Target	M	Msg/s	I	Send	I	Recv	I	Fail	V	Send	V	Recv	V	Fail	Send	Recv	Q	Rep1q	Queue	Total	Max.	CPUsec	DasdIO	Paging	Users
User Class Data:																										
>System<	.....	.	.070	.029	.029	.000				.006	.006	.006	.000	.000	.195	.000	.112	2.307	716.3	.0016					.000	38
FCONXs																										
General	.....	.	.183	.092	.092	.000				.000	.000	.000	.000	.538	.000	.000	.000	5.385	64.00	.0084	.084	.000		.000	2	
Service	.....	.	.017	.008	.009	.000				.000	.000	.000	.000	.000	.000	.000	.000	1.687	650.5	.0010	19.20	.000		.000	18	
TCPs	.....	.	.018	.000	.001	.000				.006	.010	.000	.000	.000	.000	.000	.095	2.838	1183	.0001	.000	.000		.000	11	
	.....	.	.256	.108	.108	.000				.023	.017	.000	.000	.857	.000	.000	.429	2.143	338.3	.0005	.000	.000		.000	7	
User Data:																										
AUTOBUCK	.....	N	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	64.00					
CFT2NDA	.....	N	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0					
CFT2NDA0	.....	N	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	64.00					
DISKACNT	.....	N	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	64.00					
DTCVSW1	*VSWITCH	Y	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.000	1024							
DTCVSW2	*VSWITCH	Y	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.000	1024							
EREP	.....	N	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	64.00					
FTPSEIVE	TCPIP	Y	.017	.000	.000	.000	.000	.000	.017	.000	.000	.000	.000	.000	.000	.000	1.000	0	64.00	.0001	.000	.000			.000	
FTPSRV50	TCPIP50	N	.017	.000	.000	.000	.000	.000	.017	.000	.000	.000	.000	.000	.000	.000	1.000	0	64.00	.0001	.000	.000			.000	
GCS	.....	N	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.000	64.00							
HANDYMAN	.....	N	5.556	1.111	4.444	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	0	0	0	.0024	14.40	.000				
MISCSERV	TCPIP	N	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.000	64.00							
MONWRITE	*MONITOR	N	.267	.133	.133	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	0	1.000	64.00	.0009	19.84	.000				
MPROUTE	TCPIP	Y	.733	.367	.367	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	0	3.000	64.00	.0009	.000	.000				
M8USER	.....	N	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.000	64.00							
OPERATOR	SFSFCFT	N	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2.000	64.00							
OPERSYMP	.....	N	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.000	64.00							
PERFSVM	.....	N	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	64.00							
PORTMAP	TCPIP	Y	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.000	64.00							
PVM	.....	N	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.000	4096							
RECOVERY	*BLOCKIO	N	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	5000							
RELAY	.....	N	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.000	64.00							
REXEC	TCPIP	Y	.188	.000	.007	.000	.071	.110	.000	.000	.000	.000	.000	.000	.000	.000	1.000	0	255.0	.0001	.000	.000			.000	
RSCS	.....	Y	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2.000	2048							
RSCSAUTH	.....	Y	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.000	64.00							
RSCSDNS	.....	Y	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	64.00							
RSTL3	*MONITOR	N	.367	.183	.183	.000	.000	.000	.000	.000	.000	.000	.000	1.000	.000	.000	0	10.00	64.00	.0084	.084	.000				
RXAGENT1	.....	Y	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.000	64.00							
SFSFCFT	RSTL3	Y	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	20.00	500.0							
SMTP	*SPL	Y	.017	.000	.000	.000	.000	.017	.000	.000	.000	.000	.000	.000	.000	.000	1.000	2.000	64.00	.0009	.000	.000			.000	
TCPIP	MPROUTE	Y	.990	.388	.388	.000	.143	.071	.000	.000	.000	.000	.000	6.000	.000	.000	0	8.000	1024	.0001	.000	.000			.000	
TCPIP50	FTPSRV50	Y	.017	.000	.000	.000	.017	.000	.000	.000	.000	.000	.000	.000	.000	.000	0	1.000	1024	.0014	.000	.000			.000	
VMNFS	TCPIP	Y	.010	.005	.005	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	0	3.000	1024	.0001	.000	.000			.000	
VMSEVP	*BLOCKIO	Y	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2000							
VMSEVR	RSTL3	Y	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	14.00	2000							
VMSEVS	*BLOCKIO	Y	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2000							
Select a user for user details																										
Command ==>																										
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return																					

Figure 35. Layout of User Communication Screen (FCX132)

The information shown is based on CP monitor USER domain SAMPLE data.

The sample shows a screen with the users sorted in descending sequence of the total message rate (indicated by the underscore characters above the Msg/s header). Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the SORT subcommand with the appropriate column heading as sort argument (see [“SORT” on page 295](#)), or by positioning the cursor on the header of the field which is to be sorted and then pressing Enter.

The first entry is reserved for the overall system average for all users, regardless of the current sorting sequence. It is identified by: >System<-userid.

Some additional lines with user class averages will be shown at the top when user classes have been defined (see “FCONTROL UCLASS” on page 89). The *userid* shown will then be the user class name.

For more detailed information about a specific user, do one of the following:

1. Place the cursor on the `UserId` field of the virtual machine you are interested in and press `Enter`
2. On the command line, enter the subcommand: `USER userid`

## Field Descriptions:

The following field descriptions apply to both the UCOMM and UCOMMLOG *userid* displays, with the exception of the *Userid* field which will be replaced by the field *Interval End Time* on the UCOMMLOG display.

### Userid

The user identification.

### Msg Target

The target address for the last IUCV or VMCF data transfer (not inserted for system average and user class lines)

### SVM

Indicates whether this is a service virtual machine as defined by the SVMSTAT option in the user's directory entry (Y) or not (N).

A dot is inserted if the SVM status is not available for users that logged on after starting data collection.

### Total Msg/s

The total IUCV and VMCF message rate per second

*IUCV Msg/s:*

### ISend

The rate of IUCV data transfers from this user to another user, or to CP

### IRecv

The rate of IUCV data transfers to this user by other users or CP

### IFail

The rate of IUCV data transfer failures

*VMCF Msg/s:*

### VSend

The rate of VMCF data transfers from this user to another user

### VRecv

The rate of VMCF data transfers to this user by other users

### VFail

The rate of VMCF data transfer failures

*IUCV in Q:*

### SendQ

The number of IUCV send requests queued for the user

### RecvQ

The number of IUCV requests in the user's receive queue

### ReplQ

The number of IUCV requests in the user's reply queue

### VMCF Queue

The number of VMCF messages queued for the user

### Connections Total

The number of IUCV connections for the user.

### Connections Max.

The maximum number of IUCV connections allowed.

*Per Message:* Shows load per IUCV or VMCF message (statistical information only, not necessarily related to the IUCV or VMCF data transfer operations).

### CPUsec

The total CPU time spent per IUCV or VMCF data transfer

**DasdIO**

The number of DASD I/Os executed per IUCV or VMCF data transfer

**Paging**

The number of page reads and writes per IUCV or VMCF data transfer

**Nr of Users**

The number of users which contributed to the average shown on this line (values are inserted only for lines with averages; the field remains empty for normal user data lines).

**Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages or totals will be shown for all load fields. All of a user's sessions during the selected period contribute to the averages shown. Note also that all of the activity fields (rates) are based on each user's connection time within the overall period, *not* the total elapsed time for the entire period.

## FCX133, Shared Segments Screen – NSS

**Performance Toolkit Subcommand:** “NSS” on page 243.

In the Shared Segments Screen (FCX133), information for segments for which some activity was found during the last interval is shown highlighted on the screen.

FCX133		CPU nnnn		SER nnnnn		Interval HH:MM:SS - HH:MM:SS				Perf. Monitor													
-----						<-----				Number of		Pages				<-----				Rate per Sec.			
Name	Spool	Creation	<- Users ->							<Resident>		<-Aliases-->											
File	Date	Share	N-Shr	Savd	Ndat	Resvd	Res. R<2G	Count	Lockd	XSTOR	DASD	Trans	Steal	Pgrds	Pgwrt	X-rds	X-wrt	X-mig					
>System<	12	----	----	15	0	689	1573	93	663	0	125	8	0	88	.331	.000	.000	.000	.000	.000	.000	.000	
CMS	3963	2012/10/22	27	0	1298	0	512	1298	0	1279	0	768	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	
CMSFILES	3889	2012/10/22	4	0	768	0	0	631	0	0	0	0	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	
CMSPILES	3890	2012/10/22	37	0	256	0	256	256	0	2	0	0	0	.000	.000	.000	.000	.000	.000	.000	.000	.000	
CMSVMLTB	3891	2012/10/22	37	0	256	0	0	256	0	2	0	0	0	.000	.000	.000	.000	.000	.000	.000	.000	.000	
GCS	5536	2012/10/23	2	0	120	917	256	1037	0	0	0	0	0	.000	.000	.000	.000	.000	.000	.000	.000	.000	
INSTSEG	3894	2012/10/22	37	0	768	0	0	463	0	4	0	256	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	
MONDCSS	3895	2012/10/22	2	0	0	16k	0	0	92	85	0	3	6.43	.000	.000	.000	.000	.000	.000	.000	.000	.000	
SCEE	3900	2012/10/22	3	0	256	0	0	128	0	0	0	0	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	
SCEEX	3901	2012/10/22	3	0	2304	0	0	1986	0	0	0	0	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	
ZCMS	3904	2012/10/22	5	0	1298	0	0	1015	0	0	0	0	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	
Command ==> _																							
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return																		

Figure 36. Layout of Shared Segments Screen (FCX133)

The information on this screen is based on CP monitor STORAGE domain SAMPLE data.

The sample shows a screen with the shared segments sorted in ascending segment name sequence, the initial setting. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the SORT subcommand with the appropriate column heading (see [“SORT” on page 295](#)), or by positioning the cursor on the header of the field that is to be sorted and then pressing ENTER.

The first entry is reserved for the overall averages and totals for all active segments, regardless of the current sorting sequence. It is identified by the "name" >System<.

### Field Descriptions:

## Name

The name of the named saved system (NSS) or discontinuous saved segment (DCSS).

## Spool File

The spool file number of the segment, or, on the >System< line, the number of NSSes or DCSSes over which the >System< line was calculated.

### Creation Date

The date when the segment was created. The string (Member $\pm$ ) is inserted where the information shown is for a member of a segment space.

*Users:*

## Share

The number of users that access the segment in shared mode.

## N-Shr

The number of users that access the segment in non-shared mode.

Note that the Share and N-Shr values are 0 for segment spaces. Information about the number of users is then obtained using a Q NSS NAME ..... MAP subcommand and the resulting information is inserted for each member on the following lines. The number of users is always inserted under Share for segment members.

Pages:

**Sayd**

The number of pages saved in this segment.

## NDat

The number of pages defined with a 'no data saved' page descriptor.

**Resvd**

The number of frames reserved using the SET RESERVED command.

**Res.**

The total number of pages that reside in central storage.

**R<2GB**

The number of pages that reside in central storage below the 2GB line.

**Count**

The number of SXS alias pages associated with pageable FRMTEs. They might be stealable if they are not last translated or locked.

**Lockd**

The number of locked SXS alias pages.

**XSTOR**

The number of expanded storage blocks that are occupied by pages of the segment.

**DASD**

The number of slots that are occupied on auxiliary storage.

*Rate per Sec.:*

**Trans**

The number of page translations per second.

**Steal**

The number of page frames stolen per second.

**Pgrds**

The number of page reads per second.

**Pgwrt**

The number of page writes per second.

**X-rds**

The number of page-ins per second from expanded storage.

**X-wrt**

The number of page-outs per second to expanded storage.

**X-mig**

The number of pages migrated per second from expanded storage to disk.

The string `Initial` sample is inserted for segments where no valid delta values could be calculated because only one sample has been obtained so far.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** The overall or interim period averages are displayed for all load fields. Note that all of the segments shown might not have been loaded during the entire period. However, the activity displayed (page rates, for example) is always based on the common overall elapsed time. This allows the display of each segment's contribution to the overall load, but it might not display the average load for a segment while it was active.

Usage Note: No information on segment members is available from monitor records. The fields displayed are obtained using a Q NSS NAME ..... MAP subcommand to tell you what the segment spaces are used for. Because no performance data for these members is available, the member information is always inserted immediately after the output line for the segment space, regardless of the sort sequence.



## FCX134, Shared Data Spaces Screen – DSPACESH

**Performance Toolkit Subcommand:** “DSPACESH” on page 172.

In the Shared Data Spaces Screen (FCX134), information for data spaces for which some activity was found during the last interval is shown highlighted on the screen:

FCX134	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor														
					<----- Rate per Sec. ----->				<-----Number of----->									
Pages----->	Owning Userid	Data Space Name	Users Permt	Pgstl	Pgrds	Pgwrt	X-rds	X-wrt	X-mig	Total	<--Resid--> Resid	R<2GB	<-Locked--> Lock	L<2GB	<-Aliases--> Count	Lockd	XSTOR	DASD
	SYSTEM	PTRM007A	0	.000	.000	.000	.000	.000	.000	1049k	0	0	0	0	0	0	0	0
	SYSTEM	PTRM007B	0	.000	.000	.000	.000	.000	.000	1049k	0	0	0	0	0	0	0	0
	SYSTEM	PTRM007C	0	.000	.000	.000	.000	.000	.000	1049k	0	0	0	0	0	0	0	0
	SYSTEM	PTRM007D	0	.000	.000	.000	.000	.000	.000	1049k	0	0	0	0	0	0	0	0
	SYSTEM	PTRM007E	0	.000	.000	.000	.000	.000	.000	1049k	0	0	0	0	0	0	0	0
	SYSTEM	PTRM007F	0	.000	.000	.000	.000	.000	.000	1049k	0	0	0	0	0	0	0	0
	SYSTEM	PTRM0070	0	.000	.000	.000	.000	.000	.000	1049k	0	0	0	0	0	0	0	0
	SYSTEM	PTRM0071	0	.000	.000	.000	.000	.000	.000	1049k	0	0	0	0	0	0	0	0
	SYSTEM	PTRM0072	0	.000	.000	.000	.000	.000	.000	1049k	0	0	0	0	0	0	0	0
	SYSTEM	PTRM0073	0	.000	.000	.000	.000	.000	.000	1049k	0	0	0	0	0	0	0	0
	SYSTEM	PTRM0074	0	.000	.000	.000	.000	.000	.000	1049k	0	0	0	0	0	0	0	0
	SYSTEM	PTRM0075	0	.000	.000	.000	.000	.000	.000	1049k	0	0	0	0	0	0	0	0
	SYSTEM	PTRM0076	0	.000	.000	.000	.000	.000	.000	1049k	0	0	0	0	0	0	0	0
	SYSTEM	PTRM0077	0	.000	.000	.000	.000	.000	.000	1049k	0	0	0	0	0	0	0	0
	SYSTEM	PTRM0078	0	.000	.000	.000	.000	.000	.000	1049k	0	0	0	0	0	0	0	0
	SYSTEM	PTRM0079	0	.000	.000	.000	.000	.000	.000	1049k	0	0	0	0	0	0	0	0
	SYSTEM	REAL	0	.000	.000	.000	.000	.000	.000	5243k	0	0	0	0	0	0	0	0
	SYSTEM	SYSTEM	0	.000	.002	.003	.000	.000	.000	524k	1774	101	0	0	1762	0	0	0
	SYSTEM	VDISK\$CFT2NDA0\$0291\$0001	0	.000	.000	.000	.000	.000	.000	7680	161	45	0	0	0	0	0	0
	SYSTEM	VDISK\$RSTL3\$\$\$FB91\$0002	0	.000	.000	.000	.000	.000	.000	113k	80	41	0	0	0	0	0	0
	SYSTEM	VIRTUAL\$FREE\$STORAGE	0	.000	.000	.000	.000	.000	.000	524k	13	6	0	0	0	0	0	2

Command ==>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 37. Layout of Shared Data Spaces Screen (FCX134)

The information shown is based on CP monitor STORAGE domain SAMPLE data.

The sample shows a screen with the shared data spaces sorted in ascending owner name sequence, the initial setting. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the SORT command with the appropriate column heading (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

The first entry is reserved for the overall averages for all the shared data spaces shown, regardless of the current sorting sequence. It is identified by the 'name' >System<.

### Field Description:

#### Owning Userid

userid of the owner of the shared data space.

#### Data Space Name

The name of the data space, or, on the ">System<" line, the number of data spaces over which the ">System<" line was calculated.

#### Users Permt

Number of users permitted to access the data space.

Rate per Sec.:

#### Pgstl

The number of pages stolen per second.

#### Pgrds

The number of pages read from DASD per second.

#### Pgwrt

The number of pages written to DASD per second.

#### X-rds

The number of pages moved from expanded storage to central storage, per second

**X-wrt**

The number of pages moved from central storage to expanded storage, per second.

**X-mig**

The number of pages migrated from expanded storage to DASD, per second.

*Nr of Pages:*

**Total**

The total number of pages defined for the data space.

**Resid**

The number of resident pages.

**R<2GB**

The number of pages resident in central storage below the 2GB line.

**Lock**

The number of locked pages.

**L<2GB**

The number of pages locked in central storage below the 2GB line.

**Count**

The number of SXS alias pages associated with 'pageable FRMTes'. They may be stealable if not last translated or locked.

**Lockd**

The number of locked SXS alias pages.

**XSTOR**

The number of XSTORE blocks occupied by the data space.

**DASD**

The number of slots occupied on auxiliary storage.

The string `Initial` sample will be inserted for data spaces for which no valid delta values could be calculated because only one sample has been obtained so far.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all load fields. Note that not all the data spaces shown may have been loaded during the whole period. However, the activity shown (for example, page rates) is always based on the common overall elapsed time. This allows showing each data space's contribution to the overall load, but it may not show the average load for a data space while it was active.

## FCX135, User Wait State Log Screen – USTLOG

**Performance Toolkit Subcommand:** “USTLOG” on page 353.

The User Wait State Log Screen (FCX135) shows system-wide user wait state information by time:

FCX135 Interval End Time	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor	<-SVM and->												<-%Samples showing->				<-vCPU-->	
>>Mean>>	%ACT	%RUN	%CPU	%LDG	%PGW	%IOW	%SIM	%TIW	%CFW	%TI	%EL	%DM	%IOA	%PGA	%LIM	%OTH	Q0	Q1	Q2	Q3	E0-3	Type	Cnt
08:19:16	32	5	2	0	0	0	0	68	9	3	0	10	0	0	3	0	14	0	2	75	0	MIX	69
08:20:16	34	5	1	0	0	0	0	70	8	4	0	9	0	0	2	0	14	0	2	73	0	MIX	69
08:21:16	37	5	0	0	0	0	0	82	0	4	0	8	0	0	0	0	13	0	0	79	0	MIX	70
08:22:16	36	4	1	0	0	0	0	83	0	4	0	8	0	0	0	0	13	0	0	79	0	MIX	70
08:23:16	36	5	5	0	0	0	0	77	0	4	0	8	0	0	0	0	14	0	0	79	0	MIX	70
08:24:16	31	5	0	0	0	0	0	80	2	4	0	8	0	0	0	0	12	0	0	79	0	MIX	117
08:25:16	21	5	0	0	0	0	0	82	0	4	0	8	0	0	0	0	13	0	0	78	0	MIX	117
08:26:16	26	4	0	0	0	0	0	88	0	0	0	8	0	0	0	0	10	1	24	59	0	MIX	117
08:27:16	9	10	0	0	0	0	0	67	0	0	0	20	0	0	3	0	26	0	25	35	0	MIX	117
08:28:16	11	8	0	0	0	0	0	74	0	0	0	16	0	0	2	0	20	0	24	43	0	MIX	117
08:29:16	29	4	0	0	0	0	0	90	0	0	0	6	0	0	0	0	12	4	41	39	0	MIX	117
08:30:16	13	15	0	0	0	0	0	53	0	0	0	31	0	0	0	0	36	1	0	33	0	MIX	54
08:31:16	16	17	0	0	0	0	12	40	0	6	0	24	0	0	0	0	37	0	0	42	0	MIX	54
08:32:16	17	16	0	0	0	0	10	43	0	7	0	23	0	0	0	0	36	0	0	41	0	MIX	54
08:33:16	16	18	0	0	0	0	10	40	0	7	0	24	0	0	0	0	37	0	0	41	0	MIX	54

Command ==> \_  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 38. Layout of User Wait State Log Screen (FCX135)

The information shown is based on CP monitor USER domain SAMPLE data.

The leftmost column contains the end time:

(format hh:mm:ss)

of the monitor interval for which the values were calculated.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the Interval End Time column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### %ACT

The percentage of of samples during which the user was found to be active, consuming CPU or waiting for some service.

#### %RUN

The percentage of active samples during which the user was found to be running on a logical processor.

#### %CPU

The percentage of active samples during which the user was found to be waiting to run on the logical processor.

#### %LDG

The percentage of active samples during which the user was found in loading state, that is, with a very high page read activity.

#### %PGW

The percentage of active samples during which the user was found in page wait state.

#### %IOW

The percentage of active samples during which the user was found in I/O wait state.

A virtual CPU can express the IOW state only when the virtual CPU is performing I/O using the Start Subchannel (SSCH) instruction. For information about I/O waits incurred in Diagnose I/O, refer to the %SIM field.

**%SIM**

The percentage of active samples during which the user was found in instruction simulation wait, that is, waiting for CP to complete simulation of a privileged instruction.

**%TIW**

The percentage of active samples during which the user was found in test idle wait, that is, waiting to be dropped from the dispatch list.

**%CFW**

The percentage of active samples during which the user was found in console function wait, that is, waiting for CP to run one of the console function commands that allow a user to alter the state and configuration of the virtual machine (ATTACH, DETACH, or LINK, for example).

**%TI**

The percentage of active samples during which the user was found in SVM wait and test idle wait. SVM wait state indicates that the user was waiting for the completion of a communication with a service virtual machine.

**%EL**

The percentage of active samples during which the user was found in SVM wait and in the eligible list (resource wait).

**%DM**

The percentage of active samples during which the user was found in SVM wait and in the dormant list. (This status is considered to be active because the machine is waiting for something, and is not idle.)

**%IOA**

The percentage of active samples during which the user was found waiting for something other than the above while an I/O operation started by this machine had not yet completed.

**%PGA**

The percentage of samples showing page-active wait, that is, the virtual machine has loaded a wait PSW while waiting for the completion of a page read operation. This state is to be expected for guest systems running with PAGEX ON.

**%LIM**

The percentage of active samples during which the user was found in the limit list, that is, in the dispatch list waiting for CPU but prevented from running by a maximum share setting.

**%OTH**

The percentage of active samples during which the user was found in other states.

*%Samples showing:* The columns under this header show the relative frequency with which a machine was found in the different lists and classes:

**Q0**

The percentage of active samples during which the user was found in Q0 (in the dispatch list as a class 0 user).

**Q1**

The percentage of active samples during which the user was found in Q1 (in the dispatch list as a class 1 user).

**Q2**

The percentage of active samples during which the user was found in Q2 (in the dispatch list as a class 2 user).

**Q3**

The percentage of active samples during which the user was found in Q3 (in the dispatch list as a class 3 user).

**E0-3**

The percentage of active samples during which the user was found in the eligible list (any class).

**vCPU Type**

Indicates the type of all monitored users' virtual processors:

**CP**

Central Processor

**ICF**

Internal Coupling Facility

**IFL**

Integrated Facility for Linux

**ZIIP**

IBM Z Integrated Information Processor

**MIX**

The MP users have different PU types.

This field contains a value during the observed interval.

**vCPU Cnt**

The count of all distinct defined users' virtual processors during the observed interval.

**Effect of AVERAGE and INTERIM subcommands:** None.

## FCX136, Response Time Log (All) Screen – RESPALL

**Performance Toolkit Subcommand:** “RESPALL” on page 271.

The Response Time Log (All) Screen (FCX136) shows system-wide response time information by time:

FCX136	CPU nnnn					SER nnnnn					Interval HH:MM:SS - HH:MM:SS					Perf. Monitor				
<----- Response Time (s) -----> <----- Transaction rate (/s) -----> Total <----- Percent																				
----->	<- Triv. ->		<- N-Trv ->		Quick		<- Triv. ->		<- N-Trv ->		Quick		Trans-	<-Triv.->		<-N-Trv->				
Interval	UP	MP	UP	MP	Disp	Mean	UP	MP	UP	MP	Disp	Total	actns	UP	MP	UP	MP	Qdsp		
>>Mean>>	.090	.000	.747	.205	.925	.756	.17	.00	.02	.19	1.31	1.70	2343	10.1	.0	1.4	11.4	77.1		
19:54:42	.072	.000	.832	.252	.918	.764	.16	.00	.02	.19	1.34	1.71	1334	9.1	.0	1.4	11.3	78.1		
19:55:42	.019	.000	.629	.092	1.032	.826	.12	.00	.03	.20	1.20	1.55	93	7.5	.0	2.2	12.9	77.4		
19:56:42	.037	.000	1.256	.008	.886	.724	.10	.00	.02	.20	1.25	1.57	94	6.4	.0	1.1	12.8	79.8		
19:57:42	.029	.000	.000	.463	1.216	1.074	.05	.00	.00	.18	1.15	1.38	83	3.6	.0	.0	13.3	83.1		
19:58:42	.022	.000	.000	.463	1.090	.957	.07	.00	.00	.18	1.15	1.40	84	4.8	.0	.0	13.1	82.1		
19:59:42	.021	.000	.642	.008	1.005	.797	.12	.00	.03	.20	1.22	1.57	94	7.4	.0	2.1	12.8	77.7		
20:00:42	.050	.000	.641	.009	1.058	.802	.18	.00	.03	.20	1.18	1.60	96	11.5	.0	2.1	12.5	74.0		
20:01:42	.031	.000	1.270	.464	.993	.852	.13	.00	.02	.18	1.23	1.57	94	8.5	.0	1.1	11.7	78.7		
20:02:42	.059	.000	.623	.009	.781	.622	.15	.00	.03	.20	1.30	1.68	101	8.9	.0	2.0	11.9	77.2		
20:03:42	.043	.000	.003	.009	.689	.491	.47	.00	.03	.20	1.62	2.32	139	20.1	.0	1.4	8.6	69.8		
20:04:42	.297	.000	.623	.010	.739	.566	.52	.00	.03	.20	1.43	2.18	131	23.7	.0	1.5	9.2	65.6		
See RESPUP display for UP response time details																				
Command ==>																				
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return																				

Figure 39. Layout of Response Time Log (All) Screen (FCX136)

The information shown is based on CP monitor SYSTEM domain SAMPLE data.

The leftmost column contains the end time (in *hh:mm:ss* format) of the monitor interval for which the values were calculated.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

*Response Time (s)* provides information on response times, in seconds:

#### Triv. UP

The average response time in the interval for trivial transactions of uniprocessor virtual machines, in seconds

#### Triv. MP

The average response time in the interval for trivial transactions of multiprocessor virtual machines, in seconds

#### N-Trv UP

The average response time in the interval for non-trivial transactions of uniprocessor virtual machines, in seconds

#### N-Trv MP

The average response time in the interval for non-trivial transactions of multiprocessor virtual machines, in seconds

#### Quick Disp

The average response time in the interval for virtual machines with the quick dispatch option (QUICKDSP), in seconds

#### Mean

The average response time in the interval for all transactions, in seconds

*Transaction Rate (/s)* provides information on transaction rates, in transactions per second:

**Triv. UP**

The transaction rate for trivial transactions of uniprocessor virtual machines, in transactions per second

**Triv. MP**

The transaction rate for trivial transactions of multiprocessor virtual machines, in transactions per second

**N-Trv UP**

The transaction rate for non-trivial transactions of uniprocessor virtual machines, in transactions per second

**N-Trv MP**

The transaction rate for non-trivial transactions of multiprocessor virtual machines, in transactions per second

**Quick Disp**

The transaction rate for virtual machines with the quick dispatch option (QUICKDSP), in transactions per second

**Total**

The total transaction rate for all transactions

**Total Transactns**

The total number of transactions in the interval

*Percent* shows the relative frequency of the different transaction types:

**Triv. UP**

The percentage of all transactions which were trivial UP transactions

**Triv. MP**

The percentage of all transactions which were trivial MP transactions

**N-Trv UP**

The percentage of all transactions which were non-trivial UP transactions

**N-Trv MP**

The percentage of all transactions which were non-trivial MP transactions

**Quick Disp**

The percentage of all transactions which were issued by virtual machines with the quick dispatch option

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX137, Response Time Log (UP) Screen – RESPUP

**Performance Toolkit Subcommand:** “RESPUP” on page 272.

The Response Time Log (UP) Screen (FCX137) shows system-wide response time information of uniprocessor virtual machines by time:

FCX137	CPU nnnn			SER nnnnn			Interval			HH:MM:SS - HH:MM:SS			Perf. Monitor		
Interval	<- Resp. Time -->			<- Trans. Rate -->			<-- Transactions -->			<-- % -->					
End Time	Trv_s	Ntr_s	Ave_s	Trv/s	Ntr/s	Tot/s	Triv.	N-Trv	Total	Triv	N-Tr	%			
>>Mean>>	.107	1.345	.688	14.29	12.16	26.45	282991	240733	523724	54.0	46.0				
09:05:00	.042	1.192	.562	21.68	17.89	39.57	13005	10734	23739	54.8	45.2				
09:15:00	.041	1.371	.632	19.93	15.95	35.88	11956	9572	21528	55.5	44.5				
09:25:00	.051	1.478	.677	21.19	16.57	37.76	12715	9941	22656	56.1	43.9				
09:35:00	.049	.736	.371	20.48	18.06	38.54	12288	10836	23124	53.1	46.9				
09:45:00	.045	.674	.339	19.43	17.07	36.50	11676	10258	21934	53.2	46.8				
09:55:00	.075	1.429	.721	16.62	15.14	31.76	9972	9086	19058	52.3	47.7				
10:05:00	.086	1.104	.567	17.54	15.68	33.22	10524	9406	19930	52.8	47.2				
10:15:00	.090	1.246	.630	17.33	15.21	32.54	10398	9124	19522	53.3	46.7				
10:25:00	.085	1.055	.543	17.19	15.34	32.53	10332	9218	19550	52.8	47.2				
10:35:00	.094	1.077	.556	17.92	15.88	33.81	10754	9529	20283	53.0	47.0				
10:45:00	.051	.909	.449	17.92	15.49	33.41	10754	9294	20048	53.6	46.4				
10:55:00	.054	.731	.369	17.53	15.27	32.80	10519	9160	19679	53.5	46.5				
11:05:00	.101	1.337	.678	18.94	16.57	35.51	11361	9944	21305	53.3	46.7				
11:15:00	.088	.985	.497	19.51	16.31	35.82	11723	9802	21525	54.5	45.5				
11:25:00	.108	1.753	.863	18.15	15.38	33.53	10891	9229	20120	54.1	45.9				
11:35:00	.107	1.345	.688	15.42	13.64	29.06	9252	8186	17438	53.1	46.9				
See RESPALL display for UP + MP response times															
Command ==>															
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return										

Figure 40. Layout of Response Time Log (UP) Screen (FCX137)

The information shown is based on CP monitor SYSTEM domain SAMPLE data.

The leftmost column contains the end time:

(format hh:mm:ss)

of the monitor interval for which the values were calculated.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

*Resp. Time* shows information on response times for uniprocessor virtual machines:

#### Trv\_s

The average response time in the interval for trivial UP transactions, in seconds

#### Ntr\_s

The average response time in the interval for non-trivial UP transactions

#### Ave\_s

The average response time in the interval for all UP transactions, in seconds

*Trans. Rate:*

#### Trv/s

The average transaction rate in the interval for trivial UP transactions, in transactions per second

#### Ntr/s

The average transaction rate in the interval for non-trivial UP transactions, in transactions per second



**Tot/s**

The total transaction rate in the interval for all UP transactions, in transactions per second

*Transactions:*

**Triv.**

The number of trivial UP transactions

**N-Trv**

The number of non-trivial UP transactions

**Total**

The total number of UP transactions in the interval

'%' shows the relative frequency of the different transaction types:

**Triv**

The percentage of all UP transactions which were trivial transactions

**N-Tr**

The percentage of all UP transactions which were non-trivial transactions

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX138, Minidisk Cache Log Screen – MDCACHE

**Performance Toolkit Subcommand:** [“MDCACHE” on page 233.](#)

### The Minidisk Cache Log Screen (FCX138):

FCX138	CPU nnnn	SER nnnnn	Interval	HH:MM:SS	- HH:MM:SS	Perf. Monitor					
Interval	CP	----	MDC Blocks	----	----	MDC Requests	----	----	----	----	MDC B
End Time	Par-		Max Actual	Fair Share	Reads /s	Full Hit%	Part. Hit%	Miss%	Reads /s	Read %	Hit %
>>Mean>>	458240	131072	130922	3082	409	95.8	..	...	..	...	...
09:15:00	458240	131072	131072	916	503	94.9	..	...	..	...	...
09:25:00	458240	131072	131072	862	560	95.6	..	...	..	...	...
09:35:00	458240	131072	131072	873	486	96.1	..	...	..	...	...
09:45:00	458240	131072	131072	1032	482	95.9	..	...	..	...	...
09:55:00	458240	131072	131072	970	406	94.2	..	...	..	...	...
10:05:00	458240	131072	131072	992	433	96.2	..	...	..	...	...
10:15:00	458240	131072	131072	1129	404	95.8	..	...	..	...	...
10:25:00	458240	131072	131026	1016	404	95.8	..	...	..	...	...
10:35:00	458240	131072	131055	978	430	96.5	..	...	..	...	...
10:45:00	458240	131072	131072	1065	432	95.6	..	...	..	...	...
10:55:00	458240	131072	131072	1159	403	94.3	..	...	..	...	...
11:05:00	458240	131072	131072	834	465	96.4	..	...	..	...	...
11:15:00	458240	131072	131072	970	465	96.7	..	...	..	...	...
11:25:00	458240	131072	131072	916	448	95.8	..	...	..	...	...
11:35:00	458240	131072	131072	1394	384	97.0	..	...	..	...	...
See MDCSTOR display for MDC storage utilization log											
Command ==>											
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F10=Left	F11=Right	F12=Return				

Figure 41. Layout of Minidisk Cache Log Screen (FCX138)

The information shown is based on CP monitor SYSTEM domain SAMPLE data.

The display shows minidisk cache usage by time. The leftmost column contains the end time:

```
(format hh:mm:ss)
```

of the monitor interval for which the values were calculated. Note that only part of the performance variables can be shown on the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window to the left or to the right for viewing the remaining data columns.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

## Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

## CP Partition

The size of the CP partition in expanded storage, in blocks

*MDC Blocks:*

## MDC Blocks Max

The maximum number of expanded storage blocks that can be used for minidisk caching

### MDC Blocks Actual

The number of expanded storage blocks used for minidisk caching at the end of the interval

## MDC Blocks Fair Share

The fair share limit for the minidisk cache, i.e. the max. number of blocks any one user is allowed to insert into the minidisk cache per fair share interval

*MDC Requests:*

**Reads/s**

The number of read requests per second

**Full Hit%**

The percentage of requests which were full hits, i.e. where all the requested blocks were found in the cache

**Part. Hit%**

This value is no longer available.

**Miss%**

This value is no longer available.

*MDC Blocks:*

**Reads/s**

This value is no longer available.

**Read %**

This value is no longer available.

**Hit %**

This value is no longer available.

**PGIN/s**

The rate at which minidisk cache blocks are moved from expanded storage to central storage

**PGOUT/s**

The rate at which minidisk cache blocks are moved from central storage to expanded storage

**Block Life**

The estimated average age, in seconds, of a minidisk cache block, based on Little's Law (cache size divided by PGOUT rate). Dots will be shown if the age cannot be calculated because the PGOUT rate is zero.

**Fair Share Exc/s**

The rate, per second, that blocks could not be moved into the minidisk cache because their user's fair share limit was exceeded

**In Trnsit Ins/s**

Insertions per seconds into the in-transit waiting queue (which keeps track of users requesting blocks which are already on their way to the minidisk cache)

*Invalidates:*

**Request Rate/s**

The number of requests per second to invalidate minidisk cache blocks, due to an I/O to a virtual device via a non-cachable I/O interface

**Blocks Rate/s**

The number of blocks per second which were invalidated following an invalidation request.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX139/140/141/142, CP IUCV Services Log Screen – CPSALL, CPSTOCP, CPSTOVM, CPSFAIL

**Performance Toolkit Subcommand:** “CPSALL” on page 157.

The CP IUCV Services Log Screen (FCX139/140/141/142) is shown after entering the CPSALL, CPSTOCP, CPSTOVM, CPSFAIL subcommand. Note that the layout of these four displays is almost identical. Only a CPSALL sample is included here.

FCX139	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor																	
<-Total rates/s-> <----->					Total System Services Activity, Rates/s ----->																
Interval	IUCV	VMCF	IUCV Total	*MSG	*BLO-	*MON	*SIG-	*SYM	*ACC	*LOG	*ID-	*CON	*VSW	*ASY	*SC	*VME					
End Time	+VMCF	VMCF	IUCV Total	ALL	*MSG	*RPI	CKIO	ITOR	NAL	*CCS	*SPL	PTOM	OUNT	REC	*CRM	ENT	FIG	ITCH	NCMD	LP	VENT
>>Mean>>	318.7	.2	318.4	305.8	0	.0	0	305.2	.6	0	0	0	0	0	0	0	0	0	0	0	0
19:54:42	2.5	.3	2.2	.5	0	.0	0	0	.5	0	0	0	0	0	0	0	0	0	0	0	0
19:55:42	2.2	.1	2.1	.5	0	0	0	0	.5	0	0	0	0	0	0	0	0	0	0	0	0
19:56:42	2.2	.1	2.1	.5	0	0	0	0	.5	0	0	0	0	0	0	0	0	0	0	0	0
19:57:42	2.2	.1	2.1	.5	0	0	0	0	.5	0	0	0	0	0	0	0	0	0	0	0	0
19:58:42	2.3	.1	2.1	.5	0	0	0	0	.5	0	0	0	0	0	0	0	0	0	0	0	0
19:59:42	2.2	.1	2.1	.5	0	0	0	0	.5	0	0	0	0	0	0	0	0	0	0	0	0
20:00:42	3.2	.1	3.0	.6	0	0	0	0	.6	0	0	0	0	0	0	0	0	0	0	0	0
20:01:42	940.9	.1	940.7	917.8	0	0	0	917.3	.5	0	0	0	0	0	0	0	0	0	0	0	0
20:02:42	487.6	.1	487.5	464.8	0	0	0	464.3	.5	0	0	0	0	0	0	0	0	0	0	0	0
20:03:42	2125	.1	2125	2063	0	0	0	2062	.7	0	0	0	0	0	0	0	0	0	0	0	0
20:04:42	1456	.1	1456	1380	0	0	0	1379	.6	0	0	0	0	0	0	0	0	0	0	0	0
20:05:42	1091	.1	1091	1051	0	0	0	1051	.6	0	0	0	0	0	0	0	0	0	0	0	0
20:06:42	906.8	.1	906.7	876.7	0	0	0	876.1	.6	0	0	0	0	0	0	0	0	0	0	0	0
20:07:42	1183	.1	1182	1143	0	0	0	1143	.9	0	0	0	0	0	0	0	0	0	0	0	0
20:08:42	368.0	.1	367.9	349.9	0	0	0	348.9	1.0	0	0	0	0	0	0	0	0	0	0	0	0

See CPSTOVM / CPSTOCP / CPSFAIL for other CP comm. data subsets

Command ==>>>

F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return F14=Help

Figure 42. Layout of CP IUCV Services Log Screen (FCX139/140/141/142)

The information shown is based on CP monitor SYSTEM domain SAMPLE data.

The display shows usage of CP IUCV services by time. The leftmost column contains the end time (in hh:mm:ss format) of the monitor interval for which the values were calculated.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

## Field Description:

### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

**Total rates/s:** Information on overall IUCV and VMCF activities, in transfers per second (the strings Total CP>VM/s, Total VM>CP/s and Total Fails/s will be shown instead for the CPSTOVM, CPSTOCP and CPSFAILS displays respectively).

### IUCV+VMCF

The sum of all IUCV and VMCF activity

### VMCF

The sum of all VMCF activity

### IUCV

The sum of all IUCV activity

**Total System Services Activity, Rates/s:** Provides more detailed information on the participation of specific CP IUCV system services to the total IUCV activity. The strings

- System Services CP to Users, Rates/s
- System Services Users to CP, Rates/s
- System Services Failures, Rates/s

will be shown instead for the CPSTOVM, CPSTOCP and CPSFAILS displays respectively.

**Total**

The sum of all IUCV activity by CP services

**\*MSGALL**

Overall activity for the \*MSGALL IUCV service (Message All System Service)

**\*MSG**

Overall activity for the \*MSG IUCV service (Message System Service)

**\*RPI**

Overall activity for the \*RPI IUCV service (Access Verification System Service)

**\*BLOCKIO**

Overall activity for the \*BLOCKIO IUCV service (DASD Block I/O System Service)

**\*MONITOR**

Overall activity for the \*MONITOR IUCV service (Monitor System Service)

**\*SIGNAL**

Overall activity for the \*SIGNAL IUCV service (Signal System Service)

**\*CCS**

Overall activity for the \*CCS IUCV service (Console Communication Service)

**\*SPL**

Overall activity for the \*SPL IUCV service (Spool System Service)

**\*SYMPTOM**

Overall activity for the \*SYMPTOM IUCV service (Symptom System Service)

**\*ACCOUNT**

Overall activity for the \*ACCOUNT IUCV service (Account System Service)

**\*LOGREC**

Overall activity for the \*LOGREC IUCV service (Error Logging System Service)

**\*CRM**

Overall activity for the \*CRM IUCV service (Collection Resource Management System Service)

**\*IDENT**

Overall activity for the \*IDENT IUCV service (Identify System Service)

**\*CONFIG**

Overall activity for the \*CONFIG IUCV service (Enterprise Systems Connection Manager Service)

**\*VSWITCH**

Overall activity for the \*VSWITCH IUCV service (Virtual Switch System Service)

**\*ASYNCMD**

Overall activity for the \*ASYNCMD IUCV service (Asynchronous CP Command Response System Service)

**\*SCLP**

Overall activity for the \*SCLP IUCV service (SCLP System Service)

**\*VMEVENT**

Overall activity for the \*VMEVENT IUCV service (VM Event System Service)

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX143, Paging Log Screen – PAGELOG

Performance Toolkit Subcommand: “PAGELOG” on page 244.

The Paging Log Screen (FCX143):

FCX143	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor																
<----- Expanded Storage -----> <-Real Stor-> <----- Paging to DASD -----> <Page Table>																					
Interval	Paging	PGIN	Fast-	Path	PGOUT	Total	Est.	Page	DPA	Est.	Reads	Write	Total	Shrd	Guest	Systm	Total	Reads	Writes	Reads	Writes
End Time	Blocks	/s	%	/s	/s	/s	Life	Migr	Pagable	Page	/s	/s	/s	/s	/s	/s	/s	/s	/s	/s	/s
>>Mean>>	0	.0	.0	.0	.0	.0	....	.0	77975k	31025	2818	2513	5331	34.2	437.4	29.6	467.0	13.1	10.1		
04:39:19	0	.0	.0	.0	.0	.0	....	.0	77975k	9138	5815	8533	14348	32.3	245.1	27.6	272.7	22.9	22.2		
04:44:19	0	.0	.0	.0	.0	.0	....	.0	77975k	8105	10991	9620	20611	53.4	165.8	48.7	214.5	43.5	38.4		
04:49:19	0	.0	.0	.0	.0	.0	....	.0	77975k	15281	2834	5103	7937	20.2	354.8	15.4	370.1	8.6	1.5		
04:54:19	0	.0	.0	.0	.0	.0	....	.0	77975k	69467	1039	1123	2161	56.3	186.7	51.8	238.5	37.2	5.6		
04:59:19	0	.0	.0	.0	.0	.0	....	.0	77975k	60094	1717	1298	3015	19.9	405.0	15.8	420.8	6.3	8.5		
05:04:19	0	.0	.0	.0	.0	.0	....	.0	77975k	82892	2123	940.7	3064	18.2	563.9	14.1	578.1	8.2	7.4		
05:09:19	0	.0	.0	.0	.0	.0	....	.0	77975k	35795	2599	2178	4777	21.5	712.0	17.1	729.1	10.4	7.7		
05:14:19	0	.0	.0	.0	.0	.0	....	.0	77975k	27891	4169	2796	6964	26.5	1241	22.4	1263	16.6	14.7		
05:19:19	0	.0	.0	.0	.0	.0	....	.0	77975k	39141	3347	1992	5339	24.9	695.6	20.7	716.3	15.2	10.8		
05:24:19	0	.0	.0	.0	.0	.0	....	.0	77975k	37289	2631	2091	4722	27.4	447.6	16.8	464.4	11.7	12.9		
05:29:19	0	.0	.0	.0	.0	.0	....	.0	77975k	85402	1489	913.0	2402	17.1	222.1	13.2	235.3	6.8	7.1		
05:34:19	0	.0	.0	.0	.0	.0	....	.0	77975k	73999	1082	1054	2135	12.6	162.5	8.8	171.4	4.8	4.0		
05:39:19	0	.0	.0	.0	.0	.0	....	.0	77975k	83231	1166	936.8	2103	14.2	182.0	10.2	192.1	5.3	4.7		
05:44:19	0	.0	.0	.0	.0	.0	....	.0	77975k	131k	649.8	593.9	1244	13.4	78.9	8.1	87.0	2.8	2.9		
05:49:19	0	.0	.0	.0	.0	.0	....	.0	77975k	53973	1849	1445	3294	17.2	378.0	13.4	391.4	8.5	6.4		
05:54:19	0	.0	.0	.0	.0	.0	....	.0	77975k	34778	2797	2242	5039	24.5	795.1	19.8	814.9	12.1	10.5		
Command ==> _																					
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return																					

Figure 43. Layout of Paging Log Screen (FCX143)

The information shown is based on CP monitor SYSTEM and STORAGE domain SAMPLE data.

The display shows paging activity by time. The leftmost column contains the end time:

(format hh:mm:ss)

of the monitor interval for which the values were calculated. Note that only part of the performance variables can be shown on the 80 columns of an ordinary screen: use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window to the left or to the right for viewing the remaining data columns.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### Expanded Storage:

#### Paging Blocks

The number of paging blocks in expanded storage

#### PGIN/s

The PGIN rate, i.e. the number of expanded storage blocks moved to central storage, per second

#### Fast Path %

The percentage of PGIN operations which could be handled by the fastpath routine

#### PGOUT/s

The PGOUT rate, i.e. the number of central storage pages moved to expanded storage, per second. This figure does not include PGOUT operations done for the minidisk cache.

#### Total/s

The total of the PGIN and PGOUT rates

**Est. Life sec**

The estimated life of a paging block in expanded storage, in seconds, based on Little's law (i.e. number of paging blocks divided by the PGOUT rate)

**Page Migr**

The rate at which pages are migrated from expanded storage to auxiliary storage, for dormant and non-dormant users, for shared systems and for the CP system address space

*Real Storage:*

**DPA Pgable Frames**

Number of pageable frames in the DPA

**Est. Page Life**

The estimated average age of a page in central storage, in seconds, based on Little's law (DPA size divided by the sum of PGOUT and page write rates)

*Paging to DASD:*

**Reads/s**

The page read rate per second, including directory, nucleus and warmstart I/O, but no spool I/O

**Write/s**

The page write rate per second, including directory, nucleus and warmstart I/O, but no spool I/O

**Total/s**

The total paging rate per second (sum of reads and writes)

**Shrd**

The number of pages read per second for shared segments

**Single Reads Guest**

The rate of single page reads, per second, for virtual machines

**Single Reads System**

The rate of single page reads, per second, for the system (system pages are always read as single pages)

**Single Reads Total**

The total rate of single page reads (the sum of guest reads and system reads)

*Page Table Management:*

**Reads/s**

The number of page table translations per second that resulted in a page read from auxiliary storage.

**Writes/s**

The number of page tables paged out to auxiliary storage per second.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX144, Processor Log Screen – PROCLOG

Performance Toolkit Subcommand: “PROCLOG” on page 259.

The Processor Log Screen (FCX144) displays processor activity by time.

FCX144		CPU nnnn		SER nnnnn		Interval HH:MM:SS - HH:MM:SS				Perf. Monitor													
		<--- Percent Busy --->				<--- Rates per Sec.--->				<----- PLDV ----->				<----- Paging ----->				<Co>		<Di>			
Interval	End Time	CPU	Type	Total	User	Syst	Emul	Inst Siml	DIAG	SIGP	SSCH	Pct Em-when Mast	Mean Non0 only	<-VMDBK-> Mast Stoin	To Mast	Below 2GB	PGIN	Fast Path	Page Reads	<mm> Msgs	<ag> X'9C'		
>>Mean>>	00 CP	28.2	27.3	.9	26.8	145.3	14.7	758.3	57.0	---	---	---	---	59.50	---	.0	.0	....	.0	2.6	.0		
>>Mean>>	01 CP	27.1	26.6	.5	26.3	143.4	10.0	307.7	1.3	---	---	---	---	111.2	---	.0	.0	....	.0	.2	.0		
>>Mean>>	02 CP	12.4	11.6	.8	11.0	166.2	10.9	435.1	1.4	---	---	---	---	125.4	---	.0	.0	....	.0	5.0	.0		
>>Mean>>	03 CP	12.0	11.0	1.0	10.4	182.4	12.0	392.8	2.0	---	---	---	---	173.3	---	.0	.0	....	.0	.7	.0		
>>Mean>>	04 CP	11.4	10.5	.9	9.9	180.3	12.8	396.0	1.3	---	---	---	---	151.1	---	.0	.0	....	.0	.7	.0		
>>Mean>>	05 CP	1.0	.9	.1	.9	113.6	10.7	60.4	.0	---	---	---	---	9.882	---	.0	.0	....	.0	.1	.0		
>>Mean>>	06 CP	1.1	1.0	.1	.9	129.8	17.8	71.2	.0	---	---	---	---	8.191	---	.0	.0	....	.0	.2	.0		
>>Mean>>	07 CP	.6	.5	.0	.5	20.6	.1	13.6	.0	---	---	---	---	7.891	---	.0	.0	....	.0	.0	.0		
>>Mean>>	08 CP	.7	.6	.0	.6	23.5	.1	11.1	.0	---	---	---	---	6.480	---	.0	.0	....	.0	.0	.0		
>>Mean>>	09 CP	8.2	8.1	.1	8.0	85.0	3.7	77.8	1.0	---	---	---	---	23.69	---	.0	.0	....	.0	.0	.0		
>>Mean>>	0A CP	7.5	7.3	.2	7.2	99.5	1.7	100.0	.9	---	---	---	---	29.51	---	.0	.0	....	.0	.0	.0		
>>Mean>>	0B CP	2.1	1.8	.3	1.6	119.5	2.2	167.6	.4	---	---	---	---	33.67	---	.0	.0	....	.0	.1	.0		
>>Mean>>	0C ZIIP	.0	.0	.0	.0	.4	.0	1.4	.0	---	---	---	---	.000	---	.0	.0	....	.0	.0	.0		
>>Mean>>	0D ZIIP	.0	.0	.0	.0	1.0	.0	1.7	.0	---	---	---	---	.000	---	.0	.0	....	.0	.0	.0		
>>Mean>>	0E IFL	.1	.1	.0	.1	97.7	.0	26.9	.0	---	---	---	---	.576	---	.0	.0	....	.0	.0	.0		
>>Mean>>	0F IFL	.2	.1	.1	.1	139.8	.0	36.1	.0	---	---	---	---	.685	---	.0	.0	....	.0	.0	.0		
>>Mean>>	10 IFL	.3	.2	.1	.1	194.7	.1	51.4	.0	---	---	---	---	.175	---	.0	.0	....	.0	.0	.0		
>>Mean>>	11 IFL	.0	.0	.0	.0	.0	.0	.0	.0	---	---	---	---	.000	---	.0	.0	....	.0	.0	.0		
>>Total>	12 CP	112.3	107.3	4.9	104.1	1409	96.8	2792	65.4	---	---	---	---	739.8	---	.0	.0	....	.0	9.9	.0		
>>Total>	4 IFL	.6	.4	.2	.2	432.2	.1	114.4	.0	---	---	---	---	1.436	---	.0	.0	....	.0	.0	.0		
>>Total>	1 ZIIP	.0	.0	.0	.0	1.4	.0	3.1	.0	---	---	---	---	.000	---	.0	.0	....	.0	.0	.0		
14:55:07	00 CP	29.0	27.8	1.1	27.3	102.7	1.9	1067	51.3	---	---	---	---	63.33	---	.0	.0	....	.0	.6	.0		
14:55:07	01 CP	31.8	31.2	.6	30.8	124.6	.4	327.8	.1	---	---	---	---	130.3	---	.0	.0	....	.0	.2	.0		
14:55:07	02 CP	13.9	13.1	.8	12.5	157.1	1.2	498.8	.4	---	---	---	---	107.3	---	.0	.0	....	.0	1.0	.0		
14:55:07	03 CP	13.7	12.6	1.1	11.9	195.5	4.6	388.3	1.8	---	---	---	---	211.9	---	.0	.0	....	.0	.7	.0		
14:55:07	04 CP	12.0	11.0	1.0	10.3	166.2	1.6	428.6	.6	---	---	---	---	164.0	---	.0	.0	....	.0	.5	.0		
14:55:07	05 CP	.9	.7	.2	.6	220.1	56.4	144.8	.1	---	---	---	---	15.15	---	.0	.0	....	.0	.0	.0		
14:55:07	06 CP	1.0	.8	.2	.7	350.0	139.2	239.8	.0	---	---	---	---	14.80	---	.0	.0	....	.0	.3	.0		
14:55:07	07 CP	2.3	2.2	.0	2.2	.0	.0	2.9	.0	---	---	---	---	9.766	---	.0	.0	....	.0	.0	.0		
14:55:07	08 CP	2.5	2.5	.0	2.5	.2	.1	1.8	.0	---	---	---	---	1.417	---	.0	.0	....	.0	.0	.0		
14:55:07	09 CP	2.8	2.8	.1	2.7	125.1	.0	56.5	.0	---	---	---	---	14.65	---	.0	.0	....	.0	.0	.0		
14:55:07	0A CP	2.9	2.8	.1	2.7	283.4	.1	113.8	.0	---	---	---	---	6.267	---	.0	.0	....	.0	.0	.0		
14:55:07	0B CP	.2	.2	.1	.1	216.2	.0	96.1	.0	---	---	---	---	3.833	---	.0	.0	....	.0	.0	.0		
14:55:07	0C ZIIP	.0	.0	.0	.0	.4	.0	1.4	.0	---	---	---	---	.000	---	.0	.0	....	.0	.0	.0		
14:55:07	0D ZIIP	.0	.0	.0	.0	.0	.0	1.1	.0	---	---	---	---	.000	---	.0	.0	....	.0	.0	.0		
14:55:07	0E IFL	.2	.1	.1	.1	103.3	.0	28.4	.0	---	---	---	---	.184	---	.0	.0	....	.0	.0	.0		
14:55:07	0F IFL	.2	.1	.1	.1	144.7	.1	36.3	.1	---	---	---	---	.067	---	.0	.0	....	.0	.0	.0		
14:55:07	10 IFL	.0	.0	.0	.0	.0	.0	1.2	.0	---	---	---	---	.100	---	.0	.0	....	.0	.0	.0		
14:55:07	11 IFL	.0	.0	.0	.0	.0	.0	.0	.0	---	---	---	---	.000	---	.0	.0	....	.0	.0	.0		
Command ==>>																							
F1=Help		F4=Top		F5=Bot		F7=Bkwd		F8=Fwd		F12=Return													

Figure 44. Layout of Processor Log Screen (FCX144)

**Note:** This report supports monitor data from z/VM 6.2 or earlier. When running Performance Toolkit against z/VM 6.3 or later data, the PROCLOG subcommand is available for compatibility purposes only; PROCLOG should be used instead. See “PROCLOG” on page 251 and “FCX304, Processor Log Screen – PROCLOG” on page 825.

The information shown is based on CP monitor SYSTEM and PROCESSOR domain SAMPLE data.

On this screen, the leftmost column contains the end time (in hh:mm:ss format) of the monitor interval for which the values were calculated.

The displayed data consists of three parts:

1. Overall data grouped by processor number (>>Mean>>) since the last RESET subcommand, or since data collection was activated. This section shows data for every processor that was active in the overall interval. All rates and percent loads are calculated by active period time; that is, between the first and last active times.
2. Overall data grouped by processor type (>>Total>) since the last RESET subcommand, or since data collection was activated. This section shows data for all processors types that were active in the overall interval when the count of active processors in this interval was more than one. All rates and percent loads are calculated by total interval time since the last RESET subcommand or since data collection was activated.



3. Regular data based on sample interval time when BYTIME interval is set to zero, or the BYTIME interval when it is specified.

**Field Descriptions:****Interval End Time**

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> and >>Total> entries, you can find the relevant interval in the header line.

The >>Mean>> line shows the overall average values per processor number. The >>Total> line shows the overall average values per processor type when the count of active processors in overall interval is more than one.

**CPU**

The logical processor ID in hexadecimal format.

The >>Total> line contains a decimal value that represents the total number of logical processors of this type in the partition.

**TYPE**

The processor type, as follows:

**CP**

Central Processor

**ICF**

Internal Coupling Facility

**IFL**

Integrated Facility for Linux

**ZIIP**

IBM Z Integrated Information Processor

**Percent Busy:****Total**

The percentage of total elapsed time that the processor was busy.

**User**

The percentage of total elapsed time that the processor time was attributed to specific users (includes emulation time and CP time).

**Syst**

The percentage of total elapsed time that the processor was charged to the system (CP time not attributed to specific users).

**Emul**

The percentage of total elapsed time that the processor was in emulation mode, doing work for users.

**Note:** All of the above percentages are based on *elapsed time*. For a system running in an LPAR, this could mean that it is running at maximum capacity even though none of the values have reached 100%. See [“FCX100, CPU Load and Transactions – CPU” on page 371](#) for "busy" values, which are based on the sum of "busy" and voluntary wait time, instead of on elapsed time.

**Rates per Sec.:****Inst Siml**

The number of instructions simulated per second.

**DIAG**

The number of Diagnose instructions that guests executed per second.

**SIGP**

The rate at which this CPU is the target of SIGP External Call instructions.

**SSCH**

The number of Start Subchannel or Resume Subchannel instructions that CP executed per second.

*PLDV:*

**Pct Empty**

The percentage of elapsed time that the local dispatch vector had no VMDBKs.

**Mean when Non0**

The mean number of VMDBKs in the processor local dispatch vector when it was not empty.

**VMDBK Mast only**

The mean number of VMDBKs in the master-only local dispatch vector when it was not empty.

**VMDBK Staln/s**

The rate at which VMDBKs were stolen from this processor's local dispatch vector, to be dispatched on another processor.

**To Mast/s**

The rate, per second, at which VMDBKs had to be moved to the master processor for the execution of 'master only' work.

The processor with a nonzero value in this field is the current master processor (usually the IPL processor).

*Paging:*

**Below 2GB/s**

The number of pages moved below the 2GB line for translation, per second. The field is available only for z/VM systems running in 64-bit mode.

**PGIN/s**

The number of PGINs per second (page moves from expanded to central storage).

**Fast Path %**

The percentage of PGINs which could be handled by fastpath logic.

**Page Reads/s**

The number of pages read from auxiliary storage, including directory, nucleus and warmstart I/O but no spool I/O.

*Msgs/s:*

**Comm Msgs/s**

The number of IUCV and VMCF data transfers per second (successful and unsuccessful) for this processor.

*Diag:*

**Diag X'9C'/s**

The number of times Diagnose X'9C' was issued by CP per second for this processor.

**Effect of AVERAGE and INTERIM subcommands:**

None.

## FCX145, Scheduler Queue Log Screen – SCHEDLOG

**Performance Toolkit Subcommand:** “SCHEDLOG” on page 281.

The Scheduler Queue Log Screen (FCX145) shows information on scheduler queue lengths by time:

FCX145	CPU	nnnn	SER	nnnn	Interval	HH:MM:SS	-	HH:MM:SS	Perf.	Monitor														
Total <-- Users in Dispatch List ---> Lim <- In Eligible List --> Class 1 Sum of Sum of <----- Storage (Pages)																								
----->																								
Interval	VMDBK				<- Loading -->			it	<Loading->			Elapsed	Abs.	Rel.	Total	<-----			Total	WSS	----->			
End Time	in Q	Q0	Q1	Q2	Q3	Q0	Q1	Q2	Q3	Lst	E1	E2	E3	E1	E2	E3	T-Slice	Shares	Shares	Consid	Q0	Q1	Q2	Q3
>>Mean>>	2.5	2.0	.3	.0	.2	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.145	0%	5789	5125k	28103	863	0	2527
19:54:42	1.9	1.7	.2	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.165	0%	5164	5127k	23936	38	0	0
19:55:42	2.0	1.0	1.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.156	0%	3100	5127k	3847	377	0	0
19:56:42	2.0	2.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.154	0%	6000	5127k	31976	0	0	0
19:57:42	3.0	3.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.153	3%	6000	5127k	40539	0	0	0
19:58:42	2.0	2.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.152	0%	6000	5127k	31976	0	0	0
19:59:42	2.0	2.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.153	0%	6000	5127k	31976	0	0	0
20:00:42	2.0	2.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.150	0%	6000	5127k	31977	0	0	0
20:01:42	5.0	3.0	2.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.147	0%	7700	5127k	34375	528	0	0
20:02:42	2.0	2.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.146	0%	6000	5127k	31978	0	0	0
20:03:42	3.0	3.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.137	0%	7500	5127k	34379	0	0	0
20:04:42	2.0	2.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.126	0%	6000	5127k	31978	0	0	0
20:05:42	3.0	3.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.117	0%	7500	5127k	34383	0	0	0
20:06:42	3.0	3.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.113	0%	7500	5127k	34383	0	0	0
20:07:42	5.0	3.0	2.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.114	0%	7550	5109k	34383	24766	0	0
20:08:42	2.0	2.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.101	0%	6000	5109k	31978	0	0	0
20:09:42	5.0	2.0	.0	.0	3.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.094	0%	6100	5119k	31978	0	0	36405
20:10:42	6.0	2.0	1.0	.0	3.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.086	0%	6200	5119k	31978	538	0	41940
20:11:42	2.0	2.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1.080	0%	6000	5127k	31999	0	0	0
Command ==>> _																								
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return																								

Figure 45. Layout of Scheduler Queue Log Screen (FCX145)

The information shown is based on SYSTEM domain SAMPLE data.

The leftmost column contains the end time (in *hh:mm:ss* format) of the monitor interval for which the values were calculated.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### Total VMDBK in Q

The total number of VMDBKs found in both the dispatch and eligible lists

*Users in Dispatch List* shows the number of VMDBKs in the dispatch list for each class:

#### Q0

Number of class 0 (QUICKDSP) users

#### Q1

Number of class 1 users in the dispatch list (interactive)

#### Q2

Number of class 2 users in the dispatch list

#### Q3

Number of class 3 users in the dispatch list

#### Loading Q0

Number of loading class 0 (QUICKDSP) users

#### Loading Q1

Number of loading class 1 users in the dispatch list (interactive)

#### Loading Q2

Number of loading class 2 users in the dispatch list

**Loading Q3**

Number of loading class 3 users in the dispatch list

**Limit Lst**

Number of users in the limit list.

*In Eligible List* shows the number of VMDBKs in the eligible list for each class:

**E1**

Number of class 1 users in the eligible list (interactive)

**E2**

Number of class 2 users in the eligible list

**E3**

Number of class 3 users in the eligible list

**Loading E1**

Number of loading class 1 users in the eligible list (interactive)

**Loading E2**

Number of loading class 2 users in the eligible list

**Loading E3**

Number of loading class 3 users in the eligible list

**Class 1 Elapsed T-Slice**

The length of the class 1 elapsed time slice. This is the elapsed time which a class 1 user can spend in the dispatch list, and which should allow 85% of all transactions to complete. The value is continuously adapted to changing system load conditions, and it can be a good response time indicator as long as you compare only periods with a similar 'job mix'.

**Sum of Abs. Shares**

The sum of the absolute shares set for all users which were in the dispatch list

**Sum of Rel. Shares**

The sum of the relative shares set for all users which were in the dispatch list

*Storage (Pages)* provides information on storage utilization by each class:

**Total Consid**

Number of pages to consider when selecting VMDBKs for the dispatch list.

**Total WSS**

Number of pages in all working sets of a specific class:

**Q0**

Class 0 users

**Q1**

Class 1 users

**Q2**

Class 2 users

**Q3**

Class 3 users

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX146, Auxiliary Storage Log Screen – AUXLOG

Performance Toolkit Subcommand: “AUXLOG” on page 131.

The Auxiliary Storage Log Screen (FCX146):

FCX146	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor										
Interval	<Page Slots>		<Spool Slots>		<Dump Slots>		<----- Spool Files ----->		<Average MLOAD>		<----- Paging ----->			
End Time	Total	Used	Total	Used	Total	Used	<--Created-->	<--Purged-->	Paging	Spooling	<--Overruns-->	Max		
>>Mean>>	Slots	%	Slots	%	Slots	%	Total	/s	Total	/s	msec	msec	Total	/s Queue
>>Mean>>	168716k	28	3004872	41	82555k	28	549	.07	636	.08	134k	.2	112	.01 38651
04:39:19	168716k	25	3004872	41	82555k	28	32	.11	39	.13	134k	.4	7	.02 66645
04:44:19	168716k	26	3004872	41	82555k	28	35	.12	27	.09	134k	.4	17	.06 56333
04:49:19	168716k	27	3004872	41	82555k	28	48	.16	56	.19	134k	.4	9	.03 31137
04:54:19	168716k	27	3004872	41	82555k	28	42	.14	47	.16	134k	.5	2	.01 52908
04:59:19	168716k	27	3004872	41	82555k	28	19	.06	18	.06	134k	.5	1	.00 73842
05:04:19	168716k	27	3004872	41	82555k	28	23	.08	23	.08	134k	.5	1	.00 16057
05:09:19	168716k	27	3004872	41	82555k	28	46	.15	47	.16	134k	.5	4	.01 49070
05:14:19	168716k	27	3004872	41	82555k	28	19	.06	17	.06	134k	.3	4	.01 30910
05:19:19	168716k	27	3004872	41	82555k	28	40	.13	42	.14	134k	.2	5	.02 35241
05:24:19	168716k	27	3004872	42	82555k	28	75	.25	12	.04	134k	.2	2	.01 24277
05:29:19	168716k	27	3004872	42	82555k	28	5	.02	3	.01	134k	.1	4	.01 46670
05:34:19	168716k	28	3004872	42	82555k	28	5	.02	5	.02	134k	.1	0	.00 13558
05:39:19	168716k	28	3004872	42	82555k	28	5	.02	6	.02	134k	.1	1	.00 13908
05:44:19	168716k	28	3004872	42	82555k	28	24	.08	7	.02	134k	.1	1	.00 16240

Command ==> \_  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 46. Layout of Auxiliary Storage Log Screen (FCX146)

The information shown is based on CP monitor SYSTEM domain SAMPLE data.

The display shows usage of auxiliary storage by time. The leftmost column contains the end time:

(format hh:mm:ss)

of the monitor interval for which the values were calculated.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### Page Slots:

##### Total Slots

The total number of slots allocated for PAGE areas

##### Used %

The percentage of allocated slots which are in use

#### Spool Slots:

##### Total Slots

The total number of slots allocated for SPOL areas

##### Used %

The percentage of allocated slots which are in use

#### Dump Slots:

##### Total Slots

The total number of slots allocated for SPOL areas on special DUMP disks

##### Used %

The percentage of allocated dump slots which are in use

*Spool Files:*

**Created Total**

The number of spool files created in the interval

**Created /s**

The rate at which spool files were created in the interval, in files per second

**Purged Total**

The number of spool files purged in the interval

**Purged /s**

The rate at which spool files were purged in the interval, in files per second

*Average MLOAD:* Information on the 'modified device load' value upon which CP bases its decisions when selecting paging and spooling devices (the I/O device service time for read/write requests multiplied by the number of queued requests for that same device):

**Paging msec**

The average MLOAD value for paging devices, in msec

**Spooling msec**

The average MLOAD value for spooling devices, in msec

*Paging:*

**Overruns Total**

The total number of times a paging overrun was detected (that is, the queue length has exceeded the point necessary to keep the bandwidth full).

**Overruns /s**

Paging overruns as a rate.

**Max Queue**

The maximum queue length in the interval.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX147, Virtual Disks in Storage Screen – VDISKS

**Performance Toolkit Subcommand:** “VDISKS” on page 358.

In the Virtual Disks in Storage Screen (FCX147), information for virtual disks in storage for which some activity was found during the last interval is shown highlighted on the screen.

FCX147		CPU nnnn		SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor	
Owning Userid	Devno	Virtual Blocks	Disk LINKS	IO/s V	P	--- Data Space, Rate per Sec. ---	----- Nr of Pages -----
>System<	...	21882	1	.247	-	.000 .019 .106 .000 .000 .000	20 16 0 0 265 17
APCSCIF	0291	60000	1	.000	N	.000 .000 .000 .000 .000 .000	7 0 0 0 0 0
CHG00001	0291	60000	1	2.933	Y	.000 .000 .073 .000 .000 .000	12 12 0 0 887 17
CHG00006	0291	60000	1	1.250	Y	.000 .000 .063 .000 .000 .000	10 1 0 0 434 17
CHG00008	0291	60000	1	.000	Y	.000 .000 .000 .000 .000 .000	0 0 0 0 15 17
D2501	01B1	6000	3	.000	N	.000 .000 .000 .000 .000 .000	47 0 0 0 0 0
D2501	01A1	60000	1	.013	Y	.000 .330 1.667 .000 .000 .000	265 265 0 0 2071 17
D2501	01A2	6000	1	.000	Y	.000 .000 .000 .000 .000 .000	0 0 0 0 3 17
D2501	01A6	6000	1	.000	Y	.000 .000 .000 .000 .000 .000	0 0 0 0 257 17
D2501	01A7	6000	1	.000	Y	.000 .000 .000 .000 .000 .000	0 0 0 0 257 17
D2501	01A8	6000	1	.000	Y	.000 .000 .000 .000 .000 .000	0 0 0 0 47 17
TH200006	01CC	6000	1	.000	Y	.000 .000 .000 .000 .000 .000	0 0 0 0 47 17
TH200009	01CC	6000	1	.000	Y	.000 .000 .000 .000 .000 .000	0 0 0 0 47 17
TH300002	01CC	6000	1	.000	Y	.000 .000 .000 .000 .000 .000	0 0 0 0 47 17
TH300007	01CC	6000	1	.000	Y	.000 .000 .000 .000 .000 .000	0 0 0 0 47 17
TH300010	01CC	6000	1	.000	Y	.000 .000 .000 .000 .000 .000	0 0 0 0 47 17
TH400008	01CC	6000	1	.000	Y	.000 .000 .000 .000 .000 .000	0 0 0 0 47 17

Select a user for user details  
Command ==>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 47. Layout of Virtual Disks in Storage Screen (FCX147)

The information shown is based on CP monitor STORAGE domain SAMPLE data.

The sample shows a screen with the VDISKS sorted in ascending owner name sequence, the initial setting. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the 'SORT' subcommand with the appropriate column heading (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

The first entry is reserved for the overall averages for all the shared data spaces shown, regardless of the current sorting sequence. It is identified by the 'name' >System<.

Because virtual disks in storage are built internally using shared data spaces, this performance display is very similar in its layout to the shared data spaces display, and you will also find all the VDISK data spaces again on the DSPACESH screen, but without the VDISK-specific information.

For more specific information on one VDISK user do one of the following:

- Place the cursor on the Userid field of the virtual machine you are interested in and press ENTER
- Enter the subcommand 'USER userid' on the command line

### Field Descriptions:

#### Owning Userid

The user ID of the owner of the VDISK

#### Virtual Disk:

#### Devno

The virtual device number of the VDISK

#### Blocks

The size of the VDISK, expressed as the number of 512-byte blocks for which it was defined

#### LINKS

The number of LINKS to the virtual disk

#### IO/s

The virtual I/O rate to the VDISK

#### PV

Private VDISK indicator: Y = (Yes) - a private VDISK, N = (No) - a shared VDISK

*Data Space, Rate per Sec.:*

**Pgstl**

The number of pages stolen per second

**Pgrds**

The number of pages read from DASD per second

**Pgwrt**

The number of pages written to DASD per second

**X-rds**

The number of pages moved from expanded storage to central storage, per second

**X-wrt**

The number of pages moved from central storage to expanded storage, per second

**X-mig**

The number of pages migrated from expanded storage to DASD, per second

*Nr of Pages:*

**Resid**

The number of pages resident in central storage

**IBR**

The number of IBR (invalid-but-resident) pages for this VDISK.

**Lockd**

The number of locked pages

**XSTOR**

The number of XSTOR blocks occupied by the data space. (Available only on z/VM 5.2 and later.)

**DASD**

The number of slots occupied on auxiliary storage

**Data Space Name**

The name of the data space, or, on the ">System<" line, the number of VDISKs over which the ">System<" line was calculated.

The string `Initial` sample will be inserted for VDISKs for which no valid delta values could be calculated because only one sample has been obtained so far.

**Effect of AVERAGE and INTERIM subcommands:**

Overall values or interim period averages will be shown for all load fields. Note that not all of the VDISKs might have been defined during the entire period. However, the activity shown (I/O rates, page rates) is always based on the common overall elapsed time. This allows showing each VDISK's contribution to the overall load, but it might not show the average load for a virtual disk while it was active.



## FCX148, User-Defined Performance Screen – FCONTROL DEFSCRN

**Performance Toolkit Subcommand:** “FCONTROL DEFSCRN” on page 24.

This screen can contain selected performance data fields from other standard displays. The specified fields will be copied at the end of each monitor sample interval.

Because the layout of the resulting screen is, by definition, determined by the user, any example illustrates only some of the possibilities. The following figure shows such an example:

FCX148		CPU nnnn		SER nnnnn		Interval		HH:MM:SS - HH:MM:SS		Perf. Monitor			
CPU Load										Vector Facility		Status or	
PROC	%CPU	%CP	%EMU	%WT	%SYS	%SP	%SIC	%LOGLD	%VTOT	%VEMU	REST	ded. User	
P00	46	25	21	54	8	0	97	..	not installed			Master	
P01	48	7	41	52	1	0	95	..	not installed			Alternate	
P02	42	6	36	58	1	0	96	..	not installed			Alternate	
P03	46	7	39	54	1	0	96	..	not installed			Alternate	
P04	43	6	37	57	1	0	96	..	not installed			Alternate	
P05	49	7	42	51	2	0	95	..	not installed			Alternate	
<-- Device Descr. -->													
Addr	Type	Label/ID	Mdisk Links	Pa- ths	<-Rate/s-> I/O Avoid		<----- Time (msec) -----> Pend Disc Conn Serv Resp		Req. CUWt Qued				
4B05	3745	>VTAM80	...	1	32.7	...	.1	1.8	3.2	5.1	5.1	...	.00
1947	3380-K	SPOOLB CP	0	4	13.3	.0	.2	12.0	2.4	14.6	14.6	...	.00
1A74	3380-K	CSYS04	48	4	12.1	10.6	.3	14.7	4.1	19.1	22.4	...	.04
1A67	3380-K	SPOOLA CP	0	4	11.7	.0	.2	12.4	2.4	15.0	15.0	...	.00
CHPID													
Nr.	Cur	Ave	0-10	11-20	21-30	31-40	41-50	51-60	61-70	71-80	81-90	91-100	
55	29	25	0	25	62	12	0	0	0	0	0	0	
15	22	24	0	25	75	0	0	0	0	0	0	0	
Distribution 13:53:12-14:33:12													
Userid	%CPU	TCPU	VCPU	Ratio	%Vec	Total	DASD	Avoid	UR	Pg/s	User	Status	
VSCS80	17.9	53.7	37.8	1.4	.0	.0	.0	.0	.0	.0	ESA,CL0,DISP		
ICAMWSCH	17.1	51.3	50.7	1.0	.0	3.5	2.0	1.6	.0	.3	ESA,---,DORM		
QUA	6.7	20.1	18.9	1.1	.0	9.6	9.5	4.3	.0	.0	ESA,---,DORM		
HWAL	5.9	17.8	16.9	1.1	.0	10.7	10.3	10.1	.0	.0	ESA,---,DORM		
User-defined screen - no HELP text available													
Command ==> _													
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return													

Figure 48. Example of a User-Defined Performance Screen (FCX148)

You can use the FC DEFSCRN command to define your own performance screens. You need to specify:

- The name you want to give to the screen
- Any of your own header lines
- The blocks to be inserted into the screen, where each block requires:
  - The line number of the first and last line of the block in your new screen
  - Source screen name and the first line number on the source screen where copying is to start

The subcommands used for defining the above screen have been included in the sample FCONX \$PROFILE delivered with the program, as follows:

```
FC DEFSCRN SYSSUM LINE 2 TO 9 COPY CPU FROM 1
FC DEFSCRN SYSSUM LINE 11 TO 12 COPY DEVICE FROM 2
FC DEFSCRN SYSSUM LINE 13 TO 16 COPY DEVICE FROM 5
FC DEFSCRN SYSSUM LINE 18 TO 21 COPY CHANNEL FROM 1
FC DEFSCRN SYSSUM LINE 23 TO 23 COPY USER FROM 4
FC DEFSCRN SYSSUM LINE 24 TO 27 COPY USER FROM 6
```

That is, you would select this sample log by entering the SYSSUM command on the command line. See “FCONTROL DEFSCRN” on page 24 for more information.

**Field Descriptions:**

Refer to the descriptions of the performance screens that have contributed to this combined screen.

## FCX149, Monitor Settings – MONSET

**Performance Toolkit Subcommand:** “MONSET” on page 237.

The Monitor Settings report (FCX149):

```

FCX149      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor

Initial Settings      <----- Active ----->      <--- Seconds --->
2003/03/04 08:00      Event      Sample      High      Sample      HF
Nr Domain              ---      YES      Frequency  Interval      Rate
0  SYSTEM              ---      YES      YES        300          1.00
1  MONITOR             YES      YES      ---        300          ---
2  SCHEDULER           NO       ---      ---        300          ---
3  STORAGE             YES      YES      ---        300          ---
4  USER               NO       YES      YES        300          1.00
5  PROCESSOR           YES      SOME     YES        300          1.00
6  I/O                 YES      YES      YES        300          1.00
7  SEEK               YES      ---      ---        ---          ---
8  NETWORK            YES      NO       ---        300          1.00
9  ISFC               YES      YES      ---        300          ---
10 APPLDATA           NO       NO       ---        300          ---
11 SSI                YES      YES      ---        300          ---

Additional features
6  PCIF               YES      YES      ---        300          ---

Changed Monitor Settings
Date Time      Command Line
03/04 08:01:25 MONITOR EVENT DISABLE ALL
03/04 08:01:30 MONITOR EVENT ENABLE STORAGE
03/04 08:00:45 MONITOR SAMPLE ENABLE STORAGE
03/04 08:00:55 MONITOR SAMPLE ENABLE USER ALL
03/04 08:01:05 MONITOR SAMPLE ENABLE PROCESSOR
03/04 08:01:10 MONITOR SAMPLE ENABLE I/O ALL
03/04 08:01:15 MONITOR SAMPLE RATE 6 SECONDS
03/04 08:01:20 MONITOR SAMPLE INTERVAL 5 MINUTES
03/04 08:01:35 MONITOR EVENT ENABLE PROCESSOR
03/04 08:01:40 MONITOR EVENT ENABLE I/O ALL
03/04 08:01:45 MONITOR EVENT ENABLE SEEKS ALL
03/04 08:03:50 MONITOR EVENT DISABLE ALL

Not applicable in Trend Scan mode
... (oldest lines not displayed) ...
*** Status refreshed ***

Command ==> _
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return

```

Figure 49. Layout of Monitor Settings report (FCX149)

The display shows the initial monitor settings and it also shows commands which were entered at a later time to change these settings. The information is based on CP monitor MONITOR domain SAMPLE and EVENT data.

### Field Description:

#### Initial Settings

Shows date (yyyy/mm/dd) and time (hh:mm) of the initial monitor sample for which the settings are displayed.

#### Nr

The number of the domain for which the initial settings are displayed

#### Domain

The name of the domain for which the initial settings are displayed.

#### Active:

The initial status of the different monitor domains is shown under this heading.

**Event**

Shows the initial settings for collecting event records, where:

**YES**

Indicates event record creation is enabled for the domain

**NO**

Indicates event record creation is disabled for the domain

**SOME**

Indicates that only a subset of the users or I/O devices have been enabled for event record creation.

Dashes indicate "not applicable", that is, event data cannot be created for the domain.

**Sample**

Shows the initial settings for collecting sample records, where:

**YES**

Indicates sample record creation is enabled for the domain.

**NO**

Indicates sample record creation is disabled for the domain.

**SOME**

Indicates only a subset of the users or I/O devices has been enabled for sample record creation. For the PROCESSOR domain it indicates that the CPU Measuring Facility Counters are disabled.

Dashes indicate "not applicable", that is, sample data cannot be created for the domain.

**High Frequency**

shows the initial settings for high-frequency sampling, where:

**YES**

Indicates there are high-frequency fields in the domain.

**NO**

Indicates there are high-frequency fields in the domain, but high-frequency sampling is stopped.

Dashes indicate that there are no high-frequency fields in the domain or that the corresponding sample domain is disabled.

**Seconds:**

The following fields show sampling intervals, in seconds.

**Sample Interval**

The interval for collecting sample records, in seconds

**HF Rate**

The high-frequency sampling interval, also in seconds.

**Changed Monitor Settings:**

Here, you will find a collection of the MONITOR subcommands which were entered after collection of the initial monitor settings.

**Date**

The month and day of the change (*mm/dd*).

**Time**

The time (format *hh:mm:ss*) when the MONITOR subcommand was entered.

**Command Line**

The MONITOR subcommand entered. Note that any abbreviations are automatically expanded to the full length.

**Messages:**

These messages could appear in the Changed Monitor Settings area:

**Not applicable in Trend Scan mode**

is displayed when viewing data using the TRNDSCAN command. Because the corresponding change status events are not written to the history files, they are not available for redisplay.

**... (oldest lines not displayed) ...**

is displayed when the count of the change status lines exceeds the REDISP number (set using the FC MONCOLL REDISP command).

**\*\*\* Status refreshed \*\*\***

is displayed when the new configuration record for this particular report has been received.

**Effect of AVERAGE and INTERIM subcommands:**

None.

## FCX150, Shared File System Log Screen – SFSLOG

**Performance Toolkit Subcommand:** “SFSLOG” on page 291.

The Shared File System Log Screen (FCX150):

FCX150		CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS						Perf. Monitor	
				<-- Time per File			Pool Request		<--- Ser		
Interval		FPR	FPR	Block							
End Time	Server	Count	Rate	Total	CPU	Lock	I/O	ESM	Other	Total	
>>Mean>>	RWSERV1	65820	31.4	.025	.000	.000	.019	.000	.006	1.7	
>>Mean>>	RWSERV2	66017	31.5	.023	.000	.000	.018	.000	.005	1.5	
>>Mean>>	ROSERV1	1948	.9	.011	.000	.000	.000	.000	.011	.0	
>>Mean>>	CRRSERV1	782	.4	.012	.000	.000	.010	.000	.002	.0	
06:44:04	RWSERV1	1817	30.3	.023	.002	.000	.020	.000	.001	8.9	
06:44:04	RWSERV2	1545	25.8	.021	.002	.000	.018	.000	.001	7.5	
06:44:04	ROSERV1	60	1.0	.013	.003	.000	.000	.000	.010	.3	
06:44:04	CRRSERV1	22	.4	.011	.001	.000	.010	.000	.000	.0	
06:45:04	RWSERV1	1727	28.8	.027	.002	.000	.017	.000	.008	10.9	
06:45:04	RWSERV2	1948	32.5	.028	.002	.000	.020	.000	.006	18.4	
06:45:04	ROSERV1	36	.6	.052	.003	.000	.000	.000	.049	3.5	
06:45:04	CRRSERV1	15	.3	.012	.002	.000	.008	.000	.002	.0	
06:46:04	RWSERV1	1662	27.7	.021	.002	.000	.019	.000	.000	8.0	
06:46:04	RWSERV2	2008	33.5	.019	.002	.000	.017	.000	.000	9.0	
06:46:04	ROSERV1	60	1.0	.015	.002	.000	.000	.000	.013	.2	
06:46:04	CRRSERV1	24	.4	.009	.001	.000	.008	.000	.000	.0	
Command ==> _											
F1=Help		F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F10=Left	F11=Right	F12=Return			

Figure 50. Layout of Shared File System Log Screen (FCX150)

The information shown is based on CP monitor APPLDATA and USER domain SAMPLE data.

The display shows information on the activity and performance of shared file system servers by time. The leftmost column contains the end time (format *hh:mm:ss*) of the monitor interval for which the values were calculated. Note that only part of the performance variables can be shown on the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window to the left or to the right for viewing the remaining data columns.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column. The relative activity shown is based on each server's actual connect time, not the overall elapsed time in the interval.

### Field Description:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### Server

The userid of the shared file system or CRR server to which the following figures apply.

#### FPR Count

The number of file pool requests handled by the server.

#### FPR Rate

The number of file pool requests handled per second.

#### Time per File Pool Request:

The following fields show the average time spent in a specific area per file pool request, in seconds.

#### Total

The total service time per file pool request.

**CPU**

The CPU time used by the server per file pool request.

**Lock**

The time spent in lock wait per file pool request.

**Block I/O**

The time needed for block I/O per file pool request.

**ESM**

The processing time of an external security manager per file pool request.

**Other**

Other time (the remainder of total service time minus the other times specifically mentioned above).  
If the Other field value is negative, Performance Toolkit displays dots in this field instead of a negative value.

*Server Utilization:*

The following fields show server utilization as a percentage of elapsed time during which the server was busy in the area specified.

**Total**

The total server utilization.

**CPU**

The percentage of time the server was actively using CPU.

**Page Read**

The percentage of time spent in page wait.

**Checkpoint**

The percentage of time spent for file pool checkpointing.

**QSAM**

The percentage of time the server was waiting for QSAM I/O.

*Agents:***Active**

The average number of active agents, calculated as the total service time for all file pool requests divided by elapsed time.

**Held**

The average number of held agents, calculated as the LUW (logical unit of work) holding time divided by elapsed time.

**Deadlocks**

Number of roll-backs due to deadlocks.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX151, Shared File System I/O Log Screen – SFSIOLOG

**Performance Toolkit Subcommand:** “SFSIOLOG” on page 290.

The Shared File System I/O Log Screen (FCX151):

FCX151	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS						Perf. Monitor			
<----- I/Os per File Pool Request ----->												
Interval		FPR	<--File-->		<-Catalog->		<Cntrl MD->		<--Log-->			
End Time	Server	Count	Total	Read	Write	Read	Write	Read	Write	Read	Write	
>>Mean>>	RWSERV1	65820	1.48	.24	.18	.45	.19	.00	.00	.00	.42	
>>Mean>>	RWSERV2	66017	1.49	.24	.18	.45	.19	.00	.00	.00	.42	
>>Mean>>	ROSERV1	1948	.00	.00	.00	.00	.00	.00	.00	.00	.00	
>>Mean>>	CRRSERV1	782	1.00	.00	.00	.00	.00	.00	.00	.00	1.00	
06:44:04	RWSERV1	1817	1.53	.23	.20	.44	.23	.00	.00	.00	.44	
06:44:04	RWSERV2	1545	1.59	.23	.19	.46	.23	.00	.00	.00	.47	
06:44:04	ROSERV1	60	.00	.00	.00	.00	.00	.00	.00	.00	.00	
06:44:04	CRRSERV1	22	1.00	.00	.00	.00	.00	.00	.00	.00	1.00	
06:45:04	RWSERV1	1727	1.40	.25	.18	.41	.13	.00	.01	.00	.42	
06:45:04	RWSERV2	1948	1.45	.25	.19	.43	.17	.00	.01	.00	.41	
06:45:04	ROSERV1	36	.00	.00	.00	.00	.00	.00	.00	.00	.00	
06:45:04	CRRSERV1	15	.93	.00	.00	.00	.00	.00	.00	.00	.93	
06:46:04	RWSERV1	1662	1.60	.25	.19	.48	.24	.00	.00	.00	.43	
06:46:04	RWSERV2	2008	1.41	.26	.18	.41	.19	.00	.00	.00	.38	
06:46:04	ROSERV1	60	.00	.00	.00	.00	.00	.00	.00	.00	.00	
06:46:04	CRRSERV1	24	1.00	.00	.00	.00	.00	.00	.00	.00	1.00	
Command ==> _												
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return												

Figure 51. Layout of Shared File System I/O Log Screen (FCX151)

The information shown is based on CP monitor APPLDATA and USER domain SAMPLE data.

The display shows information on I/O activity of shared file system servers by time. The leftmost column contains the end time:

```
(format hh:mm:ss)
```

of the monitor interval for which the values were calculated. Note that only part of the performance variables can be shown on the 80 columns of an ordinary screen: use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window to the left or to the right for viewing the remaining data columns.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column. The relative activity shown is based on each server's actual connect time, not the overall elapsed time in the interval.

### Field Description:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### Server

The userid of the shared file system or CRR server to which the following figures apply

#### FPR Count

The number of file pool requests handled by the server

#### I/Os per File Pool Request:

The following fields show the number of I/Os which were required for processing an average file pool request.



**Total**

The total number of I/Os required

**File Read**

The number of file read requests per file pool request

**File Write**

The number of file write requests per file pool request

**Catalog Read**

The number of catalog read requests per file pool request

**Catalog Write**

The number of catalog write requests per file pool request

**Cntrl MD Read**

The number of control MD read requests per FPR

**Cntrl MD Write**

The number of control MD write requests per FPR

**Log read**

The number of log read requests per file pool request

**Log write**

The number of log write requests per file pool request

**Blocks/BIO**

The number of blocks in an average block I/O request

**Blocks/IO**

The number of blocks in an average I/O request issued by CP on behalf of the file pool server

**SAC calls/FPR**

The number of calls to the Storage Access Component per file pool request

**SAC calls/FPR**

The number of calls to the Storage Access Component per file pool request

*Mean Time:* The following fields show the average time required for handling certain events, in seconds.

**Block I/O**

The average time per block I/O request

**Lock Wait**

The average duration of a lock wait. This is the time spent waiting for a lock held by another SFS request to become available.

**Checkpoint**

The time required for processing a checkpoint

**ESM Call**

The average time spent per call to an external security manager.

This value will be zero unless the file pool server is started with the ESECURITY parameter in the DMSPARMS file.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX152, Shared File System Requests Log Screen – SFSREQ

**Performance Toolkit Subcommand:** “SFSREQ” on page 293.

The Shared File System Requests Log Screen (FCX152):

FCX152	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS							Perf. Monitor	
<----- File Pool Request Percentages ----->											
Interval		FPR	Open	Open					Get		C
End Time	Server	Count	Read	Update	Read	Write	Close	Delete	Entry	Rename	
>>Mean>>	RWSERV1	65820	19.7	13.2	13.2	4.6	32.9	6.7	.0	.4	
>>Mean>>	RWSERV2	66017	19.8	13.1	13.2	4.6	32.9	6.7	.0	.4	
>>Mean>>	ROSERV1	1948	.0	.0	.0	.0	.0	.0	.0	.0	
>>Mean>>	CRRSERV1	782	.0	.0	.0	.0	.0	.0	.0	.0	
06:44:04	RWSERV2	1545	17.3	13.9	13.9	5.0	31.6	7.2	.0	.6	
06:44:04	ROSERV1	60	.0	.0	.0	.0	.0	.0	.0	.0	
06:44:04	CRRSERV1	22	.0	.0	.0	.0	.0	.0	.0	.0	
06:45:04	RWSERV1	1727	20.4	12.9	13.3	4.2	33.7	6.5	.0	.3	
06:45:04	RWSERV2	1948	19.5	13.2	13.8	4.8	32.4	6.8	.0	.4	
06:45:04	ROSERV1	36	.0	.0	.0	.0	.0	.0	.0	.0	
06:45:04	CRRSERV1	15	.0	.0	.0	.0	.0	.0	.0	.0	
06:46:04	RWSERV1	1662	18.5	14.0	13.1	4.8	32.1	6.7	.0	.5	
06:46:04	RWSERV2	2008	21.5	12.4	12.8	5.1	33.6	6.0	.0	.4	
06:46:04	ROSERV1	60	.0	.0	.0	.0	.0	.0	.0	.0	
06:46:04	CRRSERV1	24	.0	.0	.0	.0	.0	.0	.0	.0	
Command ==> _											
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F10=Left	F11=Right	F12=Return				

Figure 52. Layout of Shared File System Requests Log Screen (FCX152)

The information shown is based on CP monitor APPLDATA domain SAMPLE data.

The display shows information on the nature of file pool requests to shared file system servers by time. The leftmost column contains the end time:

(format hh:mm:ss)

of the monitor interval for which the values were calculated. Note that only part of the performance variables can be shown on the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window to the left or to the right for viewing the remaining data columns.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column. The relative activity shown is based on each server's actual connect time, not the overall elapsed time in the interval.

### Field Description:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### Server

The userid of the shared file system or CRR server to which the following figures apply

#### FPR Count

The number of file pool requests handled by the server

#### File Pool Request Percentages:

The following fields show which percentage of the file pool requests to a server falls into the different request categories.

**Open Read**

The percentage of file pool requests which were 'open file read' requests

**Open Update**

The percentage of file pool requests which were 'open update' requests

**Read**

The percentage of file pool requests which were 'read file' requests

**Write**

The percentage of file pool requests which were 'write file' requests

**Close**

The percentage of file pool requests which were 'close file' requests

**Delete**

The percentage of file pool requests which were 'delete file' requests

**Get Dir Entry**

The percentage of file pool requests which were 'get directory entry' requests

**Rename**

The percentage of file pool requests which were 'rename' requests

**Create Alias**

The percentage of file pool requests which were 'create alias' requests

**Creat/Delete Dir**

The percentage of file pool requests which were 'create directory' or 'delete directory' requests

**O/G/C Dir**

The percentage of file pool requests which were 'open', 'get' or 'close' directory requests

**Grant/Revoke Auth**

The percentage of file pool requests which were 'grant' or 'revoke authorization' requests

**Lock/Unlock**

The percentage of file pool requests which were 'lock' or 'unlock' requests

**Query**

The percentage of file pool requests which were 'query' requests

**Refrsh Dir**

The percentage of file pool requests which were 'refresh directory' requests

**Other**

The percentage of the remaining file pool requests which do not fall into any of the above categories.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX153, Exception Log Screen – EXCEPT

**Performance Toolkit Subcommand:** “EXCEPT” on page 178.

The Exception Log Screen (FCX153):

```
FCX153      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS      Perf. Monitor
10:50:12 FCXDEV320A Device 0853: resp. time 83ms exceeds limit 40ms (9.4 IO/s)
11:55:15 FCXPEU317A User RVS %CPU 48.0 exceeded threshold 40.0 for 2 min.
12:00:38 FCXPEU317A User RVS %CPU 47.3 exceeded threshold 40.0 for 7 min.
13:23:26 FCXPEU317A User GJA %CPU 61.7 exceeded threshold 40.0 for 2 min.
14:03:59 FCXPEU317A User PBO %CPU 57.3 exceeded threshold 40.0 for 2 min.
14:08:08 FCXPEU317A User PBO %CPU 82.1 exceeded threshold 40.0 for 2 min.
14:12:17 FCXPEU317A User PBO %CPU 72.9 exceeded threshold 40.0 for 2 min.
14:39:09 FCXPEU317A User PBO %CPU 90.1 exceeded threshold 40.0 for 2 min.
14:50:33 FCXPEU317A User PCHA %CPU 45.3 exceeded threshold 40.0 for 2 min.
15:00:12 FCXDEV319A Device 0847: queue length .66 exceeds limit of .40
15:00:12 FCXDEV320A Device 0847: resp. time 65ms exceeds limit 40ms (6.1 IO/s)
15:00:12 FCXDEV319A Device 0853: queue length .64 exceeds limit of .40
15:00:12 FCXDEV320A Device 0853: resp. time 91ms exceeds limit 40ms (0.0 IO/s)
16:11:21 FCXPEU338A User ZO in WSS loop during last 10 minutes
16:27:01 FCXPEU338A User ZO in WSS loop during last 10 minutes
16:46:42 FCXPEU317A User HMO %CPU 85.6 exceeded threshold 40.0 for 2 min.
16:50:49 FCXPEU338A User HMO in CPU loop during last 5 minutes
16:51:50 FCXPEU317A User HMO %CPU 82.7 exceeded threshold 40.0 for 7 min.
16:53:54 FCXPEU317A User PBO %CPU 65.2 exceeded threshold 40.0 for 2 min.
17:36:07 FCXPEU317A User CORADS %CPU 68.5 exceeded threshold 40.0 for 2 min.

Command ==>
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 53. Layout of Exception Log Screen (FCX153)

Provided that suitable thresholds have been defined for the system ('FC SET LIMIT ...' subcommand) the display will tell you which performance variables exceeded the set thresholds, by how much, and how long the exceptions have persisted (messages with time stamps).

The last (most recent) part of the system's exception log file will be shown first. Use the scrolling PF-keys or subcommands to browse through the remainder of the day's exception messages.

Error message FCXxxx422E No exceptions logged will be shown instead when the exception log file is empty.

Three generations of the EXCPLOG file are kept: 'yyyymmdd EXCPLOG A', 'yyyymmdd EXCPLOG1 A', and 'yyyymmdd EXCPLOG2 A', for the current day and two previous days respectively.

The files are automatically renamed at midnight, and the oldest one erased; you can copy them onto another disk if you want to keep more than these three generations.

You can control insertion of exception messages into the exception log file to some extent:

- Exception messages can be suppressed by entering the appropriate FC PROCESS ... subcommands with an output display type other than ERRMSG or CPAMSG, or with the NODISP argument.
- Because the system will log *all* alert messages, both of the ERRMSG type created by Performance Toolkit's exception monitoring function, *and* any CP 'alert' messages of the type CPAMSG, you can cause other output lines to be included in the exception log file by changing their output display type to CPAMSG.
- User-generated exception messages of any kind will also be included in the EXCPLOG file. See the “Submitting User-generated Exceptions” section in the [z/VM: Performance Toolkit Guide](#) for details.

### Working with the Exception Log Display:

When exception messages appear in the exception log, you will usually want to proceed with some kind of analysis in order to determine where the problem is and whether it still persists. Code has been included in the program to help you in this task in the case of alert messages which have been generated by the

threshold and user monitoring functions of Performance ToolkitPerformance Toolkit. To begin analysis, do the following:

1. Place the cursor on the time stamp field of the message for the problem you want to analyze.
2. Press the ENTER key.

A selection subcommand will automatically be generated for the performance data display which is most likely to help you in analyzing that particular problem. Examples include::

- The general I/O device display will be selected when the message indicated that the total I/O rate threshold, or the percentage of users in I/O wait, were exceeded
- The CP owned device display will be selected when the message indicated that either the SPOOL or PAGE space utilization threshold was exceeded
- The detailed user display will be selected when the user monitoring function indicated that a specific user was either looping or had exceeded its thresholds

The general CPU screen which will be shown for many of the 'threshold exceeded' messages anyway, will also be shown for all non-Performance ToolkitPerformance Toolkit alert messages, for which the program cannot determine the course of action to take.

Note that the time stamp fields are not intended for entering subcommands (all data entered will be ignored) and you must not enter a subcommand on the command line if you intend to use the selection subcommand. If a subcommand is found on the command line, *only* that subcommand will be run, and the cursor position ignored.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX154, System Settings Screen – SYSSET

**Performance Toolkit Subcommand:** “SYSSET” on page 314.

The System Settings Screen (FCX154):

```
FCX154      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS      Perf. Monitor

Initial Scheduler Settings: 2020/07/22 at 03:17:42
DSPSLICE (minor) 5.000 msec.      IABIAS Intensity      90 Percent
Hotshot T-slice  2.000 msec.      IABIAS Duration       2 Minor T-slices
DSPBUF Q1        32767 Openings    STORBUF Q1 Q2 Q3     300 % Main storage
DSPBUF Q1 Q2     32767 Openings    STORBUF Q2 Q3       275 % Main storage
DSPBUF Q1 Q2 Q3  32767 Openings    STORBUF Q3          250 % Main storage
LDUBUF Q1 Q2 Q3   600 % Paging exp. Max. working set    9999 % Main storage
LDUBUF Q2 Q3      500 % Paging exp. Loading user        5 Pgrd / T-slice
LDUBUF Q3         400 % Paging exp. Loading capacity    9 Paging expos.
LIMITHARD algorithm Consumption
DSPWD method      Reshuffle
Polarization      Vertical
Global Perf. Data ON
EXCESSUSE: CP MEDIUM CPUPAD: CP 8000%
              ZAAP MEDIUM          ZAAP 8000%
              IFL MEDIUM           IFL 8000%
              ICF MEDIUM           ICF 8000%
              ZIIP MEDIUM          ZIIP 8000%

Changed Scheduler Settings
Date Time Changed
07/22 07:03:33 IABIAS 90 2, DSPSLICE 5.000, Hotshot 1.999, LdUser 5, LdCap 5
STORBUF 300 250 200, LDUBUF 100 100 100, DSPBUF 32767 32767 32767
LIMITHARD algorithm Consumption
POLARIZATION: Vertical GPD: ON
EXCESSUSE: CP-MEDIUM ZAAP-MEDIUM IFL-MEDIUM ICF-MEDIUM ZIIP-MEDIUM
CPUPAD: CP-100% ZAAP-100%IFL-100% ICF-100% ZIIP-100

Command ==> _
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return
```

Figure 54. Layout of System Settings Screen (FCX154)

The information on this screen is based on CP monitor MONITOR domain SAMPLE data and on SCHEDULER domain EVENT data. Complete data can be shown only when data collection for both these domains is active.

The display shows the initial system resource manager settings and it also shows changes to SRM settings and users' SHARE and QUICKDSP settings, which were made at a later time.

### Field Descriptions:

The following fields describe the scheduler settings, as they were found initially:

#### Initial Scheduler Settings

Shows the date (*yyyy/mm/dd*) and time (*hh:mm*) of the initial monitor sample for which the settings are displayed.

#### DSPSLICE (minor)

The duration of the minor dispatch time slice, in milliseconds.

#### Hotshot T-slice

The duration of the "hot shot" time slice that a "hot shot" user is given when running an immediate #CP command.

#### DSPBUF Q1

The number of dispatch list openings for class 1 users. DSPBUF values can be set using the SET SRM DSPBUF command.

#### DSPBUF Q1 Q2

The number of dispatch list openings available to class 1 and class 2 users.

**DSPBUF Q1 Q2 Q3**

The number of dispatch list openings available to class 1, class 2, and class 3 users.

**LDUBUF Q1 Q2 Q3**

The percentage of "Load Capacity" that can be taken by class 1, class 2, and class 3 users. LDUBUF values can be set using the SET SRM LDUBUF command.

**LDUBUF Q2 Q3**

The percentage of "Load Capacity" that can be taken by class 2 and class 3 users.

**LDUBUF Q3**

The percentage of "Load Capacity" that can be taken by class 3 users.

**IABIAS Intensity**

The interactive bias intensity. It can be set using the SET SRM IABIAS command.

**IABIAS Duration**

The interactive bias duration that defines the number of minor time slices during which interactive bias is to remain effective. It can be set using the SET SRM IABIAS command.

**STORBUF Q1 Q2 Q3**

The percentage of the dynamic paging area which is available to class 1, class 2, and class 3 users. It can be set using the SET SRM STORBUF command.

**STORBUF Q2 Q3**

The percentage of the dynamic paging area that is available to class 2 and class 3 users.

**STORBUF Q3**

The percentage of the dynamic paging area that is available to class 3 users.

**Max. working set**

The maximum percentage of the dynamic paging area that any single user's working set can occupy. The value can be set using the SET SRM MAXWSS command.

**Loading user**

The number of page reads per minor time slice that will let a user be classified as a "loading user".

**Loading capacity**

The loading user capacity of the system. This value is equal to the number of paging exposures, that is, the number of paging I/O operations that can be executing concurrently.

**LIMITHARD algorithm**

Sets the algorithm that is used within the CP scheduler to enforce hard limiting of scheduled system resources. This setting only affects users with absolute maximum shares defined using the LIMITHARD option.

**DEADLINE**

The classic CP algorithm for enforcing hard limiting. It enforces limits based on a calculated deadline that defines the projected instant in time when a guest should complete a time slice.

**CONSUMPTION**

An algorithm that enforces hard limiting based on the guest's measured processor usage.

**POLARIZATION**

Indicates the partition's polarization mode:

**VERTICAL**

The partition is running in vertical polarization.

**HORIZONTAL**

The partition is running in horizontal polarization.

**GLOBAL Perf. Data**

Indicates whether the partition has access to global performance data (GPD):

**ON**

GPD is enabled.

**OFF**

GPD is disabled.

**N/A**

GPD is not applicable.

**EXCESSUSE**

The SET SRM EXCESSUSE settings, indexed by CPU type.

**CPUPAD**

The SET SRM CPUPAD settings, indexed by CPU type.

*Changed Scheduler Settings:*

A list of changes to scheduler settings and to users' SHARE and QUICKDSP settings that were made after the initial scheduler settings were collected, or after refreshing the settings at the last RESET.

**Date**

The month and day of a change, in *mm/dd* format.

**Time**

The time when the command that changed the settings was entered, in *hh:mm:ss* format.

**Changed**

The new settings after the change. One entry is created for each SET SRM command, SET SHARE command, and SET QUICKDSP command, if SCHEDULER EVENT has been enabled. The information shown is:

*For SET SRM (when running Performance Toolkit against any level of z/VM data):*

- **IABIAS** - the intensity and duration of the interactive bias.
- **DSPSLICE** - the duration of the minor time slice, in milliseconds.
- **Hotshot** - the duration of the "hot shot" time slice.
- **LdUser** - the number of page reads per minor time slice that make a user a "loading user".
- **LdCap** - the loading capacity of the system.
- **STORBUF** - the STORBUF percentages.
- **LDUBUF** - the LDUBUF percentages.
- **DSPBUF** - the DSPBUF percentages.
- **LIMITHARD algorithm** - Consumption or Deadline.

*For SET SRM (when running Performance Toolkit against z/VM 6.3 and later data only):*

- **POLARIZATION** – Vertical or Horizontal.
- **GPD** – the state of Global Performance Data: ON, OFF, or N/A.
- **EXCESSUSE** – the actual SET SRM EXCESSUSE settings.
- **CPUPAD** – the actual SET SRM CPUPAD settings.

*For SET SHARE:*

The user ID for which the relative or absolute SHARE value was set and the new share values are shown.

*For SET QUICKDSP:*

The user ID for which the QUICKDSP setting was changed and the new QUICKDSP status is shown.

**Messages:**

These messages could appear in the Changed Scheduler Settings area:

**Not applicable in Trend Scan mode**

is displayed when viewing data using the TRNDSCAN command. Because the corresponding change status events are not written to the history files, they are not available for redisplay.

**... (oldest lines not displayed) ...**

is displayed when the count of the change status lines exceeds the REDISP number (set using the FC MONCOLL REDISP command).



**\*\*\* Status refreshed \*\*\***

is displayed when the new configuration record for this particular report has been received.

**Effect of AVERAGE and INTERIM subcommands:**

None.

## FCX155, Monitor Data Statistics Screen – MONDATA

Performance Toolkit Subcommand: “MONDATA” on page 236.

The Monitor Data Statistics Screen (FCX155):

FCX155	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor						
Domain	Rec	Samp/Ev	Record Name	<--Records-->		<--Rec. Length-->			<--Total Bytes-->	
				Count	%	Min	Max	Mean	Count	Pct
<b>System</b>	1 Sample		System data	216	1.56	116	116	116	25056	.67
	2 Sample		Processor data	216	1.56	152	152	152	32832	.88
	3 Sample		Global real storage	54	.39	564	564	564	30456	.81
	4 Sample		Real storage (per proc.)	216	1.56	72	72	72	15552	.42
	6 Sample		Global auxiliary storage	54	.39	92	92	92	4968	.13
	7 Sample		Shared storage data	54	.39	72	72	72	3888	.10
	8 Sample		User data	54	.39	108	108	108	5832	.16
	9 Sample		Channel path contention	54	.39	1252	1252	1252	67608	1.81
	10 Sample		Global scheduler data	54	.39	152	152	152	8208	.22
	11 Sample		Processor communications	216	1.56	272	272	272	58752	1.57
	12 Sample		User wait states	54	.39	180	180	180	9720	.26
	13 Sample		Scheduler activity	216	1.56	56	56	56	12096	.32
	14 Sample		Minidisk Cache Data	54	.39	204	204	204	11016	.29
	15 Sample		Logical CPU utilization	54	.39	56	56	56	3024	.08
	16 Sample		CPU util. in an LPAR	1836	13.24	160	2368	503	923616	24.69
	17 Sample		CPU util. for LPAR mgmt.	54	.39	3240	3240	3240	174960	4.68
	19 Sample		System data (global)	54	.39	132	132	132	7128	.19
	20 Sample		Extended channel meas.	3132	22.58	180	180	180	563760	15.07
	21 Sample		SXS data, global	54	.39	84	84	84	4536	.12
	22 Sample		SXS data, by processor	216	1.56	88	88	88	19008	.51
	23 Sample		Formal spin lock data	54	.39	1080	1080	1080	58320	1.56
	24 Sample		Sched. Act. by processor	270	1.95	368	368	368	99360	2.66
<b>Monitor</b>	1 Event		Event profile	1	.01	48	48	48	48	.00
	4 Sample		System configuration	1	.01	420	420	420	420	.01
	5 Sample		Processor configuration	4	.03	56	56	56	224	.01
	6 Sample		Device configuration	4915	35.44	312	312	312	1533480	40.99
	7 Sample		Storage configuration	1	.01	224	224	224	224	.01
	8 Sample		Paging configuration	2	.01	48	48	48	96	.00
	9 Sample		Sample profile	1	.01	48	48	48	48	.00
	11 Sample		Interval end	109	.79	20	20	20	2180	.06
	13 Both		End of frame indicator	920	6.63	20	20	20	18400	.49
	14 Both		Domain detail	2	.01	28	28	28	56	.00
	15 Sample		Logged on users	18	.13	232	232	232	4176	.11
	16 Sample		Scheduler settings	1	.01	132	132	132	132	.00
	20 Sample		HyperPAV pool def.	6	.04	24	24	24	144	.00
	23 Sample		ISFC End Point config.	16	.12	188	188	188	3008	.08
	25 Sample		SSI Configuration	1	.01	108	108	108	108	.00
<b>I/O</b>	26 Sample		System Topology config.	1	.01	124	124	124	124	.00
	27 Sample		PCI Function config.	8	.06	140	140	140	1120	.03
	28 Sample		HyperPAV pool activity	324	2.34	48	48	48	15552	.42
	36 Event		Attach PCI Function	4	.03	44	44	44	176	.00
	37 Event		Detach PCI Function	4	.03	44	44	44	176	.00
	38 Event		Enable PCI Function	4	.03	40	40	40	160	.00
	39 Sample		PCI Function activity	64	.46	128	128	128	8192	.22
	40 Event		Disable PCI Function	4	.03	132	132	132	528	.01
	41 Event		PCI Function error	3	.02	44	44	44	132	.00
	44 Event		Modify PCI Function	4	.03	56	56	56	224	.01
...										
...										
...										
Command ==>										
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return										

Figure 55. Layout of Monitor Data Statistics Screen (FCX155)

The information shown is obtained while scanning CP monitor records during monitor data extraction by Performance Toolkit.

The display provides some statistical information on the monitor data read by Performance Toolkit. Note that a monitor record will appear in this list only if at least one such record has been found in the input data.

## Field Descriptions:

### Domain

The name of the domain. A line with the heading Overall at the bottom of the list shows the values for *all* the scanned monitor records.

### Rec

The record number.

### Samp/Ev

Indicates whether the record is part of SAMPLE or EVENT data collection, or both.

### Record name

The description of the monitor record.

### *Records:*

Information on absolute and relative monitor record counts is shown under this heading.

### Count

The number of records found.

### %

The percentage to which a record contributed to the total record count.

### *Record Length:*

Information on monitor record lengths is shown under this heading.

### Min

The length of the shortest record found for the record type.

### Max

The length of the longest record found for the record type.

### Mean

The average record length for this record type, calculated by dividing the total byte count for the record by the number of records.

### *Total Bytes:*

Information on monitor record space requirements is shown under this heading.

### Count

The number of bytes occupied by the record type.

### %

The percentage to which this record type contributed to the total monitor record space requirements.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period totals of all monitor records will be shown.

## FCX156, DASD Seeks Locations Screen – SEEKLOC

**Performance Toolkit Subcommand:** “SEEKLOC” on page 288.

The DASD Seeks Locations Screen (FCX156) shows DASD seeks statistics on the frequency with which seeks were made to specific locations for each monitored disk:

FCX156	CPU nnnn	SER nnnnn	Interval	HH:MM:SS	- HH:MM:SS	Perf.	Monitor
<----->	Device	<----->	<----->	Seeks	<----->	Percent	
	Device	Volume	Size Disc	in Time	Cach	Pct	Pct
Addr	Type	Serial	Cyls	msec	ed	Read	Non0
>> Mean <<						32	61
						146	174223
						29	24
						9	6
0600	3380-K	PROFSA	2655	21.2	Yes	13	31
0601	3380-K	VMPRA1	2655	19.3	Yes	25	38
0602	3380-K	PROFSB	2655	21.6	Yes	13	50
0603	3380-K	VMPRA3	2655	21.6	Yes	24	48
0604	3380-K	VMPRA4	2655	13.9	Yes	57	48
0605	3380-K	VMPRA5	2655	21.4	Yes	19	50
0606	3380-K	VMPRA6	2655	22.8	Yes	18	49
0607	3380-K	VMPRA7	2655	20.5	Yes	17	44
0608	3380-K	ESA111	2655	10.8	No	5	4
060A	3380-K	VMPRAA	2655	21.6	Yes	17	43
060B	3380-K	VMPRAB	2655	10.7	Yes	54	53
060C	3380-K	VMPG1N	2655	13.3	No	55	98
060D	3380-K	VMPG10	2655	13.4	No	52	99
060E	3380-K	VMPRAE	2655	21.5	Yes	20	41

Select a device for I/O device details  
 Command ==>  
 F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return

Figure 56. Layout of DASD Seeks Locations Screen (FCX156)

The sample shows a screen with disks sorted in ascending device number sequence (indicated by the underscore characters above the 'Addr' header), the initial default sequence. Other possible sort sequences are indicated by dots above the corresponding columns, they can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

The information shown is based on MONITOR and I/O domain SAMPLE data and on SEEKS domain EVENT data. The display can be built only when data collection for all these domains is active.

For more detailed seeks data for one specific disk either:

- Place the cursor on the Addr field of the disk you are interested in and press ENTER, or
- Enter the subcommand 'SEEKDET devno' on the command line

The 'Seeks Analysis Details' screen will then be shown with detailed data for the selected device (see “SEEKDET” on page 284 for more information on that performance display).

### Field Descriptions:

The following field descriptions apply to both the SEEKLOC and SEEKLLOG devno displays, with the exception of the Addr field which will be replaced by the field Interval End Time on the SEEKLLOG display.

**Device:**

#### Addr

Real device number. The first line contains the prefix >> Mean << instead of a device number. It provides averages for all disks for which seeks data were collected.

#### Device Type

Device type and model of the disk

**Volume Serial**

The volume serial of the disk. Dots can be shown initially if no I/O sample data have been collected yet for the disk.

**Size in Cyls**

The number of cylinders available on the disk.

**Disc Time msec**

The average disconnect time during the interval. Dots will be shown for disks where the disconnect time cannot be calculated, either because we have not yet received two samples which allow calculating a valid delta value, or because there has been no I/O activity during the sample period.

**Cached**

Indicates whether the disk has profited from the cache during the last sample interval, i.e. whether any read or write hits were recorded. The string Yes will be inserted if hits were found, and No will be inserted otherwise.

*Seeks:*

**Pct Read**

Percent of seeks to the disk which were for channel programs containing only READ operations.

**Pct Non0**

The percentage of seeks that caused a displacement of the access arm of at least 1 cylinder.

**Mean Non-0 Dist**

The average displacement of the access arm, taking into account only the seeks that caused a nonzero displacement.

**Total Seeks**

The total number of seeks activity samples recorded for the disk.

*Percent of Seek Targets within Percent Range of Cylinders:*

The seek targets are grouped into 10 pools, each with a size of approximately 10% of the total number of cylinders. The frequency with which seeks were made to each area is expressed as the percentage of seeks that were targeted to the area, where the seek target headings are as follows:

**0 <= <=10**

For seeks targeted to the lowest range of the disk's cylinders.

**10 < <=20**

For seeks targeted to the next-highest range of the disk's cylinders.

**20 < <=30**

For seeks targeted to the next-highest range of the disk's cylinders.

**30 < <=40**

For seeks targeted to the next-highest range of the disk's cylinders.

**40 < <=50**

For seeks targeted to the next-highest range of the disk's cylinders.

**50 < <=60**

For seeks targeted to the next-highest range of the disk's cylinders.

**60 < <=70**

For seeks targeted to the next-highest range of the disk's cylinders.

**70 < <=80**

For seeks targeted to the next-highest range of the disk's cylinders.

**80 < <=90**

For seeks targeted to the next-highest range of the disk's cylinders.

**90 <**

For seeks targeted to the highest range of the disk's cylinders.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown. Note that overall period averages will also be shown for the 'current' setting.

## FCX157, DASD Seeks Distances Screen – SEEKDIST

**Performance Toolkit Subcommand:** “SEEKDIST” on page 285.

The DASD Seeks Distances Screen (FCX157) shows DASD seeks statistics on the distribution of seeks distances for each monitored disk:

FCX157		CPU nnnn		SER nnnnn		Interval HH:MM:SS - HH:MM:SS		Perf. Monitor	
<----- Device ----->		<----- Seeks ----->		<----- Percent with ----->					
Device	Volume	Size in Disc	Cach	Pct	Pct	Non-0	Total	=	<=
Addr Type Serial		Cyls msec	ed	Read	Non0	Dist	Seeks	0	5
>> Mean <<									
0600	3380-K	PROFSA	2655	21.2	Yes	13	31	198	4222
0601	3380-K	VMPRA1	2655	19.3	Yes	25	38	171	471
0602	3380-K	PROFSB	2655	21.6	Yes	13	50	150	5539
0603	3380-K	VMPRA3	2655	21.6	Yes	24	48	279	10898
0604	3380-K	VMPRA4	2655	13.9	Yes	57	48	178	1605
0605	3380-K	VMPRA5	2655	21.4	Yes	19	50	217	5017
0606	3380-K	VMPRA6	2655	22.8	Yes	18	49	257	5816
0607	3380-K	VMPRA7	2655	20.5	Yes	17	44	217	8056
0608	3380-K	ESA111	2655	10.8	No	5	4	814	80
060A	3380-K	VMPRAA	2655	21.6	Yes	17	43	213	3072
060B	3380-K	VMPRAB	2655	10.7	Yes	54	53	117	6228
060C	3380-K	VMPG1N	2655	13.3	No	56	98	94	4136
060D	3380-K	VMPG10	2655	13.4	No	52	99	70	4187
060E	3380-K	VMPRAE	2655	21.5	Yes	19	41	206	6809
Select a device for I/O device details									
Command ==>									
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return									

Figure 57. Layout of DASD Seeks Distances Screen (FCX157)

The sample shows a screen with disks sorted in ascending device number sequence (indicated by the underscore characters above the 'Addr' header), the initial default sequence. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295), or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

Note that the information extends beyond the 80 columns of an ordinary screen: use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window so as to view the remaining data.

The information shown is based on MONITOR and I/O domain SAMPLE data and on SEEKS domain EVENT data. The display can be built only when data collection for all these domains is active.

For more detailed seeks data for one specific disk either:

- Place the cursor on the Addr field of the disk you are interested in and press ENTER, or
- Enter the subcommand 'SEEKDET devno' on the command line

The 'DASD Seeks Details' screen will then be shown with detailed data for the selected device (see “SEEKDET” on page 284 for more information on that performance display).

### Field Descriptions:

The following field descriptions apply to both the SEEKDIST and SEEKDLOG devno displays, with the exception of the Addr field which will be replaced by the field Interval End Time on the SEEKDLOG display.

Device:

#### Addr

Real device number.

The first line contains the prefix >> Mean << instead of a device number. It provides averages for all disks for which seeks data were collected.

**Device Type**

Device type and model of the disk

**Volume Serial**

The volume serial of the disk. Dots can be shown initially if no I/O sample data have been collected yet for the disk.

**Size in Cyls**

The number of cylinders available on the disk.

**Disc Time msec**

The average disconnect time during the interval. Dots will be shown for disks where the disconnect time cannot be calculated, either because we have not yet received two samples which allow calculating a valid delta value, or because there has been no I/O activity during the sample period.

**Cached**

Indicates whether the disk has profited from the cache during the last sample interval, i.e. whether any read or write hits were recorded. The string Yes will be inserted if hits were found, and No will be inserted otherwise.

*Seeks:*

**Pct Read**

Percent of seeks to the disk which were for channel programs containing only READ operations.

**Pct Non0**

Percent of seeks which caused a displacement of the access arm of at least 1 cylinder.

**Mean Non-0 Dist**

Average displacement of the access arm, taking into account only the seeks which caused a nonzero displacement.

**Total Seeks**

The total number of seeks activity samples recorded for the disk.

*Percent of Seek Distances within the Ranges:*

The seek distances are grouped into pools of geometrically increasing seek distances. The seek distance distribution is shown by displaying the percentage of seeks which falls into each seek distance pool, where the pool ranges are

**= 0**

For seeks with a seek distance of zero

**<= 5**

For seeks with a seek distance in the range of 1 up to 5 cylinders

**<= 10**

For seeks with a seek distance in the range of 6 up to 10 cylinders

...

**>800**

For seeks with a seek distance exceeding 800 cylinders.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown. Note that overall period averages will also be shown for the 'current' setting.

## FCX158, DASD Seeks Details Screen – SEEKDET

**Performance Toolkit Subcommand:** “SEEKDET” on page 284.

The DASD Seeks Details Screen (FCX158) provides detailed information on seeks activity for one specific disk:

```

FCX158          CPU nnnn    SER nnnnn   Interval HH:MM:SS - HH:MM:SS    Perf. Monitor

Seeks Analysis Details

Device number:   0615          Total nr. of SEEks :    2497
Device type    : 3380-K        Percent Read-only  :      6
Volume serial   VMPRA8        Percent non-0 seeks:     44
Cylinders      : 2655          Ave. seeks distance:    278
Cached         : Yes           Disc. time (msec)   :   20.4


Seek Range  Owner      Vaddr  Seek %Rd %Nz Dist  0  % of Total Seeks  100
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 28-   37  RVO         0191    71  21  86 1140 | > | . |
| 82-   85  MSMU       0191    80   0  63  44  | > | . |
| 133-  265 | . | . |
| 266-  398 | . | . |
| 399-  530 | . | . |
| 531-  663 | . | . |
| 664-  796 | . | . |
| 872-  872  GHA       0191     1 100 100  787 | . | . |
| 932-  937  MWR       0191   248   0  51  332 | >>> | . |
| 947-  953  EGG       0191     3   0  67   39 | . |
| 1022- 1026 PSTE      0191    68  65  44  257 | R | . |
| 1116- 1120 RCL       0191   149  15  64  553 | >> | . |
| 1195- 1327 | . | . |
| 1328- 1459 | . | . |
| 1460- 1592 | . | . |
| 1654- 1679 EKO       0191   141   9  62  377 | >> | . |
| 1769- 1780 JA        0191  1335   2  33  112 | >>>>>>>>>>>>>>>> | . |
| 1866- 1885 EGRO      0191    89  20  46   79 | > | . |
| 1991- 2123 | . | . |
| 2139- 2186 MBIR      0191   312   3  56  268 | >>> | . |
| 2256- 2388 | . | . |
| 2389- 2521 | . | . |
| 2522- 2654 | . | . |
+-----+-----+-----+-----+-----+-----+-----+-----+

Command ==> _
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return

```

Figure 58. Layout of DASD Seeks Details Screen (FCX158)

The information shown is based on MONITOR and I/O domain SAMPLE data and on SEEKS domain EVENT data. The display can be built only when data collection for all these domains is active.

Two blocks of data will be shown:

- A first fixed block with overall information for the disk, and
- A second block with variable layout, showing detailed information on the seeks activity to specific cylinder ranges and minidisks.

Seeks activity to different areas of the disk pack will also be displayed graphically as a bar chart, as part of the second block.

### Field Descriptions:

*Seeks Analysis Details:*

**Device number**

Real device number

### Device type

Device type and model of the disk



**Volume serial**

The volume serial of the disk. Dots can be shown initially if no I/O sample data have been collected yet for the disk.

**Cylinders**

The number of cylinders available on the disk

**Cached**

Indicates whether the disk has profited from the cache during the last sample interval, i.e. whether any read or write hits were recorded. The string Yes will be inserted if hits were found, and No will be inserted otherwise.

**Total nr. of seeks**

The total number of seeks activity samples recorded for the disk

**Percent Read-only**

Percent of seeks to the disk which were for channel programs containing only READ operations

**Percent non-0 seeks**

Percent of seeks which caused a displacement of the access arm of at least 1 cylinder

**Ave. seeks distance**

Average displacement of the access arm, taking into account only the seeks which caused a nonzero displacement

**Disc. time (msec)**

The average disconnect time during the interval. Dots will be shown for disks where the disconnect time cannot be calculated, either because we have not yet received two samples which allow calculating a valid delta value, or because there has been no I/O activity during the sample period.

*Detailed Data for Cylinder Ranges and Minidisks:***Seek Range**

The cylinder range to which the remaining information on the line applies.

The disk is split into 20 ranges of equal size. Where seeks are targeted to a minidisk, the lowest and highest target cylinder numbers are used to describe the minidisk extent. The ranges used for the display are either the original ranges of 1/20 of the total disk size, or (parts of) a minidisk, whichever is smaller, i.e. you may see

- More than one small minidisk shown as a separate entity although they would all have fitted into one of the 1/20 slots
- More than one line for a single minidisk if it extends beyond the boundary of one of the original 1/20 slots.

The net result should be a display which shows sufficient detail for locating very busy minidisks, or very busy areas on a single minidisk or full-pack disk.

**Owner**

The owner of the minidisk. The string System will be inserted if the cylinder range is not part of a user minidisk but of a CP area (e.g. PAGE or SPOOL area).

**Vadr**

The virtual device number of the minidisk. The field is left empty if the line does not describe a minidisk.

**Seeks**

The number of seeks samples collected which were targeted to this cylinder range

**%Rd**

The percentage of seeks to the cylinder range which were for read-only operations

**%Nz**

The percentage of seeks to the cylinder range which entailed a movement of the access arm of at least 1 cylinder

**Dist**

The average displacement of the access arm for all nonzero seeks which were targeted to this area

**% of Total Seeks**

The percentage of seeks which were targeted to the cylinder range, shown in the form of a bar chart. The characters used to form the bars are

**R**

For seeks with read-only operations

**>**

For seeks containing also write operations

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None. Average data are always displayed.

## FCX160, History Data Selection Menu Screen – HISTDATA

**Performance Toolkit Subcommand:** “HISTDATA” on page 195.

The History Data Selection Menu Screen (FCX160):

```
FCX160      File Selection: All Files with Default Filetypes      History Data

Move cursor to the history data file you are interested in and
- enter 'G' to select the file for creating graphics
- just press ENTER for viewing the data in REDISP format

S Filename  Filetype  FM  Created      S Filename  Filetype  FM  Created
. CHVM4     RMONLOG   V   2003/07/13   . CHVM5     RMONLOG   V   2003/07/13
. ACUM      HISTSUM    V   2003/07/12   . ACUM      PERFHIST  V   2003/07/12
. CHVM4     RMONLOG1   V   2003/07/12   . CHVM5     RMONLOG1   V   2003/07/12
. CHVM4     RMONLOG2   V   2003/07/11   . CHVM5     RMONLOG2   V   2003/07/11
. 071103    PERFLAG1   V   2003/07/11   . 19970711  HISTLOG1   V   2003/07/11
. 071003    PERFLAG2   V   2003/07/10   . 19970710  HISTLOG2   V   2003/07/10
. 070903    PERFLAG3   V   2003/07/09   . 19970709  HISTLOG3   V   2003/07/09
. 070803    PERFLAG3   V   2003/07/08   . 19970708  HISTLOG3   V   2003/07/08
. 070703    PERFLAG3   V   2003/07/07   . 19970707  HISTLOG3   V   2003/07/07
. 070403    PERFLAG3   V   2003/07/04   . 19970704  HISTLOG3   V   2003/07/04

Command ==> _
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 59. Layout of History Data Selection Menu Screen (FCX160)

See the "Viewing Performance History Data" section in the [z/VM: Performance Toolkit Guide](#) for more information on this screen.

## FCX161, LPAR Channel Load Screen – LCHANNEL

**Performance Toolkit Subcommand:** “LCHANNEL” on page 211.

In the LPAR Channel Load Screen (FCX161), fields which indicate that the channel has been busy are highlighted.

FCX161		CPU nnnn		SER nnnnn		Interval HH:MM:SS - HH:MM:SS						Perf. Monitor				
CHPID	Chan-Group	<%Busy>		<----- Channel %Busy Distribution		03:37:21-10:21:21 ----->										
(Hex)	Descr	Qual	Shrd	Cur	Ave	0-10	11-20	21-30	31-40	41-50	51-60	61-70	71-80	81-90	91-100	
B4	HPF	00	Yes	27	6	80	0	14	6	0	0	0	0	0	0	
B5	HPF	00	Yes	27	6	80	0	14	6	0	0	0	0	0	0	
F6	HPF	00	Yes	14	3	80	20	0	0	0	0	0	0	0	0	
F7	HPF	00	Yes	14	3	80	20	0	0	0	0	0	0	0	0	
02	ESCON	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
03	ESCON	00	No	0	0	100	0	0	0	0	0	0	0	0	0	
04	ESCON	00	No	0	0	100	0	0	0	0	0	0	0	0	0	
05	ESCON	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
06	HPF	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
0F	HPF	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
12	OSC	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
17	OSD	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
18	HPF	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
19	HPF	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
1A	OSD	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
1B	OSD	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
1C	HPF	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
1E	OSD	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
1F	OSD	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
20	FCP	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
20	FCP	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
21	FCP	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
22	FCP	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
23	FCP	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
24	HPF	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
25	HPF	00	Yes	0	0	100	0	0	0	0	0	0	0	0	0	
...																

See also CHANNEL and FCHANNEL displays

Command ==>

F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 60. Layout of LPAR Channel Load Screen (FCX161)

LPAR channel load analysis is based on CP monitor SYSTEM domain data from the Channel-Path Measurement Facility where available, otherwise data from the Extended Channel-Path Measurement Facility is used. For systems running in a logical partition, the values shown pertain to the channel activity generated by *this* partition only, not the overall channel activity. The screen shows a list of all channels for which valid information was found.

Channels will be ordered in %Busy sequence, either for the *current*, i.e. the last interval's %Busy values if 'current' data are being displayed, or for *average* %Busy values when 'average' performance data have been selected. Channels with equal %Busy figures will be sorted in ascending CHPID sequence.

### Field Descriptions for Channel Load:

#### CHPID (Hex)

The hexadecimal 'channel path identifier'.

#### Chan-Group Descr

Description of the channel model group, as found in the Store-Channel-Path description.

#### Chan-Group Qual

The hexadecimal channel model group qualifier. Note that a valid qualifier is available only on systems with the extended channel measurement facility.

#### Shrd

'Shared' indicator for the channel. It can be either:

**YES**

For channels which are shared with other logical partitions

**NO**

For channels which are dedicated to this logical partition

**%Busy**

The percentage of 'busy' conditions found, where

**Cur**

Shows the 'current' figures, calculated for the last interval

**Ave**

Shows the 'average' load, calculated for the whole period since the last RESET, or since Performance Toolkit was last invoked.

For FICON channels the 'busy' value is calculated as the percentage of actual work unit utilization compared to the maximum possible value. Channel 'busy' percentages cannot be calculated for Hipersocket channels; the string n/a is inserted instead.

Under the heading:

```
Channel %Busy Distribution hh:mm:ss-hh:mm:ss
```

You will find the channel load distribution for each of the active channels. It shows the relative frequency with which the channel busy percentage of separate measurement intervals was found in either of the 10% ranges 0-10, 11-20, ..., 91-100.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Interim averages are displayed for the INTERIM setting, and overall averages are shown for both the AVERAGE and CURRENT settings since average data is shown together with the current load. Sorting of the channels is based on the average %Busy percentage instead of the last interval's %Busy values.

## FCX162, User Resource Usage Log – USERLOG userid

**Performance Toolkit Subcommand:** “USERLOG” on page 347.

The User Resource Usage Log report (FCX162) shows a "by time" log of the selected user's resource consumption. Note that the specified user must have been selected previously for benchmarking.

FCX162		CPU nnnn		SER nnnnn		Interval HH:MM:SS - HH:MM:SS		Perf. Monitor											
Resource Usage Log for User RSTL4B																			
Interval End Time	<----- CPU Load ----->				<----- Virtual IO/s ----->				<UsrTime>		<--Spool-->		MDC		VConf Rebal Scat.				
	%CPU	TCPU	VCPU	Ratio	T/V	Total	DASD	Avoid	98	UR	Pg/s	User Status	<Minutes> Logd Act.	Total Pages		Rate SpG/s	Insert MDC/s	Share	
>>Mean>>	.03	.015	.005	2.81		.0	.0	.0	.0	.0	.0	-----,-----	28	27	.0	.0	.0	100	----
11:58:55	.01	.008	.004	2.00		.0	.0	.0	.0	.0	.0	EME,CL0,DISP	1	1	.0	.0	.0	100	.000
11:59:55	.01	.007	.004	1.75		.0	.0	.0	.0	.0	.0	EME,CL0,DISP	1	1	.0	.0	.0	100	.000
12:00:55	.01	.008	.004	2.00		.0	.0	.0	.0	.0	.0	EME,CL0,DORM	1	1	.0	.0	.0	100	.000
12:01:55	.01	.008	.004	2.00		.0	.0	.0	.0	.0	.0	EME,CL0,DISP	1	1	.0	.0	.0	100	.000
12:02:55	.01	.008	.004	2.00		.0	.0	.0	.0	.0	.0	EME,CL0,DISP	1	1	.0	.0	.0	100	.000
12:03:55	.01	.008	.004	2.00		.0	.0	.0	.0	.0	.0	EME,CL0,DORM	1	1	.0	.0	.0	100	.000
12:04:55	.02	.009	.004	2.25		.0	.0	.0	.0	.0	.0	EME,CL0,DISP	1	1	.0	.0	.0	100	.000
12:05:55	.01	.008	.004	2.00		.0	.0	.0	.0	.0	.0	EME,CL0,DORM	1	1	.0	.0	.0	100	.000
12:06:55	.01	.008	.004	2.00		.0	.0	.0	.0	.0	.0	EME,CL0,DORM	1	1	.0	.0	.0	100	.000
12:07:55	.01	.008	.004	2.00		.0	.0	.0	.0	.0	.0	EME,CL0,DORM	1	1	.0	.0	.0	100	.000
12:08:55	.01	.008	.004	2.00		.0	.0	.0	.0	.0	.0	EME,CL0,DISP	1	1	.0	.0	.0	100	.000
12:09:55	.01	.008	.004	2.00		.0	.0	.0	.0	.0	.0	EME,CL0,DISP	1	1	.0	.0	.0	100	.000
12:10:55	.01	.008	.004	2.00		.0	.0	.0	.0	.0	.0	EME,CL0,DORM	1	1	.0	.0	.0	100	.000
12:11:55	.01	.008	.004	2.00		.0	.0	.0	.0	.0	.0	EME,CL0,DISP	1	1	.0	.0	.0	100	.000
12:12:55	.01	.008	.004	2.00		.0	.0	.0	.0	.0	.0	EME,CL0,DISP	1	1	.0	.0	.0	100	.000
12:13:55	.01	.008	.004	2.00		.0	.0	.0	.0	.0	.0	EME,CL0,DISP	1	1	.0	.0	.0	100	.000
12:14:55	.01	.008	.004	2.00		.0	.0	.0	.0	.0	.0	EME,CL0,DISP	1	1	.0	.0	.0	100	.000
12:15:55	.01	.008	.004	2.00		.0	.0	.0	.0	.0	.0	EME,CL0,DISP	1	1	.0	.0	.0	100	.000
12:16:55	.01	.008	.004	2.00		.1	.0	.0	.0	.0	.0	EME,CL0,DISP	1	1	.0	.0	.0	100	.000
12:17:55	User logged off																		
12:18:55	User logged off																		
12:19:55	User logged off																		
Command ==>																			
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return																			

Figure 61. Layout of User Resource Usage Log report (FCX162)

The information shown is based on USER domain SAMPLE data.

The first entry is reserved for averages or totals (for CPU seconds) since the last RESET subcommand or since data collection was activated. This entry is designated by the >>Mean>> string in the Interval End Time column.

If a user logs off after data collection has begun, the following string is displayed (as shown in the sample report):

User logged off

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the Interval End Time field, which replaces the Userid field (and the Nr of users field, which is not included in this report), all data columns in this report are identical to the data columns in the User Resource Usage report. For more information, see “FCX112, User Resource Usage – USER” on page 431.

### Effect of AVERAGE and INTERIM subcommands:

None.

## FCX163, User Page Data Log Screen – UPAGELOG

**Performance Toolkit Subcommand:** “UPAGELOG” on page 336.

The User Page Data Log (FCX163) shows a 'by time' log of the paging activity and storage utilization for the selected user (which must have been previously selected for benchmarking):

```

FCX163      CPU nnnn  SER nnnnn Interval HH:MM:SS - HH:MM:SS      Perf. Monitor

Page Data Log for User CHCAL

Data
Interval Spaces <----- Paging Activity/s -----> <----- Number of Pages ----->
Interval Spaces <Page Rate> Page <--Page Migration--> <-Resident-> <--Locked--> Stor
End Time Owned Reads Write Steals >2GB> X>MS MS>X X>DS WSS Resrvd R<2GB R>2GB L<2GB L>2GB XSTOR DASD Size
>>Mean>> .0 .0 .0 .0 .0 .0 .0 .0 .0 786432 524288 5 786423 0 0 0 451167 3072M
06:09:19 0 .0 .0 .0 .0 .0 .0 .0 786432 524288 5 786423 0 0 0 451167 3072M
06:14:19 0 .0 .0 .0 .0 .0 .0 .0 786432 524288 5 786423 0 0 0 451167 3072M
06:19:19 0 .0 .0 .0 .0 .0 .0 .0 786432 524288 5 786423 0 0 0 451167 3072M
06:24:19 0 .0 .0 .0 .0 .0 .0 .0 786432 524288 5 786423 0 0 0 451167 3072M

Command ==>
F1=Help F4=Top F5=Bot F12=Return

```

Figure 62. Layout of User Page Data Log Screen (FCX163)

The information shown is based on USER domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Descriptions:

## Interval End Time

The end time of the monitor interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'Userid' field (and the 'Nr of users' field which is omitted from this screen), all data columns are identical to the ones of the 'User Page Data' (UPAGE) display. Please see “FCX113, User Page Data Screen – UPAGE” on page 435 for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX164, User Wait States Log Screen – USTATLOG

**Performance Toolkit Subcommand:** “USTATLOG” on page 352.

The User Wait States Log Screen (FCX164) shows a 'by time' log of the user's wait state statistics. Note that the specified user must have been previously selected for benchmarking.

FCX164	CPU nnnn	SER nnnnn	Interval	HH:MM:SS	- HH:MM:SS	Perf. Monitor																				
Wait State Data Log for User RSTL3																										
Interval	<-SVM and->																		<-Samples showing->					<-vCPU-->		
End Time	%ACT	%RUN	%CPU	%LDG	%PGW	%IOW	%SIM	%TIW	%CFW	%TI	%EL	%DM	%IOA	%PGA	%LIM	%OTH	Q0	Q1	Q2	Q3	E0-3	Type	Cnt			
>>Mean>>	10	1	0	0	0	0	0	45	0	0	0	54	0	0	0	0	46	0	0	0	0	CP	1			
17:43:38	3	0	0	0	0	0	0	0	0	0	0	100	0	0	0	0	0	0	0	0	0	CP	1			
17:44:38	7	0	0	0	0	0	0	0	0	0	0	100	0	0	0	0	0	0	0	0	0	CP	1			
17:45:38	7	0	0	0	0	0	0	0	0	0	0	100	0	0	0	0	0	0	0	0	0	CP	1			
17:46:38	7	50	0	0	0	0	0	0	0	0	0	50	0	0	0	0	50	0	0	0	0	CP	1			
17:47:38	0	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	CP	1			
17:48:38	7	0	0	0	0	0	0	0	0	0	0	100	0	0	0	0	0	0	0	0	0	CP	1			
17:49:38	3	0	0	0	0	0	0	0	0	0	0	100	0	0	0	0	0	0	0	0	0	CP	1			
17:50:38	7	0	0	0	0	0	0	0	0	0	0	100	0	0	0	0	0	0	0	0	0	CP	1			
17:51:38	7	0	0	0	0	0	0	0	0	0	0	100	0	0	0	0	0	0	0	0	0	CP	1			
17:52:38	17	20	0	0	0	0	0	40	0	0	0	40	0	0	0	0	60	0	0	0	0	CP	1			
17:53:38	17	0	0	0	0	0	0	60	0	0	0	40	0	0	0	0	60	0	0	0	0	CP	1			
17:54:38	20	0	0	0	0	0	0	67	0	0	0	33	0	0	0	0	67	0	0	0	0	CP	1			
17:55:38	17	0	0	0	0	0	0	80	0	0	0	20	0	0	0	0	80	0	0	0	0	CP	1			
17:56:38	3	0	0	0	0	0	0	100	0	0	0	0	0	0	0	0	100	0	0	0	0	CP	1			
17:57:38	7	0	0	0	0	0	0	0	0	0	0	100	0	0	0	0	0	0	0	0	0	CP	1			
17:58:38	17	0	0	0	0	0	0	60	0	0	0	40	0	0	0	0	60	0	0	0	0	CP	1			
17:59:38	7	0	0	0	0	0	0	0	0	0	0	100	0	0	0	0	0	0	0	0	0	CP	1			
18:00:38	13	0	0	0	0	0	0	50	0	0	0	50	0	0	0	0	50	0	0	0	0	CP	1			
18:01:38	17	0	0	0	0	0	0	60	0	0	0	40	0	0	0	0	60	0	0	0	0	CP	1			
18:02:38	10	0	0	0	0	0	0	33	0	0	0	67	0	0	0	0	33	0	0	0	0	CP	1			
18:03:38	13	0	0	0	0	0	0	50	0	0	0	50	0	0	0	0	50	0	0	0	0	CP	1			
18:04:38	10	0	0	0	0	0	0	33	0	0	0	67	0	0	0	0	33	0	0	0	0	CP	1			
18:05:38	13	0	0	0	0	0	0	50	0	0	0	50	0	0	0	0	50	0	0	0	0	CP	1			
18:06:38	20	0	0	0	0	0	0	67	0	0	0	33	0	0	0	0	67	0	0	0	0	CP	1			
18:07:38	3	0	0	0	0	0	0	100	0	0	0	0	0	0	0	0	100	0	0	0	0	CP	1			
18:08:38	10	0	0	0	0	0	0	33	0	0	0	67	0	0	0	0	33	0	0	0	0	CP	1			
18:09:38	10	0	0	0	0	0	0	33	0	0	0	67	0	0	0	0	33	0	0	0	0	CP	1			
18:10:38	7	0	0	0	0	0	0	50	0	0	0	50	0	0	0	0	50	0	0	0	0	CP	1			
18:11:38	13	0	0	0	0	0	0	75	0	0	0	25	0	0	0	0	75	0	0	0	0	CP	1			
18:12:38	17	0	20	0	0	0	0	60	0	0	0	20	0	0	0	0	80	0	0	0	0	CP	1			
Command ==>>																										
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return																										

Figure 63. Layout of User Wait States Log Screen (FCX164)

The information shown is based on USER domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'Userid' field (and the 'Nr of users' field which is omitted from this screen), all data columns are identical to the ones of the 'User Wait States' (USTAT) display. Please see “FCX114, User Wait States Screen – USTAT/USTATG” on page 438 for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.



## FCX165, User Response Time Log Screen – URESPLOG

**Performance Toolkit Subcommand:** “URESPLOG” on page 345.

The User Response Time Log Screen (FCX165) shows a 'by time' log of the user's response time per measured transaction. Note that the specified user must have been previously selected for benchmarking.

FCX165	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS					Perf. Monitor		
Response Time Data Log for User CHCAL										
Interval	Think	<- Response Time ->			<- Transact/Hour ->			<-- Total Trans -->		
End Time	Time	Triv.	N-Triv	Mean	Trv/h	Ntrv/h	Totl/h	Trv-T	Ntrv-T	Totl-T
>>Mean>>	2.5	...	.405	.405	0	1501	1501	0	422	422
12:36:47	2.7	...	.459	.459	0	1264	1264	0	21	21
12:37:48	2.3	...	.309	.309	0	1745	1745	0	30	30
12:38:48	2.8	...	.435	.435	0	1264	1264	0	21	21
12:39:49	2.5	...	.540	.540	0	1479	1479	0	25	25
12:40:35	4.2	...	.187	.187	0	936	936	0	12	12
12:41:47	2.5	...	.500	.500	0	1442	1442	0	29	29
12:42:47	2.1	...	.338	.338	0	1867	1867	0	31	31
12:43:48	3.0	...	.299	.299	0	1302	1302	0	22	22
12:44:39	3.0	...	.178	.178	0	1331	1331	0	19	19
12:45:45	3.1	...	.390	.390	0	1089	1089	0	20	20
12:46:47	2.3	...	.322	.322	0	1803	1803	0	31	31
12:47:49	2.8	...	.196	.196	0	1396	1396	0	24	24
12:48:48	2.8	...	.315	.315	0	1348	1348	0	22	22
12:49:49	2.1	...	.457	.457	0	1894	1894	0	32	32
Command ===> _										
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return					

Figure 64. Layout of User Response Time Log Screen (FCX165)

The information shown is based on USER domain EVENT data.

The first entry is reserved for averages and totals (for total transaction counts) since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The time of the last transaction within the monitor interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'Userid' field (and the 'Nr of users' field which is omitted from this screen), all data columns are identical to the ones of the 'User Response Time' display. Please see “FCX127, User Response Time Screen – URESP” on page 475 for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX166, User Resources per Transaction Log Screen – UTRANLOG

**Performance Toolkit Subcommand:** “UTRANLOG” on page 356.

The User Resources per Transaction Log Screen (FCX166) shows a 'by time' log of the user's average resource consumption per measured transaction. Note that the specified user must have been previously selected for benchmarking.

FCX166	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS				Perf. Monitor			
Transaction Data Log for User CHCAL										
	<-Transact->		<----- Resource Utilization per Transaction ----->							
Interval	Total	Rate	CPU	Disk	Page	Xpage	<-Spooling->		<-IUCV/VMCF->	
End Time	Count	/hour	msec	I/O	Flts	Reads	SP_RD	SP_WR	IU_BY	IU_TO
>>Mean>>	443	1271	26.2	7.7	.0	.0	.0	.4	2.9	2.9
12:37:48	30	1745	23.9	6.9	.0	.0	.0	.4	2.7	2.7
12:38:48	21	1264	27.8	9.7	.0	.0	.0	.5	3.0	3.0
12:39:49	25	1479	29.8	6.5	.0	.0	.0	.5	3.6	3.6
12:40:35	12	936	21.9	5.2	.0	.0	.0	.3	2.5	2.5
12:41:47	29	1442	29.2	8.2	.0	.0	.0	.5	3.3	3.3
12:42:47	31	1867	23.6	6.4	.0	.0	.0	.4	2.7	2.7
12:43:48	22	1302	23.5	6.0	.0	.0	.0	.3	2.8	2.8
12:44:39	19	1331	22.5	4.6	.0	.0	.0	.4	2.6	2.6
12:45:45	20	1089	31.3	5.7	.0	.0	.0	.7	3.9	3.9
12:46:47	31	1803	24.3	7.3	.0	.0	.0	.4	2.7	2.7
12:47:49	24	1396	22.7	4.9	.0	.0	.0	.4	2.7	2.7
12:48:48	22	1348	20.9	7.7	.0	.0	.0	.3	2.1	2.1
12:49:49	32	1894	27.2	11.0	.0	.0	.1	.5	2.6	2.6
Command ==>										
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return					

Figure 65. Layout of User Resources per Transaction Log Screen (FCX166)

The information shown is based on USER domain EVENT data.

The first entry is reserved for averages and totals (for total transaction count) since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The time of the last transaction within the monitor interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'Userid' field, all data columns are identical to the ones of the 'User Resources per Transaction' display. Please see “UTRANS” on page 357 for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX167, User Communication Log Screen – UCOMMLOG

**Performance Toolkit Subcommand:** [“UCOMMLOG”](#) on page 332.

The User Communication Log Screen (FCX167), which shows a 'by time' log of the selected user's IUCV and VMCF activity, will be shown if you specify a *userid* on the UCOMMLOG subcommand, where *userid* must be the user identification of a virtual machine which has previously been selected for benchmarking:

**Note:** If you specify UCOMMLOG with no *userid*, you will instead see [“FCX182, General User Communication Log Screen – UCOMMLOG”](#) on page 585.

FCX167	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS				Perf. Monitor			
Communications Data Log for User VTAM										
Interval	Msg	S								
End Time	Target	V Total	<- IUCV Msg/s -->			<- VMCF Msg/s -->			<-- IUCV in Q	
>>Mean>>	*SIGNAL	M Msg/s	ISend	IRecv	IFail	VSend	VRecv	VFail	SendQ	RecvQ R
09:45:58	*SIGNAL	Y .985	.476	.509	.000	.000	.000	.000	0	0
09:46:58	*CCS	Y 5.517	3.667	1.850	.000	.000	.000	.000	0	0
09:47:59	*CCS	Y 1.083	.483	.600	.000	.000	.000	.000	0	0
09:48:58	*CCS	Y .987	.444	.543	.000	.000	.000	.000	0	0
09:49:58	*SIGNAL	Y .203	.084	.118	.000	.000	.000	.000	0	0
09:50:58	*SIGNAL	Y .167	.067	.100	.000	.000	.000	.000	0	0
09:51:58	*SIGNAL	Y .167	.083	.083	.000	.000	.000	.000	0	0
09:52:58	*CCS	Y .317	.150	.167	.000	.000	.000	.000	0	0
09:53:58	*CCS	Y .883	.383	.500	.000	.000	.000	.000	0	0
09:54:58	*CCS	Y 1.067	.600	.467	.000	.000	.000	.000	0	0
09:55:58	*CCS	Y 1.667	.950	.717	.000	.000	.000	.000	0	0
09:56:58	*SIGNAL	Y 1.783	1.050	.733	.000	.000	.000	.000	0	0
09:57:58	*SIGNAL	Y .150	.067	.083	.000	.000	.000	.000	0	0
09:58:58	*CCS	Y .650	.350	.300	.000	.000	.000	.000	0	0
09:59:58	*CCS	Y .833	.400	.433	.000	.000	.000	.000	0	0
Command ==>										
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return					

Figure 66. Layout of User Communication Log Screen (FCX167)

The information shown is based on USER domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'Userid' field (and the 'Nr of users' field which is omitted from this screen), all data columns are identical to the ones of the 'User Communication' (UCOMM) display. Please see [“FCX129, User Resources per Transaction Screen – UTRANS”](#) on page 478 for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX168, General I/O Device Data Log Screen – DEVLOG

**Performance Toolkit Subcommand:** “DEVLOG” on page 170.

The General I/O Device Data Log Screen (FCX168) shows a "by time" log of general I/O performance data for the selected device. The device data log is available only for devices you have benchmarked. For each benchmarked device, the interval and line for that device from the FCX108 (INTERIM DEVICE) screen appears on this screen. This allows you to look at a device over time and see how it is performing.

The FCX168 screen looks like this for a HyperPAV base device or a PAV base device:

FCX168	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS		Perf. Monitor											
General I/O Data Log for Device BE00																
Interval	<- Dev. Descr.->	Mdisk	Pa-	<-Rate/s->	<-----	Time (msec)	----->	Req.	<Percent>	SEEK	Recov	<-Throttle->				
End Time	Type	Links	ths	I/O	Avoid	Pend	Disc	Conn	Serv	Resp	CUWt	Qued	Busy	READ	Cyls	SSCH
>>Mean>>	3390-9 ATP000 CP	0	4	15.1	.0	.264	13.0	.564	13.9	13.9	.087	.00	21	95	0	0
14:34:16	3390-9 ATP000 CP	0	4	17.6	.0	.270	13.5	.535	14.3	14.3	.048	.00	25	95	0	0
14:34:46	3390-9 ATP000 CP	0	4	10.0	.0	.327	17.0	.720	18.0	18.0	.185	.00	18	94	0	0
14:35:16	3390-9 ATP000 CP	0	4	13.8	.0	.283	17.1	.457	17.9	17.9	.010	.00	25	95	0	0
14:35:46	3390-9 ATP000 CP	0	4	10.2	.0	.392	29.3	.811	30.5	30.5	.339	.00	31	93	0	0
14:36:16	3390-9 ATP000 CP	0	4	15.6	.0	.258	12.2	.488	12.9	12.9	.055	.00	20	95	0	0
14:36:46	3390-9 ATP000 CP	0	4	10.7	.0	.312	16.2	.846	17.4	17.4	.222	.00	19	92	0	0
14:37:16	3390-9 ATP000 CP	0	4	17.0	.0	.280	14.7	.540	15.5	15.6	.138	.00	26	97	0	0
14:37:46	3390-9 ATP000 CP	0	4	13.1	.0	.305	14.6	.620	15.5	15.5	.043	.00	20	95	0	0
14:38:16	3390-9 ATP000 CP	0	4	19.7	.0	.247	11.9	.588	12.7	12.7	.130	.00	25	95	0	0
14:38:46	3390-9 ATP000 CP	0	4	14.8	.0	.258	11.9	.763	12.9	12.9	.203	.00	19	93	0	0
14:39:16	3390-9 ATP000 CP	0	4	17.7	.0	.232	9.60	.499	10.3	10.3	.059	.00	18	95	0	0
14:39:46	3390-9 ATP000 CP	0	4	12.0	.0	.272	11.7	.755	12.7	12.7	.090	.00	15	90	0	0
14:40:16	3390-9 ATP000 CP	0	4	19.7	.0	.206	8.54	.345	9.09	9.09	.017	.00	18	97	0	0
14:40:46	3390-9 ATP000 CP	0	4	13.3	.0	.327	16.9	.695	18.0	18.0	.087	.00	24	94	0	0
14:41:16	3390-9 ATP000 CP	0	4	16.1	.0	.270	13.1	.580	14.0	14.0	.092	.00	22	95	0	0
14:41:46	3390-9 ATP000 CP	0	4	10.4	.0	.364	21.1	.687	22.2	22.2	.067	.00	23	92	0	0
14:42:16	3390-9 ATP000 CP	0	4	15.7	.0	.279	14.1	.607	15.0	15.0	.094	.00	24	96	0	0
14:42:46	3390-9 ATP000 CP	0	4	14.3	.0	.307	15.1	.735	16.1	16.1	.050	.00	23	93	0	0
14:43:16	3390-9 ATP000 CP	0	4	16.5	.0	.260	13.1	.567	13.9	13.9	.069	.00	23	95	0	0
14:43:46	3390-9 ATP000 CP	0	4	10.5	.0	.309	14.6	.664	15.6	15.6	.038	.00	16	92	0	0
14:44:16	3390-9 ATP000 CP	0	4	28.4	.0	.098	1.24	.145	1.48	1.48	.005	.00	4	98	0	0

Command ==>

F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 67. Layout of General I/O Device Data Log Screen (FCX168) - Example 1

The FCX168 screen looks like this for a HyperPAV alias device:

FCX168	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor															
General I/O Data Log for Device BE10																				
Interval	<- Dev. Descr.->	Mdisk Pa-	<-Rate/s->	<----- Time (msec) ----->	Req. <Percent>	SEEK	Recov	<-Throttle->												
End Time	Type	Label/ID	Links	ths	I/O	Avoid	Pend	Disc	Conn	Serv	Resp	CUWt	Qued	Busy	READ	Cyls	SSCH	Set/s	Dly/s	SSID
>>Mean>>	HPAV	.....	...	4	10.4	.0	.258	10.1	3.42	13.7	13.7	2.18	.00	14	61	0	0	...	...	0600
14:34:16	HPAV	.....	...	4	13.0	.0	.255	11.1	2.49	13.9	13.9	1.49	.00	18	67	0	0	...	...	0600
14:34:46	HPAV	.....	...	4	8.0	.0	.286	13.3	4.21	17.8	17.8	2.94	.00	14	53	0	0	...	...	0600
14:35:16	HPAV	.....	...	4	11.8	.0	.234	9.55	2.67	12.5	12.5	1.57	.00	15	59	0	0	...	...	0600
14:35:46	HPAV	.....	...	4	8.4	.0	.287	12.0	4.59	16.8	16.8	2.95	.00	14	53	0	0	...	...	0600
14:36:16	HPAV	.....	...	4	12.0	.0	.248	10.3	2.93	13.5	13.5	1.81	.00	16	62	0	0	...	...	0600
14:36:46	HPAV	.....	...	4	10.1	.0	.277	10.7	4.40	15.3	15.3	2.99	.00	16	55	0	0	...	...	0600
14:37:16	HPAV	.....	...	4	12.6	.0	.239	9.96	2.61	12.8	12.8	1.54	.00	16	65	0	0	...	...	0600
14:37:46	HPAV	.....	...	4	10.1	.0	.282	10.3	4.35	14.9	14.9	3.09	.00	15	59	0	0	...	...	0600
14:38:16	HPAV	.....	...	4	13.5	.0	.242	8.66	2.72	11.6	11.6	1.59	.00	16	66	0	0	...	...	0600
14:38:46	HPAV	.....	...	4	8.6	.0	.275	10.4	4.32	15.0	15.0	2.57	.00	13	53	0	0	...	...	0600
14:39:16	HPAV	.....	...	4	12.8	.0	.236	8.47	2.90	11.6	11.6	1.64	.00	15	65	0	0	...	...	0600
14:39:46	HPAV	.....	...	4	9.0	.0	.277	10.4	4.00	14.7	14.7	2.38	.00	13	51	0	0	...	...	0600
14:40:16	HPAV	.....	...	4	9.8	.0	.235	8.99	3.08	12.3	12.3	2.09	.00	12	67	0	0	...	...	0600
14:40:46	HPAV	.....	...	4	8.9	.0	.302	13.4	4.62	18.3	18.3	3.42	.00	16	58	0	0	...	...	0600
14:41:16	HPAV	.....	...	4	11.4	.0	.266	11.2	3.19	14.6	14.6	2.15	.00	17	65	0	0	...	...	0600
14:41:46	HPAV	.....	...	4	9.6	.0	.289	10.6	4.36	15.3	15.3	2.84	.00	15	53	0	0	...	...	0600
14:42:16	HPAV	.....	...	4	13.0	.0	.256	9.97	2.76	13.0	13.0	1.58	.00	17	65	0	0	...	...	0600
14:42:46	HPAV	.....	...	4	9.1	.0	.272	9.60	4.37	14.2	14.2	2.91	.00	13	53	0	0	...	...	0600
14:43:16	HPAV	.....	...	4	13.0	.0	.253	9.48	2.99	12.7	12.7	1.88	.00	17	66	0	0	...	...	0600
14:43:46	HPAV	.....	...	4	8.0	.0	.268	9.17	4.37	13.8	13.8	2.70	.00	11	54	0	0	...	...	0600
14:44:16	HPAV	.....	...	4	5.2	.0	.127	1.60	1.36	3.09	3.09	.742	.00	2	79	0	0	...	...	0600
Command	==>																			
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return															

Figure 68. Layout of General I/O Device Data Log Screen (FCX168) - Example 2

The information shown is based on MONITOR domain and I/O domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand or since data collection was activated. It is designated by the >>Mean>> entry in the Interval End Time column.

## Field Descriptions:

### Interval End Time

The end time of the monitor interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the heading line.

*Dev. Descr.:*

### Type

The device type and model. Special cases are as follows:

- For PAV alias devices of an IBM DASD subsystem, the device type is replaced by the following string:

```
-> xxxx
```

where xxxx is the device number of the base RDEV.

- For HyperPAV aliases, the device type is replaced by the following string:

```
HPAV
```

### Label/ID

The label or ID, as follows (note that items 3, 4, and 5 apply to DASDs that are not attached to guests):

1. For a device that is attached to a guest, this field contains:

```
>userid
```

where *userid* is the user ID of the guest to which the device is attached.

2. For a device that is not a DASD, this field contains a label that consists of dots.
3. For a HyperPAV alias device, this field contains a label that consists of dots.
4. For a PAV alias device, this field contains the disk label of the base RDEV.
5. For all other DASDs, this field contains the disk label, and, if the device is in the CP-owned list, the label will have a suffix of CP.

### Mdisk Links

For DASDs that are:

- Attached to SYSTEM, this field contains the number of minidisk links.
- Not attached to SYSTEM, this field always contains a 0.

When the device is other than a non-PAV DASD RDEV, a PAV base RDEV, or a HyperPAV base RDEV, dots are displayed in this field.

### Paths

The number of available channel paths to the device. If the number of channel paths is low, and their usage is high, this could cause excessive reconnect misses, which would result in abnormally high disconnect times.

For emulated FBA devices (EDEV), this number is always 1, even when multiple channel connections to the device are available.

*Rate/s:*

### I/O

The number of SSCH and RSCH operations executed per second from this system during the selected interval.

A suffix of T will be appended to the I/O rate numbers of all devices with an I/O rate that is currently being "throttled".

**Avoid**

The number of real I/O operations per second that were avoided because the blocks were found in the minidisk cache.

*Time (msec):*

**Pend**

The "function pending" time. This is the average amount of time that an I/O request remained pending in the I/O subsystem due to path-busy conditions.

This value and the Disc, Conn, Serv, Busy, and Resp values are calculated based on information obtained from the subchannel measurement block. This information is important especially for disks. Dots are inserted for devices where the information is not available or is not calculated.

**Disc**

The "device disconnected" time. This is the average amount of time that the device remained disconnected from the channel while executing I/O requests. For DASDs, it includes the time spent on executing the SEEK and SET SECTOR orders and also any re-connect misses. High values could indicate overloaded paths or many long SEEKS to data at opposite ends of the device.

The value also includes the "device-active-only" time where that information is available.

**Conn**

The "device connected" time. This is the average time that the device was connected to the channel while executing I/O requests, that is, receiving commands or actually transferring data. For DASDs, it also includes the SEARCH time.

**Serv**

The service time for the device. This is the sum of the function pending, connected, and disconnected times.

**Resp**

The device response time. This is the device service time plus the time during which an I/O request was waiting to be started.

**CUWt**

The control unit queuing time, that is, the additional wait time caused when the target device was already busy with an I/O operation that was started from another system.

**Req. Qued**

The average I/O queue length in the selected interval.

The value is taken from CP monitor high-frequency sample data; it is an indicator for I/O contention on the device.

The I/O queue length (Req. Qued) is not meaningful for devices that are dedicated (attached) to guest systems. The guest system will not try to start another I/O operation before an active one has ended, but rather will do its own queuing. Guest system queuing cannot be shown by host performance monitors, so they will have to be analyzed using a performance monitor on the guest system.

The queue length reported here is the number of real I/O operations (channel programs) that CP's real I/O layer is waiting to start to the device or volume. This includes I/O operations performed for all reasons: for minidisk I/O, for paging I/O, and for spooling I/O, for example. For devices in the CP-owned list, additional queueing information can be found on one of these screens:

- [“FCX109, CP-Owned Device Screen – DEVICE CPOWNED” on page 416](#)
- [“FCX170, CPOWNED Device Screen – CPOWNLOG” on page 553.](#)

*Percent:*

**Busy**

The percentage of time during which the device was busy.

**READ**

The percentage of I/O requests that were for read operations only.

**SEEK Cyls**

The average number of cylinders skipped per seek order.

**Recov SSCH**

The number of SSCHs (*not* the rate) that were executed for error recovery purposes.

*Throttle:*

**Set/s**

The throttle I/O rate set for the device.

This field contains dots on the >>Mean>> line.

**Dly/s**

The rate at which arriving I/O operations were delayed due to throttling.

**SSID**

For DASDs that have an SSID, this field contains the SSID. Otherwise, this field contains dots.

**Usage note:**

Statistics depicted on the rows of this screen are calculated using only the measurement data for the listed device itself. The effects of PAV or HyperPAV are not included. For a depiction of volume-level statistics that account for the effect of PAV or HyperPAV, use the FCX329 screen (VOLUME) or the FCX330 screen (VOLLOG).

**Effect of AVERAGE and INTERIM subcommands:** None.

## FCX169, Control Unit Cache Log Screen – CACHDLOG

**Performance Toolkit Subcommand:** “CACHDLOG” on page 144.

The Control Unit Cache Log Screen (FCX169) shows 'by time' information on the cache activity and effectiveness for the selected disk. Note that the disk must be connected to a cached control unit which has previously been selected for benchmarking.

FCX169	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS					Perf. Monitor		
Control Unit Cache Log for Device 0EC2										
Interval		Overall CU-Cache Performance						Split		
End Time	VOLSER	IO/S	%READ	%RDHIT	%WRHIT	ICL/S	BYP/S	IO/S	%READ	%RDHIT
>>Mean>>	ESAPK2	3.9	49	92	0	.0	.0	3.9	49	92 (N)
>>Mean>>		.0	0	0				.0	0	0 (S)
>>Mean>>		.0	0	0				.0	0	0 (F)
10:35:00	ESAPK2	2.4	49	94	0	.0	.0	'NORMAL'	I/O	only
10:36:00	ESAPK2	1.4	50	97	0	.0	.0	'NORMAL'	I/O	only
10:37:00	ESAPK2	.4	42	88	0	.0	.0	'NORMAL'	I/O	only
10:38:00	ESAPK2	1.5	49	93	0	.0	.0	'NORMAL'	I/O	only
10:39:00	ESAPK2	1.2	43	87	0	.0	.0	'NORMAL'	I/O	only
10:40:00	ESAPK2	3.2	35	94	0	.0	.0	'NORMAL'	I/O	only
10:41:00	ESAPK2	4.4	30	92	0	.0	.0	'NORMAL'	I/O	only
10:42:00	ESAPK2	5.8	43	94	0	.0	.0	'NORMAL'	I/O	only
10:43:00	ESAPK2	5.3	32	95	0	.0	.0	'NORMAL'	I/O	only
10:44:00	ESAPK2	5.3	40	94	0	.0	.0	'NORMAL'	I/O	only
10:45:00	ESAPK2	3.2	50	78	0	.0	.0	'NORMAL'	I/O	only
10:46:00	ESAPK2	6.1	49	92	0	.0	.0	'NORMAL'	I/O	only
Command ==> _										
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return					

Figure 69. Layout of Control Unit Cache Log Screen (FCX169)

The information shown is based on MONITOR and I/O domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'Dir' and 'Addr' fields, all data columns are identical to the ones of the 'CU-cached Disks Screen'. Please see “FCX111, CU-Cached Disks Screen – CACHDBSE” on page 428 for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.



## FCX170, CPOWNED Device Log Screen – CPOWNLOG

**Performance Toolkit Subcommand:** “CPOWNLOG” on page 156.

The CPOWNED Device Log Screen (FCX170) shows the availability and utilization of CP system areas on the selected disk, which must have been previously selected for benchmarking, and it is available only for CP-owned disks that contain some such areas.

```
FCX170      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor
CPOWNED Device Log for Device 5CCC

<- Device -->
Interval  Volume Area  Area  Used  <----- Rate/s ----->  I/O  Serv MLOAD Block %Used I
End Time Devtyp Serial Type Extent %  <---Page---> <---Spool--->  SSCH Inter Queue Time MLOAD Resp Page %Used I
          3390-9 FC5CCC DUMP 11793600 0  P-Rds P-Wrt S-Rds S-Wrt Total +RSCH feres Lngth /Page Time Resp Time Page Size Alloc M
>>Mean>> 3390-9 FC5CCC DUMP 11793600 0  .0 .0 .0 .0 .0 .0 .0 0 0 0 ... .. T
06:45:35 3390-9 FC5CCC DUMP 11793600 0  .0 .0 .0 .0 .0 .0 .0 0 0 0 ... .. T
06:46:35 3390-9 FC5CCC DUMP 11793600 0  .0 .0 .0 .0 .0 .0 .0 0 0 0 ... .. T
06:47:35 3390-9 FC5CCC DUMP 11793600 0  .0 .0 .0 .0 .0 .0 .0 0 0 0 ... .. C
06:48:35 3390-9 FC5CCC DUMP 11793600 0  .0 .0 .0 .0 .0 .0 .0 0 0 0 ... .. T

Command ==> _
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return
```

Figure 70. Layout of CPOWNED Device Log Screen (FCX170)

The information shown is based on MONITOR and STORAGE domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand or since data collection was activated. It is designated by the string >>Mean>> in the Interval End Time column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the heading line.

Device:

#### Devtyp

Device type and model of the disk

#### Volume Serial

Disk label

#### Area Type

Describes what the system area is used for. The area type indicated can be:

##### DIRECT

For directory cylinders

##### DUMP

For a spooling area on DUMP disks

##### PAGE

For a paging area

##### SPOOL

For a general spooling area

##### T-DISK

For temporary minidisk space

#### Area Extent

Start and end cylinders for the system area, or start and end slot number for FBA disks. Where the cylinder or block number range would require more space than is available for this field, the number of slots in the area is inserted instead.

#### Used %

Percentage of slots currently in use in this area. In the case of TDSK and DIRECTORY space, the "used" percentage is based on the number of used cylinders, not blocks.

Valid information for the remainder of the performance fields can be shown only for PAGE, SPOOL, and DUMP areas.

*Rate/s:*

**P-Rds**

The average page read rate during the last interval. This field and the "Page Write" (P-Wrt) field are inserted only for the first page area on a specific disk, but show the *total* activity on *all* PAGE areas on the same disk pack.

**P-Wrt**

The average page write rate during the last interval.

**S-Rds**

The average spool read rate during the last interval. This field and the "Spool Write" (S-Wrt) field are inserted only for the first spool area on a specific disk, but show the *total* activity on *all* SPOOL areas on the same disk pack.

**S-Wrt**

The average spool write rate during the last interval.

**Total**

The total page and spool read and write activity during the selected interval.

**SSCH+RSCH**

The number of SSCH and RSCH operations executed per second from this system during the selected interval.

**I/O Interferes**

The number of I/O requests that had interference in the last interval (for current data) or the average number of I/O requests which had interference per monitor sample interval if average data are shown. Some examples of I/O interference are:

- CP I/O done to the volume on behalf of privileged CP commands, for example: QUERY PATHS, QUERY CACHE, HYPERSWAP, and SNAPDUMP.
- CP I/O done on behalf of MONITOR data collection, for example: Cache statistics in the Device Activity record.
- CP I/O done on behalf of CP error recovery, for example: State change processing.
- User I/O to non-paging/non-spool extents on the volume.

**Queue Lngth**

The number of page and spool reads and writes that are currently active or queued in CP's paging layer. This is a single sample if "current" data are being shown; it is the average of all samples since the last reset if "average" data are shown.

Note that this value reflects queue lengths for CP paging and spooling activity only; it does *not* indicate anything about the queueing of other I/O to the same device (that is, by users).

For more information about I/O queueing at the device, see the Req. Qued field of [“FCX108, General I/O Device Screen – DEVICE or DEVICE DASD”](#) on page 412.

**Serv Time/Page**

The smoothed service time for processing a page, in milliseconds (msec).

**MLOAD Resp Time**

The smoothed response time for CP I/O operations to this area, in milliseconds (msec).

The value is normalized to reflect the response time for servicing a single page. It also takes into account the PIOBK queue length for the device.

**Block Page Size**

The average number of pages in the blocks available on the device.

**%Used for Alloc**

Percentage of times the volume was chosen for allocation when looked at while choosing a volume.

**IOM**

The I/O protocol that is used to communicate with the device:

**C**

For command mode

**T**

For transport mode

**E**

If the device is an EDEV.

The >>Mean>> line will contain the value if unchanged over the entire time or a dot if this value has changed.

**Effect of AVERAGE and INTERIM subcommands:**

None.

## FCX171, DASD Seeks Locations Log Screen – SEEKLLOG

**Performance Toolkit Subcommand:** “SEEKLLOG” on page 287.

The DASD Seeks Locations Log Screen (FCX171) shows a 'by time' log of seeks locations for all I/O activity to the selected disk (which must have been previously selected for benchmarking).

FCX171		CPU nnnn		SER nnnnn		Interval HH:MM:SS - HH:MM:SS				Perf. Monitor			
SEEKS Locations Log for Device 0EC2													
<----- Device ----->					<----- Seeks ----->					<----- Per Within P ----->			
Interval	Device	Volume	Size	Disc	Cach	Pct	Pct	Mean	Total	0<=	10<	20<	Per
End Time	Type	Serial	Cyls	in Time	ed	Read	Non0	Dist	Seeks	<=10	<=20	<=30	<
>>Mean>>	3380-E	ESAPK2	1770	6.2	Yes	54	15	64	30079	7	2	85	
10:35:00	3380-E	ESAPK2	1770	6.4	Yes	50	10	88	146	3	0	97	
10:36:00	3380-E	ESAPK2	1770	5.3	Yes	53	10	80	80	3	1	94	
10:37:00	3380-E	ESAPK2	1770	9.1	Yes	48	29	164	21	3	2	93	
10:38:00	3380-E	ESAPK2	1770	6.7	Yes	51	18	90	89	2	0	98	
10:39:00	3380-E	ESAPK2	1770	6.3	Yes	44	13	135	72	2	0	98	
10:40:00	3380-E	ESAPK2	1770	8.7	Yes	46	22	79	169	19	14	26	
10:41:00	3380-E	ESAPK2	1770	8.1	Yes	33	14	64	237	63	14	22	
10:42:00	3380-E	ESAPK2	1770	6.2	Yes	44	10	44	346	6	11	83	
10:43:00	3380-E	ESAPK2	1770	8.0	Yes	37	15	80	290	51	0	42	
10:44:00	3380-E	ESAPK2	1770	6.9	Yes	46	18	65	290	64	16	15	
10:45:00	3380-E	ESAPK2	1770	7.9	Yes	51	20	213	189	29	0	71	
10:46:00	3380-E	ESAPK2	1770	5.5	Yes	50	9	28	367	19	0	42	
Command ==> _													
F1=Help		F4=Top		F5=Bot		F7=Bkwd		F8=Fwd		F10=Left		F11=Right F12=Return	

Figure 71. Layout of DASD Seeks Locations Log Screen (FCX171)

The information shown is based on MONITOR and I/O domain SAMPLE data and on SEEKS domain EVENT data.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

Note that the information extends beyond the 80 columns of an ordinary screen: use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window so as to view the remaining data.

### Field Description:

#### Interval End Time

The time when the last seek operation was monitored in the interval, or the end time of the sample interval when no seeks activity was recorded in an interval.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'Addr' field, all data columns are identical to the ones of the 'DASD Seeks Locations Screen'. Please see “FCX156, DASD Seeks Locations Screen – SEEKLOC” on page 532 for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX172, DASD Seeks Distances Log Screen – SEEKDLOG

**Performance Toolkit Subcommand:** “SEEKDLOG” on page 286.

The DASD Seeks Distances Log Screen (FCX172) shows a 'by time' log of the average seeks distances measured for all I/O activity to the selected disk (which must have been previously selected for benchmarking).

FCX172	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS						Perf. Monitor			
SEEKS Distances Log for Device 0EC2												
<----- Device ----->					<----- Seeks ----->					<----- Perc		
Interval	Device	Volume	Size	Disc				Mean				
End Time	Type	Serial	in	Time	Cach	Pct	Pct	Non-0	Total	=	<=	<=
>>Mean>>	3380-E	ESAPK2	1770	6.2	Yes	54	15	64	30079	85	6	1
10:35:00	3380-E	ESAPK2	1770	6.4	Yes	50	10	88	146	90	3	3
10:36:00	3380-E	ESAPK2	1770	5.3	Yes	53	10	80	80	90	3	4
10:37:00	3380-E	ESAPK2	1770	9.1	Yes	48	29	164	21	71	0	10
10:38:00	3380-E	ESAPK2	1770	6.7	Yes	51	18	90	89	82	4	1
10:39:00	3380-E	ESAPK2	1770	6.3	Yes	44	13	135	72	88	0	0
10:40:00	3380-E	ESAPK2	1770	8.7	Yes	46	22	79	169	78	11	0
10:41:00	3380-E	ESAPK2	1770	8.1	Yes	33	14	64	237	86	9	0
10:42:00	3380-E	ESAPK2	1770	6.2	Yes	44	10	44	346	90	7	0
10:43:00	3380-E	ESAPK2	1770	8.0	Yes	37	15	80	290	85	7	0
10:44:00	3380-E	ESAPK2	1770	6.9	Yes	46	18	65	290	82	10	0
10:45:00	3380-E	ESAPK2	1770	7.9	Yes	51	20	213	189	80	4	2
10:46:00	3380-E	ESAPK2	1770	5.5	Yes	50	9	28	367	91	5	3
Command ==> _												
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F10=Left	F11=Right	F12=Return					

Figure 72. Layout of DASD Seeks Distances Log Screen (FCX172)

The information shown is based on MONITOR and I/O domain SAMPLE data and on SEEKS domain EVENT data.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

Note that the information extends beyond the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window so as to view the remaining data.

### Field Description:

#### Interval End Time

The time when the last seek operation was monitored in the interval, or the end time of the sample interval when no seeks activity was recorded in an interval.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'Addr' field, all data columns are identical to the ones of the 'DASD Seeks Distances Screen'. Please see “FCX157, DASD Seeks Distances Screen – SEEKDIST” on page 534 for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX173, Benchmark Log Selection Menu Screen – BENCHMARK

**Performance Toolkit Subcommand:** [“BENCHMARK”](#) on page 138.

After benchmarking has been activated, you can select the Benchmark Log Selection Menu Screen (FCX173):

```
FCX173      CPU nnnn  SER nnnnn  BENCHMARK Log Data  Perf. Monitor
-----
Object      Log File
S ID        Name      Description
. LXA000001 UCOMMLOG  User IUCV and VMCF communications log
. LXA000001 UPAGELOG  User paging load log
. LXA000001 URESPLOG  User response time log
. LXA000001 USERLOG   User resource consumption log
. LXA000001 USRTMLOG  User processor time Log
. LXA000001 USTATLOG  User wait state log
. LXA000001 UTRANLOG  User resources per transaction log
.
.
. BE10      DEVLOG   General I/O device performance log
. BE10      HPFLOG   HPF I/O device activity log
.
.
. BE1F      DEVLOG   General I/O device performance log
. BE1F      HPFLOG   HPF I/O device activity log
. BE1F      VOLLOG   DASD Volume performance log

Select an object log with cursor and hit ENTER
Command ==>
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 73. Layout of Benchmark Log Selection Menu Screen (FCX173)

The reports offered for selection are dependent on the objects that you have benchmarked using the FCONTROL BENCHMARK command. For a list of the types of objects that can be benchmarked and the log files that can be created, see [“FCONTROL BENCHMARK”](#) on page 15.

**Field Descriptions:**

**S**

The cursor selection field.

**Object ID**

The object for which performance is logged. An object can be a user ID, a device number of an I/O device, an ISFC logical link name, an ISFC end point ID, or a PCI function ID, for example.

**Log File Name**

The command with which the log data can be selected for display.

**Description**

The log file description.

See the "Benchmarking Objects" section in Chapter 5 of the [z/VM: Performance Toolkit Guide](#) for more information about the FCX173 screen.

## FCX174, User Transaction Details Screen – UTRANDET

**Performance Toolkit Subcommand:** “UTRANDET” on page 355.

The User Transaction Details Screen (FCX174):

```
FCX174      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS    Perf. Monitor

User Transaction Details for USER1

Absolute share per processor:   3%           Max. abs. share per processor: 10%

      C   Resp  <----- Resource Utilization per Transaction ----->
      P   Time  CPU   Disk   Page   Xpage  <-Spooling->  <-IUCV/VMCF->
Time Stamp U (sec) msec  I/O   Flts  Reads  SP_RD  SP_WR  IU_BY  IU_TO
09:59:50.161 00 Add to elig. list: Class 1, E1= 1 E2= 0 E3= 0, Prior.: .134
09:59:50.161 00 Add to disp. list: Class 1, Q0= 5 Q1=11 Q2=12 Q3= 2
                  D-List Priority: .000 ( .008), IABIAS
                  Tot. rel. share=39400, Tot. abs. share= 3%, ETS=0.545
                  PGS: Total available=03663, Total WSS=10100, Our WSS=01100
                      RES=01122, LIM=50683, XSTORE=00140, Reads=0000/s
                  Last SEND to *MONITOR
09:59:50.173 -   5 SEEs to MDISK 0191:  0% Reads, Av. Displ.   3 Cyl.
09:59:50.173 00 Drop from disp. list:
                  Scheduler State: Idle, MTS:    1
                  Last SEND to *MONITOR
>>Trans. End 0  .01    3    0    0    0    0    0    0
10:00:16.225 00 Add to elig. list: Class 1, E1= 1 E2= 0 E3= 0, Prior.: .182
10:00:16.225 00 Add to disp. list: Class 1, Q0= 6 Q1=14 Q2= 6 Q3= 4
                  D-List Priority: .008 ( .012), IABIAS
                  Tot. rel. share=43920, Tot. abs. share= 3%, ETS=0.571
                  PGS: Total available=03169, Total WSS=10778, Our WSS=01101
                      RES=01122, LIM=50682, XSTORE=00140, Reads=0000/s
                  IUCV/VMCF int. pending, Last SEND to *MONITOR
10:00:16.797 00 Drop from disp. list: SIM-Wait
                  Scheduler State: Susp. short-time event, MTS:    3
                  Last SEND to *MONITOR
10:00:16.826 00 Add to elig. list: Class 0, E1= 0 E2= 0 E3= 0, Prior.: 99999
10:00:16.826 00 Add to disp. list: Class 0, Q0= 7 Q1=12 Q2= 7 Q3= 4
                  D-List Priority: .000 ( .008), Lock Shot, IABIAS
                  Tot. rel. share=51700, Tot. abs. share= 3%, ETS=3.446
                  PGS: Total available=03050, Total WSS=10045, Our WSS=01122
                      RES=01125, LIM=00001, XSTORE=00140, Reads=0000/s
                  Last SEND to *MONITOR
10:00:16.826 00 Drop from disp. list:
                  Scheduler State: Selected for processing, MTS:    0
                  Last SEND to *MONITOR

Command ==>
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 74. Layout of User Transaction Details Screen (FCX174)

The information shown is based on CP monitor USER domain EVENT data and on SCHEDULER domain EVENT data.

Note that EVENT data collection for all users can cause considerable overhead on a large system. Activate this part of data collection only for periods where you really need the information for analysis, or only for a subset of users.

The sample shows a screen where both USER and SCHEDULER domain event data were enabled for the user. The display can also be selected when only USER domain event data are available, but the queue add/drop entries will then be missing.

### Field Descriptions:

#### xxxxxxx share per processor

Indicates the virtual machine's share, as specified in its directory or by command. xxxxxxxx is Relative or Absolute.

**Max. xxx. share per processor**

Indicates the virtual machine's maximum share, as specified in its directory or by command. xxx. is rel. or abs. to indicate whether a relative or absolute share maximum has been set.

The maximum share string is omitted where no maximum share has been set for a virtual machine.

**Time Stamp**

The exact time, in hours, minutes, seconds, and milliseconds when a particular event was recorded.

The string >>Trans. End is inserted instead for the line that contains the resource consumption data for the preceding transaction when scheduler domain data were also processed for the user.

**CPU**

The processor ID of the virtual CPU in hexadecimal format.

**Resp Time (sec)**

The response time for the transaction. This field, and also all of the remaining field descriptors in the header lines, applies only to the transaction end data lines.

*Resource Utilization per Transaction:*

Usage figures for a variety of resources are shown under this heading.

**CPU msec**

The CPU time used per transaction, in milliseconds.

**Disk I/O**

The number of disk I/O requests per transaction.

**Page Flts**

The number of page faults per transaction.

**Xpage reads**

The number of pages moved from expanded storage to central storage per transaction.

**Spooling SP\_RD**

The number of UR read operations from spool per transaction.

**Spooling SP\_WR**

The number of UR write operations to spool per transaction.

**IU\_BY**

The number of IUCV and VMCF data transfers from this user to another virtual machine, or to CP, per transaction.

**IU\_TO**

The number of IUCV and VMCF data transfers to this user, per transaction.

*Scheduler activities:*

The three possible entries are:

1. Add to elig. list
2. Add to disp. list
3. Drop from disp. list

Descriptions of each entry follow:

Add to elig. list entries have the following general layout:

```
Add to elig. list: Class n, E1= m E2= m E3= m, Prior.: p.ppp
```

where:

**Class**

Indicates the class (0-3) which was assigned to the transaction.

**E1=**

Indicates the total number of class 1 users in the eligible list.



**E2=**

Indicates the total number of class 2 users in the eligible list.

**E3=**

Indicates the total number of class 3 users in the eligible list.

**Prior.**

Indicates the eligible list priority assigned to the user.

Add to disp. list entries have the following general layout:

```
Add to disp. list: Class n, Q0= m Q1= m Q2= m Q3= m
D-List Priority: .nnn ( .nnn), aaaaaaaa
Tot. rel. share=nnnnn, Tot. abs. share=nnn%, ETS=n.nnn
PGS: Total available=nnnnn, Total WSS=nnnnn, Our WSS=nnnnn
RES=nnnnn, LIM=nnnnn, XSTORE=nnnnn, Reads=nnnn/s
bbbbbbbb
```

where:

**Class**

Indicates the class (0-3) which was assigned to the transaction.

**Q0=**

Indicates the total number of class 0 users in the dispatch list (quick dispatch).

**Q1=**

Indicates the total number of class 1 users in the dispatch list.

**Q2=**

Indicates the total number of class 2 users in the dispatch list.

**Q3=**

Indicates the total number of class 3 users in the dispatch list.

**D-List Priority**

Indicates the user's actual dispatch list priority, and (in parentheses) the dispatch list priority before adjustments for factors such as interactive bias and paging bias.

**aaaaaaaa**

Can indicate one or more of the following conditions:

- Hot Shot
- Lock Shot
- Hot Shot Req.
- IABIAS
- Loading

**Tot. rel. share**

Is the sum of the relative shares of all users in the dispatch list.

**Tot. abs. share**

Is the sum of the absolute shares of all users in the dispatch list.

**ETS**

Is the current length of the elapsed time slice for the user's class.

**PGS:**

Indicates that the following fields contain information on storage usage and paging.

**Total available**

Is the total number of pages available to the dispatcher.

**Total WSS**

Is the total working set size, in pages, of all users in the dispatch list.

**Our WSS**

Is the size of the monitored user's working set, in pages.

**RES**

Is the number of the monitored user's resident pages.

**LIM**

Is the page growth limit.

**XSTORE**

Is the number of allocated XSTORE blocks.

**Reads**

Is the page read rate.

**bbbbbbbb**

Can provide more special status information. A combination of the following strings can be shown:

**SVM Wait**

If the user was in SVM wait for IUCV/VMCF.

**IUCV/VMCF int. pending**

If an IUCV or VMCF external interrupt was pending.

**Last SEND to**

Provides the target of the last IUCV or VMCF SEND operation.

Drop from disp. list entries have the following general layout:

```
Drop from disp. list: ccccccc
Scheduler State: ddddddd, MTS:  nn
eeeeeee
```

where:

**cccccccc**

Can be a combination of the following:

- Pre-empted (due to large storage requirements)
- CF-Wait
- SIM-Wait
- I/O-Wait

**Scheduler State**

One of the following states is inserted for *ddddddd*:

- Review suspended
- Review idle
- Selected for processing
- Ready for dispatcher
- Test-idle
- Susp. short time event
- Idle

**MTS**

Shows the number of minor time slices which have expired since the beginning of the transaction.

**eeeeeeee**

Can be a combination of the following:

**SVM Wait**

If the user was in service machine wait for IUCV/VMCF.

**IUCV/VMCF int. pending**

If an IUCV or VMCF external interrupt was still pending.

**Last SEND to**

Provides the target of the last IUCV or VMCF SEND operation.

### Seeks information:

Where seeks data collection has been enabled for disks with some I/O activity from the monitored user, additional entries with seeks information might be included in the log, as follows:

```
nnnn SEEKs to MDISK xxxx: rr% Reads, Av. Displ. ccccc Cyl.
```

where:

**nnnn**

Is the number of seek operations to the minidisk since the last queue drop.

**xxxx**

Is the virtual device number of the minidisk.

**rr**

Is the percentage of read operations.

**cccc**

Is the average access arm displacement in cylinders.

**Note:** Minidisk seeks statistics show the I/O activity to a minidisk for the whole period since the last queue drop. They are inserted when a queue drop monitor record is processed and show the same time stamp as that queue drop, that is, this is *not* the time of the actual I/O activity.

### Effect of AVERAGE and INTERIM subcommands:

None.

## FCX175, Scheduler Data Log Screen – SCHDATA

**Performance Toolkit Subcommand:** “SCHDATA” on page 280.

The Scheduler Data Log Screen (FCX175) shows a 'by time' log of some of the smoothed scheduler data fields which are used by scheduler algorithms:

FCX175	CPU	nnnn	SER	nnnnn	Interval	HH:MM:SS	-	HH:MM:SS	Perf.	Monitor
Scheduler	<Page	Read>	Page	<-Resource->	<-- Aver. User Load -->	<Expansion>				
Interval	Rate	msec	Steal	<-Weights-->	<-Class 1->	All Classes	<-Factors->			
End Time	Pgs/s	/Pg	Value	Stor. Paging	Pgs/s	WSS	Pgs/s	WSS	C11	C12 C13
>>Mean>>	203	18	0	16	1	0	317	1	316	2.0 2.0 2.0
09:34:01	209	17	0	14	0	0	318	1	312	2.0 2.0 2.0
09:35:01	199	17	0	14	1	0	330	1	326	2.0 2.0 2.0
09:36:01	200	17	0	17	2	0	334	1	317	2.0 2.0 2.0
09:37:01	182	17	0	14	0	0	320	1	307	2.0 2.0 2.0
09:38:01	198	17	0	15	0	0	316	1	333	2.0 2.0 2.0
09:39:01	202	17	0	18	0	0	328	1	324	2.0 2.0 2.0
09:40:01	203	17	0	17	2	0	324	1	313	2.0 2.0 2.0
09:41:01	194	18	0	18	0	0	316	1	306	2.0 2.0 2.0
09:42:01	195	19	0	16	0	0	306	1	317	2.0 2.0 2.0
09:43:01	192	19	0	15	0	0	325	1	343	2.0 2.0 2.0
09:44:01	200	19	0	14	1	0	320	1	315	2.0 2.0 2.0
09:45:01	205	19	0	13	0	0	317	1	289	2.0 2.0 2.0
09:46:01	194	19	0	15	2	0	319	1	306	2.0 2.0 2.0
09:47:01	202	19	0	18	1	0	320	1	302	2.0 2.0 2.0
09:48:01	214	19	0	19	2	0	333	1	330	2.0 2.0 2.0

Command ==>>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 75. Layout of Scheduler Data Log Screen (FCX175)

The information shown is based on SCHEDULER domain EVENT data.

The leftmost column contains the time stamp of the last data record

(format hh:mm:ss)

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Scheduler Interval End Time' column.

### Field Description:

#### Scheduler Interval End Time

The end time of the last scheduler interval or the BYTIME interval which fell into the last monitor sample interval.

Note that scheduler data records are generated in an interval which is independent of the monitor sample interval. Since the length of the interval is currently 30 seconds, more than one such data record will usually be generated in one monitor sample interval. Average values are then calculated for all records which were obtained during the last monitor sample interval.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Page Read shows the following information on page reads:

#### Rate Pgs/s

The page read rate, per second

#### msec/Pg

The elapsed time per page read

#### Page Steal Value

Corresponds to the STEAL parameter in the INDICATE LOAD subcommand

### *Resource Weights:*

#### **Stor.**

The resource weight for storage. Used by the scheduler to decide how much of a bottleneck the storage resources are.

#### **Paging**

The resource weight for paging. Used by the scheduler to decide how much of a bottleneck the paging resources are.

*Average User Load* shows paging and storage consumption for the average user:

#### **Class 1 Pgs/s**

The average page rate of class 1 users at the end of their Q1 dispatch list stay

#### **Class 1 WSS**

The average working set size of class 1 users at the end of their Q1 dispatch list stay

#### **All Classes Pgs/s**

The average page rate of all users at the end of their dispatch list stay

#### **All Classes WSS**

The average working set size of all users at the end of their dispatch list stay

### *Expansion Factors:*

#### **CI1**

The ratio of time it is taking for a short transaction to complete, compared with the amount of time it would take if it were running alone on the system.

#### **CI2**

The ratio of time it is taking for a medium transaction to complete, compared with the amount of time it would take if it were running alone on the system.

#### **CI3**

The ratio of time it is taking for a long transaction to complete, compared with the amount of time it would take if it were running alone on the system.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX176, Cached Control Units Screen – CTLUNIT

**Performance Toolkit Subcommand:** “CTLUNIT” on page 164.

The Cached Control Units Screen (FCX176) shows overall load information on cached control units, the use made, and the effectiveness of the cache. Control units that have been active during the last measuring interval will be highlighted on the screen.

FCX176		CPU nnnn		SER nnnnn		Interval HH:MM:SS		Perf. Monitor	
Sub-sys- tem	Control	<-----Cache Size----->		<-----DASD Load Data ----->		<-----Cache Utilization Data ----->		<-----Percentages ----->	
		<Volatile-->	<Non-Volat-->	<--Total-->	<I/O rates>	<Read Rates/s-->	<--TotWrt-->	<--TotWrt-->	<--TotWrt-->
ID	Unit	Conf Avail	NV-Cf NV-Av	Cache SCMBK	Busy Pend Disc Conn Serv Resp	N-Seq Seq	FW Read %DFW %CFW	RdHt Wrt DFW CFW	
0362	2107-E8	192k	172k 6291k 6291k	.0 6.7	0 .099 .000 .080 .179 .179	.0 .0	.0 .0	.0 .0	.0 .0
0363	2107-E8	192k	172k 6291k 6291k	.0 6.7	0 .100 .000 .084 .184 .184	.0 .0	.0 .0	.0 .0	.0 .0
0364	2107-E8	192k	172k 6291k 6291k	.0 6.7	0 .098 .000 .080 .178 .178	.0 .0	.0 .0	.0 .0	.0 .0
0600	2107-E8	192k	172k 6291k 6291k	631.4 419.2	9 .168 1.73 1.41 3.31 8.19 487.4	27.7 116.3	82 100	0 93 100 100	.0 .0
0601	2107-E8	192k	172k 6291k 6291k	624.5 412.3	9 .168 1.92 1.37 3.45 8.75 483.4	26.0 115.0	82 100	0 92 100 100	.0 .0
0602	2107-E8	192k	172k 6291k 6291k	632.8 420.5	8 .164 1.44 1.40 3.01 7.17 489.3	27.4 116.1	82 100	0 93 100 100	.0 .0
0603	2107-E8	192k	172k 6291k 6291k	627.1 415.8	9 .166 1.69 1.41 3.27 8.15 486.6	25.4 115.1	82 100	0 93 100 100	.0 .0
0604	2107-E8	192k	172k 6291k 6291k	634.7 423.0	8 .164 1.43 1.38 2.97 6.68 491.4	27.5 115.8	82 100	0 93 100 100	.0 .0
0605	2107-E8	192k	172k 6291k 6291k	634.9 423.4	8 .166 1.60 1.41 3.17 7.85 492.3	26.4 116.2	82 100	0 93 100 100	.0 .0
0606	2107-E8	192k	172k 6291k 6291k	629.2 417.5	8 .167 1.58 1.41 3.16 7.49 485.6	27.9 115.6	82 100	0 92 100 100	.0 .0
0607	2107-E8	192k	172k 6291k 6291k	630.9 419.8	8 .166 1.62 1.39 3.18 7.82 487.0	28.0 115.9	82 100	0 93 100 100	.0 .0
0608	2107-E8	192k	172k 6291k 6291k	632.9 419.9	8 .164 1.46 1.39 3.01 6.82 486.5	31.5 114.9	82 100	0 93 100 100	.0 .0
0609	2107-E8	192k	172k 6291k 6291k	633.2 421.0	9 .167 1.70 1.45 3.32 8.63 490.2	25.9 117.1	82 100	0 92 100 100	.0 .0
060A	2107-E8	192k	172k 6291k 6291k	627.1 417.7	8 .167 1.58 1.46 3.21 7.92 482.2	28.7 116.2	81 100	0 93 100 100	.0 .0
060B	2107-E8	192k	172k 6291k 6291k	626.9 415.6	8 .165 1.56 1.39 3.11 7.60 486.6	25.0 115.4	82 100	0 93 100 100	.0 .0
060C	2107-E8	192k	172k 6291k 6291k	630.6 416.9	8 .166 1.60 1.39 3.16 7.54 486.7	27.1 116.8	81 100	0 93 100 100	.0 .0
060D	2107-E8	192k	172k 6291k 6291k	631.4 419.8	9 .168 1.79 1.44 3.39 8.64 488.4	26.5 116.5	82 100	0 92 100 100	.0 .0
060E	2107-E8	192k	172k 6291k 6291k	627.8 416.8	8 .166 1.55 1.44 3.15 7.66 485.6	26.6 115.7	82 100	0 93 100 100	.0 .0
060F	2107-E8	192k	172k 6291k 6291k	629.4 416.8	10 .170 2.04 1.47 3.68 10.0 484.6	28.3 116.5	81 100	0 92 100 100	.0 .0
0610	2107-E8	192k	172k 6291k 6291k	.0 .3	0 .129 .000 .248 .376 .376	.0 .0	.0 .0	.0 .0	.0 .0

Command ==> \_  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 76. Layout of Cached Control Units Screen (FCX176)

The information shown is based on MONITOR and I/O domain data.

The sample shows a screen with performance information for several 2107 control units, sorted in ascending order by subsystem ID (as indicated by the underscore characters above the Subsystem ID heading). Other possible sort sequences are indicated by dots above the corresponding columns. These sequences can be set using the SORT subcommand with the column heading (see “SORT” on page 295) or by positioning the cursor on the heading of the field that is to be sorted and then pressing Enter.

### Field Descriptions for Overall Performance:

#### Subsystem ID

Is the control unit's subsystem identifier

#### Control Unit

Is the control unit type and model.

Different models of IBM DASD subsystems will be recognized and displayed as such only if they are operating in enhanced mode; otherwise they will be shown as 3990-3 control units.

Cache Size provides information about the amount of storage that is available for caching:

#### Conf

Shows the size of configured cache storage, in megabytes.

#### Avail

Shows the amount of cache storage that is available, in megabytes.

#### NV-Cf

Shows the size of configured non-volatile storage, in kilobytes.

#### NV-Av

Shows the amount of non-volatile storage that is available, in kilobytes

DASD Load Data provides overall load information about the performance of all disks that are connected to the cached control unit:

**I/O rates Cache**

The rate at which channel command word (CCW) chains are arriving at the cached control unit, per second. A channel program consists of one or more CCW chains, often delimited by Locate Record CCWs. Also, chains could be arriving from more than one partition.

**I/O rates SCMBK**

The rate at which the reporting partition is sending channel programs to the cached control unit, per second. Sometimes also known as the Start Subchannel rate.

**Pct Busy**

The average busy percentage for all connected disks.

**Pend**

The average function pending time (msec). This is the average time from when the channel subsystem starts the I/O to when the channel subsystem receives the control unit's initial response (IR).

**Disc**

The average disconnected time (msec). This is the average time that the device remained disconnected from the channel while executing I/O requests. For DASDs, it includes the time spent for executing the SEEK and SET SECTOR orders and also any re-connect misses. High values usually indicate cache misses in the control unit.

This value also includes the device-active-only time where that information is available.

**Conn**

The average connected time (msec). This is the average time that the device was connected to the channel while executing I/O requests, that is, receiving commands or transferring data. It also includes the SEARCH time.

**Serv**

The average service time, that is, the sum of function pending, disconnected, and connected time.

**Resp**

The device response time. This is the device service time plus the amount of time CP waited to start the I/O.

**Note:** Averages for the control unit queueing time are not included on this display due to space restrictions. See the CUWt field of [“DEVICE” on page 169](#) for queueing times of individual DASDs.

*Cache Utilization Data* provides overall information about the activity and efficiency of the cached control unit.

*Total Rates/s:*

**Read N-Seq**

The arrival rate of read-normal chains or read Cache Fast Write (CFW) chains, per second.

**Read Seq**

The arrival rate of read-sequential chains, per second.

**Write FW**

The arrival rate of write DASD Fast Write (DFW) chains or write CFW chains, per second.

*Percentages:*

**Read**

The percentage of arriving chains that are read chains.

**/TotWrt %DFW**

The percentage of arriving write chains that are write DFW chains.

**/TotWrt %CFW**

The percentage of arriving write chains that are write CFW chains.

**Hits RdHt**

The percentage of arriving read chains satisfied without accessing a disk drive module (DDM).

**Hits Wrt**

The percentage of arriving write DFW chains or write CFW chains satisfied without accessing a DDM.

**Hits DFW**

The percentage of arriving write DFW chains satisfied without accessing a DDM.

**Hits CFW**

The percentage of arriving write CFW chains satisfied without accessing a DDM.

**Usage note:**

The entire analysis is based on counters that are updated within the cached control units. These counters reflect the load and behavior of specific disks, but they don't show where the I/O request came from. If several systems share a control unit, the total of all activity on the control unit is displayed, not just the I/O requests from the system that is doing the analysis.

The same issue also exists for systems that are running second-level: an analysis by the second-level system will include any I/O activity to the same control units from the first-level system.

**Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages will be shown for all load fields.



## FCX177, Cache Extended Functions Screen – CACHEXT

**Performance Toolkit Subcommand:** “CACHEXT” on page 146.

The Cache Extended Functions Screen (FCX177) shows performance data for disks connected to a cached control unit capable of handling extended functions. It provides information about the use made and the effectiveness of the cache. Devices that have been active during the last measuring interval are highlighted on the screen.

FCX177		CPU nnnn		SER nnnnn		Interval HH:MM:SS		Perf. Monitor																				
<--Device Descr.-->		Stg	C	D	D	<-----Rate/s----->			<-----Percent----->										<--Staging-->			DFW	BYP	Time				
Addr	Type	VOLSER	ID	C	W	L	ST	Cache	SCMBK	N-Seq	Read	Write	Seq	FW	Read	Tot	RdHt	Wrt	DFW	CFW	%DFW	%CFW	Noxm	%Seq	DeSt	Byps	+ICL	Outs
C50D	3390-9	ASP111	0607	A	A	-	00	35.0	22.4	27.1	2.1	5.9	83	60	52	100	100	..	100	0	93	9	87	0	0	0	0	
C004	3390-9	ATP032	0602	A	A	-	00	34.9	22.3	26.8	2.1	6.0	83	62	54	100	100	..	100	0	88	9	90	0	0	0	0	
C701	3390-9	AEP127	0609	A	A	-	00	34.7	22.3	26.8	2.3	5.6	84	59	52	100	100	..	100	0	93	10	91	0	0	0	0	
C707	3390-9	AVP133	0609	A	A	-	00	34.8	22.3	26.9	2.1	5.8	83	60	52	100	100	..	100	0	93	10	91	0	0	0	0	
CD05	3390-9	AEP215	060F	A	A	-	00	35.0	22.3	26.8	2.2	6.0	83	59	51	100	100	..	100	0	94	9	89	0	0	0	0	
CD0A	3390-9	AJP220	060F	A	A	-	00	35.0	22.3	27.0	2.2	5.8	83	62	54	100	100	..	100	0	87	9	88	0	0	0	0	
C207	3390-9	AOP063	0604	A	A	-	00	34.8	22.2	26.9	2.1	5.8	83	59	50	100	100	..	100	0	95	9	89	0	0	0	0	
C80A	3390-9	ANP150	060A	A	A	-	00	34.8	22.2	26.8	2.4	5.7	84	61	54	100	100	..	100	0	89	9	88	0	0	0	0	
C901	3390-9	AEP155	060B	A	A	-	00	34.8	22.2	26.9	2.0	5.9	83	59	51	100	100	..	100	0	94	10	89	0	0	0	0	
C90A	3390-9	ASP164	060B	A	A	-	00	34.8	22.2	26.6	2.3	5.9	83	61	53	100	100	..	100	0	90	10	90	0	0	0	0	
Select a device for I/O device details																												
Command ==>																												
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return																												

Figure 77. Layout of Cache Extended Functions Screen (FCX177)

The information shown is based on MONITOR and I/O domain data.

Figure 77 on page 569 shows a screen with performance information for 3390 disks that have been sorted in descending order of the total SCMBK I/O rate value, as indicated by the underscore characters above the Total SCMBK heading. Other possible sort sequences are indicated by dots above the corresponding columns. These sequences can be set using the SORT subcommand with the column heading (see “SORT” on page 295) or by positioning the cursor on the heading of the field that is to be sorted and then pressing Enter.

For more detailed information about the load of one specific disk, perform one of the following actions:

- Place the cursor on the Addr field of the device you are interested in and press Enter
- Enter the subcommand DEVICE *devno* on the command line.

### Field Descriptions:

*Device Descr.* provides some general information about the disks:

#### Addr

The disk device number.

#### Type

The device type and model.

For PAV alias devices of an IBM DASD subsystem, the device type is replaced by the string:

```
-> xxxx
```

where xxxx is the device number of the base device.

#### VOLSER

The volume label of the device.

#### Stg Ctlr ID

The storage controller subsystem identifier.

#### CAC

The caching status:

**A**  
Caching is activated.

**D**  
Caching is deactivated.

**P**  
Deactivation is pending.

#### **DFW**

The status of DASD fast write:

**A**  
DASD fast write is activated.

**D**  
DASD fast write is deactivated.

**P**  
DASD fast write deactivate is pending.

#### **DUL**

The status of dual copy:

**1**  
The disk is the primary device of a dual copy pair; dual copy is active.

**2**  
The disk is the secondary device of a dual pair; dual copy is active.

**D**  
Dual copy has been suspended.

**P**  
Dual copy deactivation is pending.

A dash indicates that dual copy is not active for the disk.

#### **ST**

The caching status flag in hexadecimal representation:

```
Bits 0-1
00 Caching is activated
01 Reserved
10 Deactivation is pending
11 Caching is deactivated

Bits 2-3
00 DASD fast write is allowed
01 Reserved
10 DASD fast write deactivation is pending
11 DASD fast write is deactivated

Bit 4: The primary device of a duplex pair, if set
Bit 5: The secondary device of a duplex pair, if set

Bits 6-7
00 A duplex pair is available
01 A duplex pair is pending (the copy to establish the duplex pair is in progress)
10 Failed duplex (the original is on the primary device)
11 Suspended duplex (the original is not on the primary device)
```

*Rate/s* shows cache activity, as the number of I/O operations per second:

#### **Total Cache**

The rate at which channel command word (CCW) chains are arriving at the device, per second. A channel program consists of one or more CCW chains, often delimited by Locate Record CCWs. Also, chains could be arriving from more than one partition.

**Total SCMBK**

The rate at which the reporting partition is sending channel programs to the device, per second. Also referred to as the Start Subchannel rate.

**Read N-Seq**

The arrival rate of read-normal chains or read Cache Fast Write (CFW) chains, per second.

**Read Seq**

The arrival rate of read-sequential chains, per second.

**Write FW**

The arrival rate of write DASD Fast Write (DFW) or write CFW chains, per second.

*Percent Hits* shows cache efficiency percentages:

**Read**

The percentage of arriving chains that are read chains.

**Tot**

The percentage of arriving chains satisfied without accessing a disk drive module (DDM).

**RdHt**

The percentage of arriving read chains satisfied without accessing a DDM.

**Wrt**

The percentage of arriving write DFW chains or write CFW chains satisfied without accessing a DDM.

**DFW**

The percentage of arriving write DFW chains satisfied without accessing a DDM.

**CFW**

The percentage of arriving write CFW chains satisfied without accessing a DDM.

*Percent /TotWrt* shows DASD and cache fast write usage percentages:

**/TotWrt %DFW**

The percentage of arriving write chains that are write DFW chains.

**/TotWrt %CFW**

The percentage of arriving write chains that are write CFW chains.

*Percent Staging* shows staging/destaging percentages, based on total cache I/O activity:

**Norm**

The percentage of non-sequential DASD to cache transfer operations.

**%Seq**

The percentage of sequential DASD to cache transfer operations.

**DeSt**

The percentage of cache to DASD transfer operations.

*Percent* shows the percentage of cache bypass operations, based on total cache I/O activity:

**DFW Byps**

The percentage of DASD fast write operations that were forced to bypass the cache and access DASD directly due to non-volatile storage constraints.

**BYP+ICL**

The percentage of I/O operations that bypassed caching voluntarily.

**Time Outs**

The number of timeouts that occurred while waiting for data from the control unit (status = -2). If no timeouts have occurred, this field is blank.

**Usage Notes:**

1. The analysis is based on counters that are updated within the cached control units. These counters reflect the load and behavior of specific disks, but they don't indicate where the I/O came from. If several systems share a device, the total of all activity on the disk is displayed, not just the I/O operations from the system that is doing the analysis.

The same problem exists for systems that are running second-level: an analysis by the second-level system will include any I/O activity to the same devices from the first-level system.

2. The I/O rates include alias I/O when the listed device on the CACHEEXT screen is a PAV or HyperPAV base device with non-zero alias activity.

**Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages will be shown for all load fields.

## FCX178, Minidisk Cache Storage Log Screen – MDCSTOR

**Performance Toolkit Subcommand:** “MDCSTOR” on page 234.

The Minidisk Cache Storage Log Screen (FCX178) shows minidisk cache storage utilization data by time:

FCX178	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor							
<----- Main Storage Frames -----> <----->											
Interval	<--Actual-->			Min	Max	Page	Steal				
End Time	Ideal	<2GB	>2GB	Set	Set	Del/s	Invokd/s	Bias	Ideal	Actu	
>>Mean>>	94948	42088	519166	0	655360	299	.240	1.00	16384	162	
11:23:58	102479	41600	380397	0	655360	244	.200	1.00	16384	162	
11:24:58	102522	41655	387876	0	655360	245	.200	1.00	16384	163	
11:25:58	102216	41689	394707	0	655360	271	.250	1.00	16384	161	
11:26:58	101257	41730	402826	0	655360	279	.217	1.00	16384	163	
11:27:58	101179	41743	413336	0	655360	322	.250	1.00	16384	162	
11:28:58	103179	41761	423091	0	655360	222	.183	1.00	16384	161	
11:29:58	100248	41769	429965	0	655360	340	.250	1.00	16384	163	
11:30:58	99880	41790	438826	0	655360	273	.233	1.00	16384	162	
11:31:58	99335	41874	447344	0	655360	300	.233	1.00	16384	162	
11:32:58	98832	41938	456691	0	655360	257	.200	1.00	16384	161	
11:33:58	98168	41974	464636	0	655360	300	.233	1.00	16384	163	
11:34:58	98062	41992	473121	0	655360	251	.217	1.00	16384	163	
11:35:58	97119	41998	482695	0	655360	308	.250	1.00	16384	162	
11:36:58	96477	41986	491686	0	655360	307	.250	1.00	16384	161	
11:37:58	95730	41963	500476	0	655360	300	.233	1.00	16384	162	
See MDCACHE display for minidisk cache activity log											
Command ==>											
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F10=Left	F11=Right	F12=Return				

Figure 78. Layout of Minidisk Cache Storage Log Screen (FCX178)

The information shown is based on CP monitor SYSTEM domain SAMPLE data.

The leftmost column contains the end time:

```
(format hh:mm:ss)
```

of the monitor interval for which the values were calculated. Note that only part of the performance variables can be shown on the 80 columns of an ordinary screen: use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window to the left or to the right for viewing the remaining data columns.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Main Storage Frames:

#### Ideal

The ideal number of main storage page frames in the minidisk cache (based on average age of a page in DPA, determined by arbiter)

#### Actual <2GB

The actual number of main storage page frames used for the minidisk cache below the 2GB line

#### Actual >2GB

The actual number of main storage page frames used for the minidisk cache above the 2GB line. The field is available only for z/VM systems running in 64-bit mode.

#### Min Set

The minimum number of main storage pages to be used for the minidisk cache

**Max Set**

The maximum number of main storage pages to be used for the minidisk cache

**Page Del/s**

The number of main storage pages deleted from cache per second

**Steal Invokd/s**

The number of times the steal function was invoked to steal main storage pages from cache per second

**Bias**

The bias for minidisk cache usage of real storage. A number less than 1 is a bias against minidisk cache; a number greater than 1 is a bias for minidisk cache.

*Expanded Storage Blocks:***Ideal**

The ideal number of expanded storage blocks in the minidisk cache (based on average age of an XSTORE block, as determined by arbiter)

**Actual**

The actual number of expanded storage blocks used for the minidisk cache

**Min Set**

The minimum number of expanded storage blocks to be used for the minidisk cache

**Max Set**

The maximum number of expanded storage blocks to be used for the minidisk cache

**Page Del/s**

The number of XSTORE pages deleted from cache per second

**Steal Invokd/s**

The number of times the steal function was invoked to steal expanded storage pages from cache per second

**Bias**

The bias for minidisk cache use of expanded storage. A number less than 1 is a bias against minidisk cache; a number greater than 1 is a bias for minidisk cache.

**Age**

The average age of paging XSTORE blocks

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX179, System Facilities Log Screen – SYSLOG

**Performance Toolkit Subcommand:** “SYSLOG” on page 312.

The System Facilities Log Screen (FCX179) shows a 'by time' log of system facilities utilization.

FCX179	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS						Perf. Monitor			
<----- Fast CCW Translations/s ----->												
<----- for DASD Devices -----> <-- for Network Devices -->												
Interval	<----- Command ----->					<-Transport>				<-UserExits->		
End Time	Done	Abort	Notelig	Total	Write	Done	Done	Abort	Notelig	Total	usec	%Busy
>>Mean>>	42593	58.39	.000	42652	.000	.000	.000	.000	.000	.000	82046k	136.7
15:05:00	40244	85.73	.000	40329	.000	.000	.000	.000	.000	.000	77255k	128.8
15:06:00	47829	28.53	.000	47858	.000	.000	.000	.000	.000	.000	112m	187.4
15:07:00	38633	6.283	.000	38639	.000	.000	.000	.000	.000	.000	63811k	106.4
15:08:00	34118	62.82	.000	34181	.000	.000	.000	.000	.000	.000	99663k	166.1
15:09:00	46434	51.93	.000	46485	.000	.000	.000	.000	.000	.000	53017k	88.4
15:10:00	45696	108.1	.000	45805	.000	.000	.000	.000	.000	.000	60776k	101.3
15:11:00	48882	51.70	.000	48934	.000	.000	.000	.000	.000	.000	89347k	148.9
15:12:00	42880	68.13	.000	42948	.000	.000	.000	.000	.000	.000	126m	210.2
15:13:00	45538	46.70	.000	45584	.000	.000	.000	.000	.000	.000	31613k	52.7
15:14:00	35680	73.90	.000	35754	.000	.000	.000	.000	.000	.000	106m	177.4

Command ==>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 79. Layout of System Facilities Log Screen (FCX179)

The information shown is based on SYSTEM domain SAMPLE data.

The leftmost column contains the end time (format *hh:mm:ss*) of the interval for which the values were calculated.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Descriptions:

#### Interval End Time

The end time of the sample interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

**Fast CCW Translations/s:** These fields show information on the activity of the fast CCW translation code. For z/VM V.4.1 and later, two different sets of CCW translation rates are provided:

- For DASD Devices and
- For Network Devices,

showing the CCW translation activity for CCW chains directed at disks and network devices respectively. Only CCW translations for DASD devices are reported on prior VM releases.

#### Done

Successfully translated CCWs per second

#### Abort

Aborted translation attempts per second

#### Notelig

Number of CCWs which were found to be not eligible for translation, per second

#### Total

Total number of CCWs handled per second

#### Transport Write

Number of times a write channel program for a DASD device was presented for HPF DCW translation.

**Transport Done**

Number of times HPF DCW translation processing completed successfully for a DASD device.

**User Exits**

Time spent processing user exit code.

**usec**

Microseconds of CPU time, per minute, spent in user exits during the by-time/average interval.

**%Busy**

Percent-busy running user exits in the by-time/average interval, where a value of 100 represents one physical processor's worth of CPU power.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.



## FCX180, System Configuration – SYSCONF

---

**Performance Toolkit Subcommand:** [“SYSCONF” on page 310.](#)

The System Configuration report (FCX180):

```

FCX180      CPU nnnn SER nnnnn Interval HH:MM:SS - HH:MM:SS      Perf. Monitor

System History for SYSTEM1
CP Level          z/VM Version 7.2.0, Service Level 0000 (64-bit)
CP Load Module    HCPA0TP1 Generated 2020/07/21 at 19:30:45
Last Termination Time 2020/07/21 at 19:32:07
Last Termination Code .....
Last IPL Time      2020/07/21 at 19:33:18
Time Zone Offset   -04:00:00 From Greenwich Mean Time
Checkpoint Volume  FCRES
Warmstart Volume   FCRES
PDR Volume         .....
Capacity Indic./Reason 100/000 (Running at nominal capacity)
Unified Resource Mgr  Not Available
Ensemble ID        .....
Channel subsystem ID 02
Enhanced CMF        Available
Polarization mode    Vertical
Topology check/changes 13989/1
Plant and Sequence Code 02 82F17 CEC 2964-NC9 MCI 757
Multithreading       Enabled

<----- Threads ----->
H/W Requested System Activated
Max Threads          1      MAX      1      2
CP core              1      MAX      1      1
IFL core              2      MAX      2      1
ICF core              2      MAX      1      1
ZIIP core             2      MAX      1      1

Server Time Protocol (STP) facility configuration
XRC_TEST enabled     No      XRC_OPTIONAL enabled     No
STP H/W feature installed Yes    STP H/W feature enabled     No
STP Timestamping enabled No      STP Timezone usage enabled     No
STP is active         No      STP is suspended           No
STP susp. message issued No
STP TOD clock offset +00:00:00.0000000000

Initial Status on 2020/07/22 at 03:17, Processor 2964-NC9
Total Conf Sthy Resvd Ded Shrd
Real Proc: Cap 492.0000 128 57 0 71
Sec. Proc: Cap 492.0000 71 71 0 0
Log. CP : CAF 456 26 26 0 0 0 26
Log. IFL : CAF 156 10 10 0 0 0 0
Log. ICF : CAF 500 2 2 0 0 0 0
Log. ZIIP: CAF 666 2 2 0 0 0 0

<----- Processor -----> Core/
Num Serial-Nr Type Status Thread
00 262F17 CP Master 00/0
02 262F17 CP Alternate 01/0
04 262F17 CP Alternate 02/0
06 262F17 CP Alternate 03/0
08 262F17 CP Alternate 04/0
0A 262F17 CP Alternate 05/0
0C 262F17 CP Alternate 06/0
0E 262F17 CP Alternate 07/0
:
3A 262F17 IFL Alternate 1D/0
3C 262F17 IFL Alternate 1E/0
3E 262F17 IFL Alternate 1F/0
40 262F17 IFL Alternate 20/0
42 262F17 IFL Alternate 21/0
44 262F17 IFL Alternate 22/0
46 262F17 IFL Alternate 23/0
48 262F17 ICF Alternate 24/0
4A 262F17 ICF Alternate 25/0
4C 262F17 ZIIP Alternate 26/0
4E 262F17 ZIIP Alternate 27/0
Processor Configuration Mode: VM

Changed System Status
Date Time Changed
07/22 03:31:19 SMT Configuration change 2 has started
07/22 03:31:19 Processor 35 (Core 1A, Thread 1) varied online
07/22 03:31:19 Processor 37 (Core 1B, Thread 1) varied online
07/22 03:31:19 Processor 39 (Core 1C, Thread 1) varied online
07/22 03:31:19 Processor 3B (Core 1D, Thread 1) varied online
07/22 03:31:19 Processor 3D (Core 1E, Thread 1) varied online
07/22 03:31:19 Processor 3F (Core 1F, Thread 1) varied online
07/22 03:31:19 Processor 41 (Core 20, Thread 1) varied online
07/22 03:31:19 Processor 43 (Core 21, Thread 1) varied online
07/22 03:31:19 Processor 45 (Core 22, Thread 1) varied online
07/22 03:31:19 Processor 47 (Core 23, Thread 1) varied online
07/22 03:31:19 SMT Configuration change 2 has finished

<----- Threads ----->
H/W Requested System Activated
CP core          1      MAX      1      1
IFL core          2      MAX      2      2
ICF core          2      MAX      1      1
ZIIP core         2      MAX      1      1

Command ==>
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

```

Figure 80. Layout of System Configuration report (FCX180)

The information shown is based on MONITOR domain SAMPLE data. It is initially provided by a MONITOR domain record when Performance Toolkit connects to the \*MONITOR service. If the program has been set up to collect the monitor data from another machine's disk file, the program tries to obtain the monitor

domain data from the first monitor sample in the disk file. The data cannot be shown if the required information is not found in that first sample.

### Field Descriptions:

*System History for system-ID:*

The string for *system-ID* is omitted in MONSCAN mode for z/VM systems where the system identification is not available.

One of the following strings might be appended in realtime monitoring mode if the system is not running first-level or is running in a logical partition:

```
(running 2nd level)
(running in LPAR)
(running 2nd level in LPAR)
```

The correct level is displayed if it is higher than second-level.

### CP Level

The CP level, with version, release, modification level, and service level.

### CP Load Module

The name of the CP load module used to IPL the z/VM system.

### Generated

The date and time of system generation of the CP load module.

### Last Termination Time

The date and time of last termination, if available. Dots are shown otherwise.

### Last Termination Code

The abend code of last termination, if available. Dots are shown otherwise.

### Last IPL Time

The date and time of the last IPL.

### Time Zone Offset

The time zone difference to Greenwich Mean Time.

### Checkpoint Volume

The disk label of the volume that contains the checkpoint cylinders.

### Warmstart Volume

The disk label of the volume that contains the warm start cylinders.

### PDR Volume

The disk label of the volume that contains the persistent data record of the SSI configuration.

### Capacity Indic./Reason

The capacity-adjustment indication (CAI) and capacity-change reason (CCR) values. When the CAI is 0, the indication is not reported. When the CAI is in the range 1-99, some amount of reduction is indicated. When the CAI is 100, the machine is operating at its normal capacity.

### Unified Resource Mgr

Indicates whether the system is managed by the Unified Resource Manager. Possible values are:

**Available**

**Disabled by CP**

**Disabled by SMAPI Server**

**Not Available**

**Unknown**

**Waiting for SMAPI Server**

### Ensemble ID

The Unified Resource Manager ensemble identifier.

### Channel subsystem ID

The channel subsystem (CSS) identifier (in hexadecimal).

**Enhanced CMF**

Indicates whether the Enhanced Channel Measurement Facility is installed and available.

**Polarization mode**

The polarization mode: Vertical or Horizontal.

**Topology check/changes**

The cumulative count of checks and detected changes of topology configuration.

**Plant and Sequence Code**

The hardware CPC identification (*AA BBBB CPC XXXX-YYY MCI ZZZ*), where:

**AA**

is the plant of manufacture.

**BBBB**

is the sequence code.

**XXXX**

is the machine type number.

**YYY**

is the model identifier.

**ZZZ**

is the model-capacity identifier.

When running on second-level z/VM systems (or higher), you might see dots in this field.

**Multithreading**

The initial multithreading mode. Possible values are:

**Enabled**

The processor or hardware configuration supports multithreading. Multithreading is enabled in the system configuration file.

**Facility is not installed**

The multithreading facility is not installed on the hardware or the logical partition, or the z/VM system is running second-level.

**No CP support**

The z/VM 6.3 system does not have CP APAR VM65586 installed.

**Disabled**

The processor or configuration supports multithreading, but multithreading is disabled.

Multithreading might be disabled for one of these reasons:

**As requested**

A MULTITHREADING statement in the system configuration file specified that multithreading is disabled on the system.

**No MULTITHREADING statement**

The MULTITHREADING statement is not specified in the system configuration file.

The following strings might also be shown as explanations if multithreading is disabled:

**Corequisite facilities *nn* were not available**

**ESA/390 LPAR**

**Horizontal polarization**

**IPL CPU address is too high**

**LINUX only LPAR with CP processors**

**SET SRM REBALANCE is active**

**SIGP failure**

*Threads:*

The following fields describe the initial multithreading configuration.

**H/W**

The maximum number of threads supported by the hardware, by core type.

**Requested**

The requested threads for this core type. The MAX keyword specifies that the maximum value supported by the machine is to be used.

**System**

The maximum number of threads supported by z/VM, by core type.

**Activated**

The number of activated threads, by core type.

*Server Time Protocol (STP) facility configuration:*

Shows the Server Time Protocol facility configuration settings.

**XRC\_TEST enabled**

Indicates whether XRC\_TEST is enabled in the configuration file.

**XRC\_OPTIONAL enabled**

Indicates whether XRC\_OPTIONAL is enabled in the configuration file.

**STP H/W feature installed**

Indicates whether the STP hardware feature is installed.

**STP H/W feature enabled**

Indicates whether the STP hardware feature is enabled.

**STP Timestamping enabled**

Indicates whether STP timestamping is enabled in the configuration file.

**STP Timezone usage enabled**

Indicates whether STP time zone usage is enabled in the configuration file.

**STP is active**

Indicates whether STP is active.

**STP is suspended**

Indicates whether STP is suspended.

**STP susp. message issued**

Indicates whether an STP suspension message has been issued.

**STP TOD clock offset**

Indicates the STP TOD clock offset in *+HH:MM:SS.ssssssssss* format. The fractional portion (sssssssss) provides a precision of 0.1 nanoseconds, which is sufficient to display any TOD clock value.

*{Initial|Refresh} Status on yyyy/mm/dd at hh:mm, Processor nnnn-nnn:*

Indicates when the status was obtained, that is, when monitoring was last activated (or reactivated), and the processor type and model. *Initial* Status is displayed the first time status information is shown. *Refresh* Status is displayed when a new configuration record arrives. The real, logical, and virtual CPU configurations, are displayed (when available), as follows.

**Real Proc**

The number of real processors on the machine and the capability of these processors.

**Sec. Proc**

The number of secondary processors on the machine and the capability of these processors.

**Virt Proc**

The number of virtual processors defined for a second-level guest z/VM system and the capability of these processors.

**Log. CP**

The number of logical processors of type CP in the partition and the capability adjustment factor (CAF) for this type.

**Log. IFL**

The number of logical processors of type IFL in the partition and the CAF for this type.

**Log. ICF**

The number of logical processors of type ICF in the partition and the CAF for this type.

**Log. ZIIP**

The number of logical processors of type ZIIP in the partition and the CAF for this type.

**Cap**

The real CPU capability of one of the CPUs in the configuration. This value is used as an indication of the capability of the CPU relative to the capability of other CPU models.

**CAF**

The processor capability adjustment factor for logical or virtual CPUs. This value specifies the amount of underlying level-1, level-2, or level-3 configuration capability that is allowed to be used for this configuration. The maximum level is 1000. Lower values indicate available CPU capabilities that are correspondingly lower.

**Total**

The total number of CPUs in the configuration. This value includes all CPUs in configured, standby, and reserved states.

**Conf**

The number of CPUs that are in a configured state.

**Stby**

The number of CPUs that are in a standby state.

**Resvd**

The number of CPUs that are in a reserved state.

**Ded**

The number of CPUs in the logical configuration that are dedicated CPUs.

**Shrd**

The number of CPUs in the logical configuration that are shared CPUs.

*Processor:*

Additional processor details follow for CPUs in a configured state.

**Num**

The logical processor ID in hexadecimal format.

**Serial-Nr**

The processor serial number.

**Type**

The processor type. Possible values are:

**CP**

Central Processor

**ICF**

Internal Coupling Facility

**IFL**

Integrated Facility for Linux

**ZIIP**

IBM Z Integrated Information Processor

**Status**

The processor status. Possible values are:

**Alternate****Dedicated****Master****Offline**

**Online**  
**Parked**

**Core/Thread**

The core ID and thread ID of the current processor in this format: *cc/t*, in hexadecimal.

**Processor Configuration Mode**

Displays the current LPAR (or virtual, for second-level) configuration mode:

**ESA390**  
**LINUX**  
**VM**

*Changed System Status:*

Shows the time and nature of any system status changes.

**Date**

The date of the status change (in *mm/dd* format).

**Time**

The time of the status change (in *hh:mm* format).

**Changed**

Describes how the status changed. For example:

```
Adding memory: Perm=xxxxG, Reconf=yyyyG (by uuuuuuuu)
Crypto facility for processor nn varied online/offline
Currently configured: Perm=xxxxG, Reconf=yyyyG
Memory add halted: Real=xxxxG (by uuuuuuuu)
Memory added: Real=xxxxG, SXS=yyyyG
Memory changed: StandBy= xxM, Reserved= yyM, Offline= zz
Memory converted: Real=xxxxG
Memory remove halted: Real=xxxxG (by uuuuuuuu)
Memory removed: Real=xxxxG
Processor nn (Core cc, Thread t) varied online/offline
Processor nn varied online/offline
Removing memory: Reconf=xxxxG (by uuuuuuuu)
SMT Configuration change n has finished
      <----- Threads ----->
      H/W Requested System Activated
CP core      1      MAX      1      1
IFL core     2      MAX      2      1
ICF core     2      MAX      1      1
ZIIP core    2      MAX      1      1
SMT Configuration change n has started
System polarization mode changed to xxxxxx (Detected)
System Topology changed. Check/Change count is xxx/yyy
The capacity has been changed: CAI=xxx, CCR=yyy
```

**Messages:**

These messages could appear in the Changed System Status area:

**Not applicable in Trend Scan mode**

is displayed when viewing data using the TRNDSCAN command. Because the corresponding change status events are not written to the history files, they are not available for redisplay.

**... (oldest lines not displayed) ...**

is displayed when the count of the change status lines exceeds the REDISP number (set using the FC MONCOLL REDISP command).

**\*\*\* Status refreshed \*\*\***

is displayed when the new configuration record for this particular report has been received.

**Usage Note:**

The phrase **Real=xxxxxG** in the Changed System Status area represents the amount of main storage that was added, removed, or converted as the result of a SET STORAGE command.

**Effect of AVERAGE and INTERIM subcommands:**

None.



## FCX182, General User Communication Log Screen – UCOMMLOG

**Performance Toolkit Subcommand:** [“UCOMMLOG” on page 332.](#)

The General User Communication Log Screen (FCX182), which shows 'by time' information on the overall use made of IUCV and VMCF communication facilities, will be shown if you enter the UCOMMLOG subcommand with no *userid*:

**Note:** If you specify UCOMMLOG *userid*, you will instead see [“FCX167, User Communication Log Screen – UCOMMLOG” on page 547.](#)

FCX182	CPU	nnnn	SER	nnnnn	Interval HH:MM:SS - HH:MM:SS				Perf. Monitor			
Interval	Total	<- IUCV Msg/s -->				<- VMCF Msg/s -->			<-- IUCV in Q -->			VMCF
End Time	Msg/s	ISend	IRecv	IFail	VSend	VRecv	VFail	SendQ	RecvQ	ReplQ	Queue	
>>Mean>>	253.7	110.5	122.8	.259	10.00	10.06	.000	.002	.102	.000	.001	
07:20:15	81.78	30.18	37.27	.203	6.973	7.157	.000	.003	.111	.000	.000	
07:25:15	97.35	37.16	45.33	.230	7.273	7.357	.000	.003	.117	.000	.000	
07:30:15	103.1	37.09	47.67	.210	9.070	9.080	.000	.003	.108	.000	.000	
07:35:15	118.9	46.44	55.31	.247	8.443	8.470	.000	.003	.103	.000	.000	
07:40:15	161.9	61.42	74.97	.263	12.63	12.66	.000	.002	.108	.000	.000	
07:45:15	156.5	56.28	72.43	.253	13.78	13.80	.000	.002	.102	.000	.000	
07:50:15	200.1	80.62	94.29	.280	12.43	12.46	.000	.002	.091	.000	.000	
07:55:15	197.0	78.07	92.78	.267	12.93	12.96	.000	.002	.097	.000	.000	
08:00:15	269.1	105.2	127.2	.303	18.22	18.26	.000	.002	.104	.000	.000	
08:05:15	319.5	128.6	152.6	.377	18.88	19.08	.000	.001	.099	.000	.001	
08:10:15	287.9	108.6	137.2	.330	20.84	20.86	.000	.001	.089	.000	.001	
08:15:15	401.1	157.2	192.3	.417	25.54	25.57	.000	.002	.080	.000	.001	
08:20:15	766.9	340.5	373.7	.380	26.12	26.21	.000	.001	.078	.000	.005	
08:25:15	565.4	244.3	272.7	.420	23.99	24.08	.000	.001	.079	.000	.007	
08:30:15	379.3	154.9	180.2	.400	21.90	21.91	.000	.001	.074	.000	.000	
08:35:15	613.2	264.6	297.2	.423	25.53	25.57	.000	.001	.086	.000	.000	
Command ==> _												
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return							

Figure 81. Layout of General User Communication Log Screen (FCX182)

The information shown is based on USER domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'Userid' field (and the 'Nr of users' field which is omitted from this screen), for the message target and SVM fields, and for the total and maximum IUCV connections fields which do not make much sense for the system communications log, all data columns are identical to the ones of the 'User Communication' (UCOMM) display. Please see [“FCX132, User Communication Screen – UCOMM” on page 484](#) for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX183, DASD Performance Log Screen – DASDLOG

**Performance Toolkit Subcommand:** “DASDLOG” on page 166.

The DASD Performance Log Screen (FCX183) shows overall performance data for all disks connected to the system, by time, with one entry per sample interval.

FCX183	CPU nnnn			SER nnnnn		Interval HH:MM:SS - HH:MM:SS					Perf. Monitor							
Interval	Mdisk	Pa-	<-Rate/s->		<----- Time (msec) ----->					Req.	<Percent>		SEEK	Recov	<-Throttle->			
End Time	Links	ths	I/O	Avoid	Pend	Disc	Conn	Serv	Resp	CUWt	Qued	Busy	READ	Cyls	SSCH	Set/s	Dly/s	
>>Mean>>	.1	2.4	314	541	.1	.0	.3	.4	.4	.0	.00	0	40	271k	0	...	.0	
19:54:44	.1	2.4	110	1.1	.1	.0	.1	.2	.2	.0	.00	0	3	15	0	...	.0	
19:55:44	.1	2.4	111	1.3	.1	.0	.1	.2	.2	.0	.00	0	4	18	0	...	.0	
19:56:44	.1	2.4	110	.1	.1	.0	.1	.2	.2	.0	.00	0	0	11	0	...	.0	
19:57:44	.1	2.4	110	.2	.1	.0	.1	.2	.2	.0	.00	0	0	11	0	...	.0	
19:58:44	.1	2.4	110	.2	.1	.0	.1	.2	.2	.0	.00	0	0	11	0	...	.0	
19:59:44	.1	2.4	111	1.3	.1	.0	.1	.2	.2	.0	.00	0	4	20	0	...	.0	
20:00:44	.1	2.4	116	1.3	.1	.0	.1	.2	.2	.0	.00	0	2	11	0	...	.0	
20:01:44	.1	2.4	116	451	.1	.0	.1	.2	.2	.0	.00	0	0	128	0	...	.0	
20:02:44	.1	2.4	113	234	.2	.0	.1	.3	.3	.0	.00	0	3	118	0	...	.0	
20:03:44	.1	2.4	123	1097	.2	.0	.1	.3	.3	.0	.00	0	2	189	0	...	.0	
20:04:44	.1	2.4	115	630	.2	.0	.1	.3	.3	.0	.00	0	0	178	0	...	.0	
20:05:44	.1	2.4	115	583	.2	.0	.1	.3	.3	.0	.00	0	0	152	0	...	.0	
20:06:44	.1	2.4	117	445	.1	.0	.1	.2	.2	.0	.00	0	0	59	0	...	.0	
20:07:44	.1	2.4	172	543	.1	.0	.2	.3	.3	.0	.00	0	11	53	0	...	.0	
20:08:44	.1	2.4	171	238	.1	.0	.2	.3	.3	.0	.00	0	0	22	0	...	.0	
20:09:44	.1	2.4	214	87.0	.1	.0	.2	.3	.3	.0	.00	0	33	66	0	...	.0	
20:10:44	.1	2.4	429	403	.2	.0	.4	.6	.6	.0	.00	0	4	17	0	...	.0	
20:11:44	.1	2.4	117	56.6	.1	.0	.1	.2	.2	.0	.00	0	2	36	0	...	.0	
20:12:44	.1	2.4	111	105	.1	.0	.1	.2	.2	.0	.00	0	4	26	0	...	.0	
20:13:44	.1	2.4	161	83.2	.1	.0	.2	.3	.3	.0	.00	0	0	30	0	...	.0	
20:14:44	.1	2.4	317	209	.2	.0	.2	.4	.4	.0	.00	0	29	221	0	...	.0	
20:15:44	.1	2.4	1090	1.4	.0	.0	.1	.1	.1	.0	.00	0	98	5	0	...	.0	
20:16:44	.1	2.4	805	4130	.1	.0	.2	.3	.3	.0	.00	0	96	20	0	...	.0	
20:17:44	.1	2.4	1213	3843	.1	.0	.3	.4	.4	.0	.00	0	70	261k	0	...	.0	
20:18:43	.1	2.4	875	2737	.1	.0	.3	.4	.4	.0	.00	0	74	46	0	...	.0	
20:19:44	.1	2.4	1258	3400	.1	.0	.5	.6	.6	.0	.00	0	55	169	0	...	.0	
20:20:44	.2	2.4	604	379	.2	.0	.7	.9	.9	.0	.00	0	3	31	0	...	.0	
20:21:43	.2	2.4	875	625	.2	.0	.5	.7	.7	.0	.00	0	1	21	0	...	.0	
20:22:44	.2	2.4	937	960	.2	.0	.5	.7	.7	.0	.00	0	2	2182k	0	...	.0	
20:23:44	.3	2.4	1046	1477	.2	.0	.4	.6	.6	.0	.00	0	1	40	0	...	.0	
Command ==>																		
F1=Help		F4=Top		F5=Bot		F7=Bkwd		F8=Fwd		F12=Return								

Figure 82. Layout of DASD Performance Log Screen (FCX183)

The information shown is based on MONITOR and I/O domain SAMPLE data.

The leftmost column contains the end time (in *hh:mm:ss* format) of the monitor interval for which the values were calculated.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'Addr', 'Type' and 'Label/ID' fields, all data columns are identical to the ones of the 'I/O Device' (DEVICE) display. Please see “FCX110, I/O Device Details Screen – DEVICE” on page 420 for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX184, Cache Extended Functions Log Screen – CACHELOG

**Performance Toolkit Subcommand:** "CACHELOG" on page 145.

The Cache Extended Functions Log Screen (FCX184) shows "by time" information on the cache activity and effectiveness for the selected disk. Note that the disk must be connected to a cached control unit which has previously been selected for benchmarking.

FCX184	CPU nnnn		SER nnnnn	Interval HH:MM:SS - HH:MM:SS					Perf. Monitor					
Cache Ext. Function Log for Device nnnn														
<Dev. Descr.>		Stg	C	D	D	<----- Rate/s			<----->		<----->			
Interval		Ctlr	A	F	U	Total	Total	Read	Read	Write				
End Time	Type	VOLSER	ID	C	W	L	ST	Cache	SCMBK	N-Seq	Seq	FW	Read	Tot
>>Mean>>	3390-2	SPE704	0006	A	A	-	00	23.4	23.4	14.9	.0	8.5	64	85
09:06:35	3390-2	SPE704	0006	A	A	-	00	32.9	32.9	22.1	.0	10.8	67	87
09:11:35	3390-2	SPE704	0006	A	A	-	00	33.6	33.6	20.1	.0	13.5	60	87
09:16:35	3390-2	SPE704	0006	A	A	-	00	40.7	40.7	23.0	.0	17.7	57	87
09:21:35	3390-2	SPE704	0006	A	A	-	00	24.6	24.7	15.4	.0	9.3	62	84
09:26:35	3390-2	SPE704	0006	A	A	-	00	29.2	29.2	16.0	.0	13.1	55	85
09:31:35	3390-2	SPE704	0006	A	A	-	00	32.9	32.9	22.0	.0	10.9	67	90
09:36:35	3390-2	SPE704	0006	A	A	-	00	11.6	11.6	8.9	.0	2.7	77	85
09:41:35	3390-2	SPE704	0006	A	A	-	00	16.0	16.0	9.4	.0	6.6	59	80
09:46:35	3390-2	SPE704	0006	A	A	-	00	13.5	13.5	9.2	.0	4.3	68	86
09:51:35	3390-2	SPE704	0006	A	A	-	00	23.2	23.4	13.5	.0	9.7	58	82
09:56:35	33902	SPE704	0006	A	A	-	00	26.1	25.9	12.8	.0	13.4	49	80
10:01:37	33902	SPE704	0006	A	A	-	00	27.9	28.0	17.5	.0	10.4	63	82
10:06:35	33902	SPE704	0006	A	A	-	00	23.2	23.3	14.1	.0	9.1	61	92
10:11:35	33902	SPE704	0006	A	A	-	00	12.7	12.6	10.0	.0	2.7	79	80
Command ==>														
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return									

Figure 83. Layout of Cache Extended Functions Log Screen (FCX184)

The information shown is based on MONITOR and I/O domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the Interval End Time column.

### Field Descriptions:

*Device Descr.* provides some general information about the disks:

#### Interval End Time

The end time of the monitor interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### Type

The device type and model.

For PAV alias devices of an IBM DASD subsystem, the device type is replaced by the string:

```
-> xxxx
```

where xxxx is the device number of the base device.

#### VOLSER

The volume label of the device.

#### Stg Ctlr ID

The storage controller subsystem identifier.

#### CAC

The caching status, as follows:

- A** Caching is activated.
- D** Caching is deactivated.
- P** Deactivation is pending.

#### **DFW**

The status of DASD fast write, as follows:

- A** DASD fast write is activated.
- D** DASD fast write is deactivated.
- P** DASD fast write deactivate is pending.

#### **DUL**

The status of dual copy, as follows:

- 1** The disk is the primary device of a dual copy pair; dual copy is active.
- 2** The disk is the secondary device of a dual pair; dual copy is active.
- D** Dual copy has been suspended.
- P** Dual copy deactivation is pending.

A dash indicates that dual copy is not active for the disk.

#### **ST**

The caching status flag in hexadecimal representation, where:

```

Bits 0-1
'00' Caching is activated
'01' Reserved
'10' Deactivation is pending
'11' Caching is deactivated

Bits 2-3
'00' DASD fast write is allowed
'01' Reserved
'10' DASD fast write deactivation is pending
'11' DASD fast write is deactivated

Bit 4: The primary device of a duplex pair, if set
Bit 5: The secondary device of a duplex pair, if set

Bits 6-7
'00' A duplex pair is available
'01' A duplex pair is pending (the copy to establish the duplex pair is in progress)
'10' Failed duplex, the original is on the primary device
'11' Suspended duplex, the original is not on the primary device

```

*Rate/s* shows cache activity, as the number of I/O operations per second:

#### **Total Cache**

The total I/O rate for the disk as it is perceived by the cached control unit, that is, where multiple systems are connected to one control unit, the total I/O activity from *all* systems is shown.

#### **Total SCMBK**

The total I/O rate as indicated by subchannel measurement block data for the system that does the monitoring, that is, only the I/O activity for that system is included. This field includes the I/O that is

performed directly using the base device and the I/O that is performed against that base device by an alias device, if any.

**Read N-Seq**

The non-sequential read rate (read normal + read CFW).

**Read Seq**

The sequential read rate (read sequential).

**Write FW**

The total write fast write rate (CFW + DFW).

*Percent Hits* shows cache efficiency percentages:

**Read**

The percentage of cachable read operations, based on total I/O activity.

**Tot**

The percentage of total hits (read + CFW + DFW), based on the sum of all cachable read and write operations.

**RdHt**

The percentage of read hits, based on the sum of all cachable read operations.

**Wrt**

The percentage of write hits (DASD FW + cache FW), based on the sum of all DASD and cache fast write write operations.

**DFW**

The percentage of DASD fast write hits, based on the sum of all DASD fast write write operations.

**CFW**

The percentage of cache fast write hits, based on the sum of all cache fast write write operations.

*Percent /TotWrt* shows DASD and cache fast write usage percentages:

**/TotWrt %DFW**

Percentage of DASD fast write operations, based on total write activity

**/TotWrt %CFW**

Percentage of cache fast write operations, based on total write activity

*Percent Staging* shows staging/destaging percentages, based on total cache I/O activity:

**Norm**

The percentage of non-sequential DASD to cache transfer operations

**%Seq**

The percentage of sequential DASD to cache transfer operations

**DeSt**

The percentage of cache to DASD transfer operations

**Percent** of cache bypass operations, based on total cache I/O activity:

**DFW Byps**

The percentage of DASD fast write operations that were forced to bypass the cache and access DASD directly due to non-volatile storage constraints

**BYP+ICL**

The percentage of I/O operations that bypassed caching voluntarily

**T**

The timeout indicator. An asterisk (\*) appears in this field if timeouts have occurred. This field is left blank when no timeouts were found.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX185, I/O Configuration Changes Screen – IOCHANGE

**Performance Toolkit Subcommand:** “IOCHANGE” on page 201.

The I/O Configuration Changes Screen (FCX185) shows changes made to the initial I/O configuration after activating monitor data collection:

```
FCX185      CPU nnnn  SER nnnnn      Status  HH:MM:SS      Perf. Monitor

I/O Configuration Change Log
Log started at 07:21:26 on 2016/09/14
I/O event data initially enabled for ALL devices.
PCIF class event data initially enabled.
See MONSET display for later changes.

Date Time      Changed
09/16 07:03:09 Device 4113 attached to DICVSW2
09/16 07:03:09 Device F020 attached to DICVSW2
09/16 07:03:09 HIPER-BRDG device F020 activated for DTCVSW2 as F020
09/16 07:03:09 QDIO device 4113 activated for DTCVSW2 as 4113
09/16 07:03:09 HPF features have been changed for 010D 'EM'-'>'E'
09/16 07:03:09 Device 4113 attached to DICVSW2
09/16 07:03:12 HyperPAV pool 0 created for C700
09/16 07:03:12 PAV features have been changed for C700 NOPAV -> HYPAV

Command ===>
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return
```

Figure 84. Layout of I/O Configuration Changes Screen (FCX185)

The information shown is based on I/O domain EVENT data.

### Field Descriptions:

*I/O Configuration Change Log:*

#### Log started

Time and date when data collection was started.

#### I/O event data ...

The initial data collection setting for the I/O event domain.

I/O configuration changes can be logged only for periods during which I/O event data were collected for the affected device or devices.

#### PCIF class event data ...

The initial data collection setting for the PCIF class of the I/O event domain.

*Configuration Changes:* Shows date, time, and nature of any I/O configuration changes.

#### Date

Month and day (format: *mm/dd*) of the status change

#### Time

Time (format: *hh:mm:ss*) of the status change

#### Changed

Describes how the status changed, as follows:

- Device xxxx varied offline
- Device xxxx varied online
- Device xxxx attached to *userid*
- Device xxxx detached from *userid*
- Shared device xxxx attached to system
- Subchannel to device xxxx varied online

- Subchannel to device *xxxx* varied offline
- Measurement facility set on for device *xxxx*
- Measurement facility set off for device *xxxx*
- Device *xxxx* deleted
- Device *xxxx* modified
- Channel path *yy* deleted
- Throttling rate for device *xxxx* set to *n.nn/s*
- Throttling rate for device *xxxx* set off
- NIC *userid vxxx* for Virtual Network *userid lanname* is Up
- NIC *userid vxxx* for Virtual Network *userid lanname* is Down
- *qdio* device *xxxx* activated for *userid* as *vxxx*
- *qdio* device *xxxx* deactivated
- HPF features have been changed for *aaaa bb->cc*
- PAV features have been changed for *aaaa ddddd->eeee*
- HyperPAV pool *fffff* created for *aaaa*

where:

***xxxx***

Is the device number of affected I/O devices

***yy***

Is the channel path ID

***userid***

Is the user identification of virtual machines: to or from which an I/O device has been attached or detached; the owner of the specified NIC virtual card or attached QDIO device; the virtual network owner.

***qdio***

Is the type of network device: HiperSocket, QDIO, OSN, INMN, IEDN, HIPER-BRDG, HIPER-IEDN, or Unknown. Note: after deactivation of QDIO device, the IOCHANGE log contains three records about detached devices: *xxxx*, *xxxx+1*, and *xxxx+2*.

***vxxx***

Is a virtual device number in the current configuration of the specified *userid*.

***lanname***

Is a virtual network name.

***aaaa***

Is the SSID for which the HPF-installed features or PAV operating stage codes have been changed or a new HyperPAV pool has been created.

***bb cc***

Are the HPF-installed feature bit codes. The following codes are supported:

**D**

HPF feature has been disabled.

**E**

HPF feature has been enabled.

**M**

Multitrack operations have been enabled.

***dddd***

Is the previous PAV operating state code. The following codes are supported:

**NOPAV**

CU is operating in No PAV mode.

**PAV**

CU is operating in Basic PAV mode.

**HYPAV**

CU is operating in HyperPAV mode.

**eeee**

Is the current PAV operating state code. The following codes are supported:

**NOPAV**

CU is operating in No PAV mode.

**PAV**

CU is operating in Basic PAV mode.

**HYPAV**

CU is operating in HyperPAV mode.

**ffff**

Is the HyperPAV pool number.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.



# FCX186, Remote Basic Mode Screen – BASMODE

**Performance Toolkit Subcommand:** “BASMODE” on page 137.

The Remote Basic Mode Screen (FCX186) allows entering CP and CMS commands and can, therefore, be selected only by users which have received 'command' authorization on the remote system (CMD specification in the requestor's entry of file FCONRMT AUTHORIZ on the remote system).

```
FCX186          BASIC Mode Screen Copy of User OPERATNS          CHVM4

+63 08:48:13 FCXUSL317A User HEIS %CPU 63.6 exceeded threshold 30.0 for 5 min.
+64 08:58:13 FCXUSL317A User HEIS %CPU 82.5 exceeded threshold 30.0 for 5 min.
FCXAPP527I User FCONMNT connected on path 0003
FCOMMNT : MENU
FCOMMNT : BASM
FCOMMNT : CMS ID
cms id
OPERATNS AT CHVM4      VIA RSCS1      04/03/96 09:01:16 CES      WEDNESDAY
R;
FCOMMNT : CP Q STOR
cp q stor
STORAGE = 0015424K

Command ==> _
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 85. Layout of Remote Basic Mode Screen (FCX186)

**Description:**

The information shown on this screen is a simple copy of the data on the server machine's basic mode screen. You can:

- Enter CP commands, with the CP prefix
- Enter CMS commands, with the CMS prefix, and you can
- Directly view the result of these commands.

Several people may select the remote basic mode screen and enter commands concurrently.

**Restrictions:**

As mentioned above, what you see is a copy of the contents of the server machine's basic mode screen, and only line mode output can be intercepted and displayed.

- *Do not attempt to activate any full screen application.* Doing so will place the server machine in a full screen mode whose output cannot be intercepted and transmitted to you, nor can you pass back commands for controlling that full screen application.

You will have to logon to the server machine in order to exit from the full screen application if you have inadvertently started one.

- The server machine's screen size may be different from your own. If it is larger, you may have to use the forward/backward scroll keys for viewing all the data on the screen.
- Changes on the server machine's screen will not automatically be reflected on your copy, except immediately after entering a command.

Press the ENTER key without entering a command in order to force a screen refresh.

## FCX187, User-Defined Performance Data Logs Screen – FC DEFLOG

**Performance Toolkit Subcommand:** “FCONTROL DEFLOG” on page 21.

This screen can contain selected performance data fields from other standard displays. The specified fields will be copied at the end of each monitor sample interval.

Because the layout of the resulting screen is, by definition, determined by the user, any example only illustrates some of the possibilities. The following figure shows such an example:

```
FCX187      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS      Perf. Monitor

Sample log with some fields copied from the SYSTEM screen
Interval
End Time   SieEx SieIn  Subp-free-<2G Subp-free->2G Totl-free-<2G Totl-free->2G
08:40:44   3268 2934   6448    107 126837 2114   6448    107 126857 2114
08:41:44   3240 2873   6457    108 81690 1362   6457    108 81704 1362
08:42:44   3271 2886   6428    107 81772 1363   6428    107 81786 1363
08:43:44   3745 2949   6342    106 81685 1361   6342    106 81698 1362
08:44:44   3700 3035   6442    107 81842 1364   6442    107 81856 1364
08:45:44   3259 2959   6422    107 81640 1361   6422    107 81655 1361
08:46:44   3358 2874   6435    107 81751 1363   6435    107 81765 1363
08:47:44   3672 3003   6364    106 81691 1362   6364    106 81705 1362
08:48:44   3734 2949   6348    106 81804 1363   6348    106 81818 1364
08:49:44   3210 2881   6343    106 81569 1359   6343    106 81583 1360
08:50:44   3467 2889   6369    106 81717 1362   6369    106 81731 1362
08:51:44   3259 2959   6345    106 81581 1360   6345    106 81595 1360
08:52:44   3390 2897   6388    106 81661 1361   6388    106 81674 1361
08:53:44   3459 2968   6344    106 81618 1360   6344    106 81633 1361
08:54:44   3424 2886   6356    106 81685 1361   6356    106 81698 1362
08:55:44   3313 2928   6345    106 81624 1360   6345    106 81639 1361
User-defined screen - no HELP text available
Command ==>
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 86. A Sample User-Defined Performance Data Logs Screen (FCX187)

Use the FC DEFLOG subcommand to define:

- The name you want to give to the log
- Any of the header lines 1 to 3
- The fields to be inserted into the log, where each requires:
  - Start column and length for the field in your new log
  - Source screen name and source field location (line number and start column)
  - Optionally, a short description to be inserted into the corresponding columns of the last header line (header line 3).

The subcommands used for defining the above log have been included in the sample FCONX \$PROFILE initialization file delivered with the program, as follows:

```
FC DEFLOG MYLOG H1 Sample log with some fields copied from the SYSTEM screen
FC DEFLOG MYLOG DESC Sample DEFLOG - fields from the SYSTEM screen
FC DEFLOG MYLOG COL 12 LEN 5 COPY SYSTEM LINE 14 COL 34 NAME SieEx
FC DEFLOG MYLOG COL 18 LEN 5 COPY SYSTEM LINE 14 COL 75 NAME SieInter
FC DEFLOG MYLOG COL 25 LEN 13 COPY SYSTEM LINE 19 COL 26 NAME Subp-free-<2G
FC DEFLOG MYLOG COL 39 LEN 13 COPY SYSTEM LINE 20 COL 26 NAME Subp-free->2G
FC DEFLOG MYLOG COL 53 LEN 13 COPY SYSTEM LINE 19 COL 67 NAME Totl-free-<2G
FC DEFLOG MYLOG COL 67 LEN 13 COPY SYSTEM LINE 20 COL 67 NAME Totl-free->2G
```

You would select this sample log by entering the command MYLOG on the command line. See “FCONTROL DEFLOG” on page 21 for a detailed description of the command.

The leftmost column contains the end time:

```
(format hh:mm:ss)
```

of the monitor interval for which the values were calculated.

**Field Description:****Interval End Time**

The end time of the monitor interval for which the values were calculated. This is the only field that will *always* be inserted; you cannot overwrite it with other data.

**Usage Notes:****1. Current and average data:**

The fields in a user-defined log are displayed just as they are displayed in the source report from which they are retrieved. If the source report supports "current" and "average" views, the data that is displayed will be the value that is reflected in the view that is active at the time the data was gathered.

**2. Data availability:**

Data will only be available on a user-defined log if the report from which it is retrieved can be generated.

## FCX188, Multitasking Users Screen – MTUSERS

**Performance Toolkit Subcommand:** “MTUSERS” on page 240.

The Multitasking Users Screen (FCX188) shows information on CMS multitasking activity in virtual machines by user.

Information for users with some multitasking activity in the interval is shown highlighted on the screen.

FCX188	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS				Perf. Monitor			
Userid	<-----Threads----->		<-----Max.----->		<-----Max.----->		<-----Max.----->		<-----Max.----->	
	<-Creates-->	<-Deletes-->	<-Switch/s>	Threads	<Concurrent>	Blocked	Proc.	Thrds	POSIX	Nr of
	Cre/s	ms/Cre	Del/s	ms/Del	Slow	Fast			Create	Fails
>System<	.002	.000	.001	.005	.001	.008	9.7	2.8	12.7	.0
WRH	.030	.000	.010	.013	.010	.113	6.0	2	8	0
BITMAN	.023	.000	.013	.001	.013	.130	9.0	2	11	0
SQLMACHF	.007	.000	.003	.000	.003	.023	13.0	3	17	0
CALSERV	.000	...	.000	...	.000	.000	10.0	3	13	0
EDLSFS	.000	...	.000	...	.000	.000	10.0	3	13	0
EDLSFS1	.000	...	.000	...	.000	.000	10.0	3	13	0
EDLSFS2	.000	...	.000	...	.000	.000	10.0	3	13	0
EDLSFS3	.000	...	.000	...	.000	.000	10.0	3	13	0
EDLSFS4	.000	...	.000	...	.000	.000	10.0	3	13	0
EPLREF	.000	...	.000	...	.000	.000	10.0	3	13	0
FLDTEST	.000	...	.000	...	.000	.000	10.0	3	13	0
IPSU	.000	...	.000	...	.000	.000	11.0	3	15	0
KAHLON	.000	...	.000	...	.000	.000	6.0	2	8	0
MKC	.000	...	.000	...	.000	.000	6.0	2	8	0
SCHANCK	...	...	...	...	...	...	...	...	...	...

Select a user for user details  
Command ==>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return

Figure 87. Layout of Multitasking Users Screen (FCX188)

The information shown is based on APPLICATION domain SAMPLE data.

The sample shows a screen with the users sorted in descending sequence of the threads creation rate (indicated by the underscore characters above the Cre/s header). Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the SORT subcommand with the appropriate column heading as the sort argument (see “SORT” on page 295), or by positioning the cursor on the header of the field that is to be sorted and then pressing Enter.

The first entry is reserved for the overall system average of all multitasking users for which data could be collected, regardless of the current sorting sequence. It is identified by: >System<-userid.

Some additional lines with user class averages will be shown at the top when user classes have been defined (see “FCONTROL UCLASS” on page 89). The *userid* shown will then be the user class name.

For more detailed information about a specific user, do one of the following:

1. Place the cursor on the Userid field of the virtual machine you are interested in and press Enter
2. On the command line, enter the subcommand: USER *userid*

### Field Descriptions:

The following field descriptions apply to the MTUSERS and MTUSRLOG *userid* displays, with the exception of the Userid field, which is replaced by the Interval End Time field on the MTUSRLOG display.

#### Userid

The user identification of the virtual machine for which the line contains performance data.

*Threads*: Shows thread activity data.

#### Creates Cre/s

The thread creation rate.

**Creates ms/Cre**

The average elapsed time, in milliseconds, per thread creation.

**Deletes Del/s**

The thread deletion rate.

**Deletes ms/Del**

The average elapsed time, in milliseconds, per thread deletion.

**Switch/s slow**

The regular path switch rate.

**Switch/s fast**

The fast path switch rate.

**Threads Blocked**

The number of threads that are currently blocked.

*Max. Concurrent:* Shows some maximum values.

**Proc.**

The highest number of processes that were concurrently defined.

**Thrds**

The highest number of threads that were concurrently defined.

**POSIX Create Fails**

The number of times that a POSIX thread creation failed due to an attempt to exceed the maximum allowable POSIX processes.

**Nr of Users**

The number of users that contributed to the average shown on this line. Values are inserted only for lines with averages; this field is blank for user data lines.

**Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages or totals will be shown. All of a user's sessions during the selected period contribute to the averages shown. Note also that all of the activity fields (rates) are based on each user's connection time within the overall period, *not* the total elapsed time for the entire period.

## FCX189, Multitasking User Log Screen – MTUSRLOG

**Performance Toolkit Subcommand:** “MTUSRLOG” on page 241.

When requesting a Multitasking User Log Screen (FCX189) for a user, the *userid* must be that of a virtual machine which has previously been selected for benchmarking. It provides a 'by time' log of the user's multitasking activity.

FCX189	CPU nnnn	SER nnnnn	Interval	HH:MM:SS	- HH:MM:SS	Perf. Monitor
Multitasking Data Log for User BITMAN						
Interval	<----- Threads ----->					
End Time	<-Creates-->	<-Deletes-->	<-Switch/s>	Threads	<---Max.--->	POSIX
>>Mean>>	Cre/s	ms/Cre	Del/s	ms/Del	Slow	Fast
09:27:12	.023	.000	.013	.001	.013	.130
09:32:15	.007	.000	.007	.001	.007	.033
09:37:15	.007	.000	.007	.001	.007	.033
09:42:16	.163	.000	.143	.001	.143	1.485
09:47:16	.290	.000	.290	.001	.290	2.969
09:52:15	.113	.000	.113	.001	.113	1.268
09:57:16	.063	.000	.063	.001	.063	.693
10:02:16	.010	.000	.010	.002	.010	.193
10:07:10	.003	.000	.003	.002	.003	.071
10:12:10	1.229	.000	1.213	.000	1.213	12.56
10:17:10	84.49	.000	84.49	.000	84.49	622.6
10:22:11	5.347	.000	5.357	.000	5.357	45.32
10:27:10	2.171	.000	2.165	.006	2.165	21.70
10:32:10	3.464	.000	3.464	.000	3.464	31.97
10:37:11	2.611	.000	2.611	.010	2.611	23.55
10:42:10	.084	.000	.090	.001	.090	.735
10:47:15	.000	...	.000	...	.000	.000
Command ==> _						
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return						

Figure 88. Layout of Multitasking User Log Screen (FCX189)

The information shown is based on APPLICATION domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'Userid' field (and the 'Nr of users' field which is omitted from this screen), all data columns are identical to the ones of the 'Multitasking Users' (MTUSERS) display. Please see “FCX188, Multitasking Users Screen – MTUSERS” on page 596 for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX190, Byte File System File Activity Log Screen – BFSFILES

**Performance Toolkit Subcommand:** “BFSFILES” on page 139.

The Byte File System File Activity Log Screen (FCX190) shows information on the file and directory activity of byte file system servers by time (only shared file system servers with some BFS activity are included):

FCX190		CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS					Perf. Monitor			
		<----- File Activity/sec ----->										
Interval		<--Open-->							Check	Change	Change	Lo
End Time	Server	Create	New	Old	Read	Write	Close	Access	Mode	Owner	Ob	
>>Mean>>	SCHANCK3	0	0	0	0	0	0	0	0	0	0	
>>Mean>>	IPSU	0	.000	.001	.007	.001	.003	.000	.000	0		
>>Mean>>	SERVERV	0	.000	0	.000	.000	.000	0	0	0		
>>Mean>>	SERVER8	0	0	0	0	0	0	0	0	0		
>>Mean>>	VMDCEBFS	0	.000	.000	.002	.000	.000	.000	.000	.000		
>>Mean>>	SQLMACHF	0	0	0	0	0	0	0	0	0		
>>Mean>>	SCHANCK4	0	0	0	0	0	0	0	0	0		
09:06:48	SCHANCK3	0	0	0	0	0	0	0	0	0		
09:06:48	IPSU	0	.001	.002	.018	.001	.007	.000	.000	0		
09:06:48	SERVERV	0	0	0	0	0	0	0	0	0		
09:06:48	SERVER8	0	0	0	0	0	0	0	0	0		
09:06:48	VMDCEBFS	0	.000	.001	.001	.000	.001	0	0	0		
09:11:49	SCHANCK3	0	0	0	0	0	0	0	0	0		
09:11:49	IPSU	0	0	0	.000	.000	.000	0	0	0		
09:11:49	SERVERV	0	0	0	0	0	0	0	0	0		
09:11:49	SERVER8	0	0	0	0	0	0	0	0	0		
09:11:49	VMDCEBFS	0	0	0	.000	.002	.000	0	0	0		
09:16:50	SCHANCK3	0	0	0	0	0	0	0	0	0		
09:16:50	IPSU	0	0	0	0	0	0	0	0	0		
09:16:50	SERVERV	0	0	0	0	0	0	0	0	0		
09:16:50	SERVER8	0	0	0	0	0	0	0	0	0		
09:16:50	VMDCEBFS	0	.000	.000	.005	.000	.000	.000	0	0		
Command ==> _												
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F10=Left	F11=Right	F12=Return					

Figure 89. Layout of Byte File System File Activity Log Screen (FCX190)

The information shown is based on CP monitor APPLDATA domain SAMPLE data.

The leftmost column contains the end time:

```
(format hh:mm:ss)
```

of the monitor interval for which the values were calculated. Note that only part of the performance variables can be shown on the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window to the left or to the right for viewing the remaining data columns.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column. The relative activity shown is based on each server's actual connect time, not the overall elapsed time in the interval.

### Field Description:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### Server

The userid of the shared file system or CRR server to which the following figures apply

#### *File Activity/sec:*

The following fields show file activity rates by different categories.

##### **Create**

The total file create rate (sum of create regular, create character and create block special file requests)

##### **Open New**

The total 'open new file' request rate (read and write)

##### **Open Old**

The total 'open old file' request rate (read and write)

##### **Read**

The 'read file' request rate

##### **Write**

The 'write file' request rate

##### **Close**

The 'close file' request rate

##### **Check Access**

The 'check file accessibility' request rate (to check the accessibility of a byte file object)

##### **Change Mode**

The 'change mode' request rate (to change the mode associated with a byte file object)

##### **Change Owner**

The 'change owner' request rate (to change the owner of a byte file object (UID/GID))

##### **Lookup Object**

The 'lookup' request rate (to lookup a byte file object)

##### **Make Object**

The 'makecat' request rate (to make a byte file object)

##### **Rename Object**

The 'rename' request rate (to rename a byte file object)

#### *Directory Activity/s:*

The following fields show directory activity rates by different categories.

##### **Create**

The 'create directory' request rate

##### **Open**

The 'open directory' request rate

##### **Read**

The 'read directory' request rate (to read directory entries)

##### **Close**

The 'close directory' request rate

##### **Remove**

The 'remove directory' request rate

The string 'No data for server' will be inserted if no new monitor records were processed for a server (logged off or monitoring disabled).

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.



# FCX191, Byte File System PIPEs and Linking Activity Log Screen – BFSPIPES

**Performance Toolkit Subcommand:** “BFSPIPES” on page 141.

The Byte File System PIPEs and Linking Activity Log Screen (FCX191) shows information on the PIPEs and linking activity of byte file system servers by time (only shared file system servers with some BFS activity are included):

FCX191	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS					Perf. Monitor		
<----- PIPE Activity / s ----->										
Interval	Server	Create	Open	Read	Write	Close	Access	Update TimeSt	Status	<---- C Hard S
>>Mean>>	SCHANCK3	0	0	0	0	0	0	0	0	0
>>Mean>>	IPSU	0	0	0	0	0	0	0	0	0
>>Mean>>	SERVERV	0	0	0	0	0	0	0	0	0
>>Mean>>	SERVER8	0	0	0	0	0	0	0	0	0
>>Mean>>	VMDCEBFS	0	0	0	0	0	0	0	0	0
>>Mean>>	SQLMACHF	0	.000	0	.000	0	0	0	0	0
>>Mean>>	SCHANCK4	0	.000	0	.000	.000	0	0	0	0
09:06:48	SCHANCK3	0	0	0	0	0	0	0	0	0
09:06:48	IPSU	0	0	0	0	0	0	0	0	0
09:06:48	SERVERV	0	0	0	0	0	0	0	0	0
09:06:48	SERVER8	0	0	0	0	0	0	0	0	0
09:06:48	VMDCEBFS	0	0	0	0	0	0	0	0	0
09:11:49	SCHANCK3	0	0	0	0	0	0	0	0	0
09:11:49	IPSU	0	0	0	0	0	0	0	0	0
09:11:49	SERVERV	0	0	0	0	0	0	0	0	0
09:11:49	SERVER8	0	0	0	0	0	0	0	0	0
09:11:49	VMDCEBFS	0	0	0	0	0	0	0	0	0
09:16:50	SCHANCK3	0	0	0	0	0	0	0	0	0
09:16:50	IPSU	0	0	0	0	0	0	0	0	0
09:16:50	SERVERV	0	0	0	0	0	0	0	0	0
09:16:50	SERVER8	0	0	0	0	0	0	0	0	0
09:16:50	VMDCEBFS	0	0	0	0	0	0	0	0	0
Command	==>	_								
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F10=Left	F11=Right	F12=Return			

Figure 90. Layout of Byte File System PIPEs and Linking Activity Log Screen (FCX191)

The information shown is based on CP monitor APPLDATA domain SAMPLE data.

The leftmost column contains the end time:

(format hh:mm:ss)

of the monitor interval for which the values were calculated. Note that only part of the performance variables can be shown on the 80 columns of an ordinary screen: use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window to the left or to the right for viewing the remaining data columns.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column. The relative activity shown is based on each server's actual connect time, not the overall elapsed time in the interval.

## Field Description:

### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

**Server**

The userid of the shared file system or CRR server to which the following figures apply

### *PIPE Activity / s:*

The following fields show BFS PIPE activity in different areas.

#### **Create**

The 'create named pipe (FIFO)' request rate

#### **Open**

The sum of the 'pipe open for read' and 'pipe open for write' request rates (to open a named pipe for reading or writing)

#### **Read**

The 'pipe read' request rate (byte file FIFO file pool requested to read from a named pipe)

#### **Write**

The 'pipe write' request rate (byte file FIFO file pool requested to write to a named pipe)

#### **Close**

The 'pipe close' request rate (byte file FIFO file pool requested to close a named pipe)

#### **Access**

The 'pipe access' request rate (byte file file pool requested to verify the access authorization to a named pipe)

#### **Update TimeSt**

The 'pipe Utime' request rate (byte file file pool requested to update the timestamps associated with a named pipe)

#### **Status**

The 'pipe stat' request rate (byte file FIFO file pool requested to obtain current status information about a named pipe)

### *Linking Activity / s:*

#### *Create:*

#### **Hard**

The 'create link' request rate (byte file file pool requested to create a hard link)

#### **Symbol**

The 'create symbolic link' request rate (byte file file pool requested to create a symbolic link)

#### **Extern**

The 'create external link' request rate (byte file file pool requested to create an external link)

#### **Read**

The 'read link contents' request rate (byte file file pool requested to read the contents of a link)

#### **Unlink**

The 'unlink' request rate (byte file file pool requested to remove a byte file object)

#### **Unlinkd Cleanup**

The 'unlinked file cleanup' request rate (unlinked files removed during FILESERV START)

The string 'No data for server' will be inserted if no new monitor records were processed for a server (logged off or monitoring disabled).

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX192, Byte File System Locking Activity Log Screen – BFSLOCKS

**Performance Toolkit Subcommand:** “BFSLOCKS” on page 140.

The Byte File System Locking Activity Log Screen (FCX192) shows information on the locking activity of byte file system servers by time (only shared file system servers with some BFS activity are included). The information shown is based on CP monitor APPLDATA domain SAMPLE data.

FCX192		CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS				Perf. Monitor			
		<----->				Locking Activity / sec <----->					
Interval	Server	Lock	Byte Test	Lock UnLock	Waits	Retry	Exceed	File	Direct	Lock Conflict	
>>Mean>>	SCHANCK3	0	0	0	0	0	0	0	0	0	
>>Mean>>	IPSU	.003	0	0	0	0	0	0	0	0	
>>Mean>>	SERVERV	.000	0	0	0	0	0	0	0	0	
>>Mean>>	SERVER8	0	0	0	0	0	0	0	0	0	
>>Mean>>	VMDCEBFS	.000	.000	.000	0	0	0	0	0	0	
>>Mean>>	SQLMACHF	0	0	0	0	0	0	0	0	0	
>>Mean>>	SCHANCK4	0	0	0	0	0	0	0	0	0	
09:06:48	SCHANCK3	0	0	0	0	0	0	0	0	0	
09:06:48	IPSU	.009	0	0	0	0	0	0	0	0	
09:06:48	SERVERV	0	0	0	0	0	0	0	0	0	
09:06:48	SERVER8	0	0	0	0	0	0	0	0	0	
09:06:48	VMDCEBFS	0	0	0	0	0	0	0	0	0	
09:11:49	SCHANCK3	0	0	0	0	0	0	0	0	0	
09:11:49	IPSU	0	0	0	0	0	0	0	0	0	
09:11:49	SERVERV	0	0	0	0	0	0	0	0	0	
09:11:49	SERVER8	0	0	0	0	0	0	0	0	0	
09:11:49	VMDCEBFS	0	0	0	0	0	0	0	0	0	
09:16:50	SCHANCK3	0	0	0	0	0	0	0	0	0	
09:16:50	IPSU	0	0	0	0	0	0	0	0	0	
09:16:50	SERVERV	0	0	0	0	0	0	0	0	0	
09:16:50	SERVER8	0	0	0	0	0	0	0	0	0	
09:16:50	VMDCEBFS	.000	0	0	0	0	0	0	0	0	
Command ==>											
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return											

Figure 91. Layout of Byte File System Locking Activity Log Screen (FCX192)

The leftmost column contains the end time (format *hh:mm:ss*) of the monitor interval for which the values were calculated. Note that only part of the performance variables can be shown on the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window to the left or to the right for viewing the remaining data columns.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column. The relative activity shown is based on each server's actual connect time, not the overall elapsed time in the interval.

### Field Description:

#### Interval End Time

The end time of the monitor interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### Server

The userid of the shared file system or CRR server to which the following figures apply

*Locking Activity / sec:* The following fields show locking activity and lock conflict rates.

*Byte Lock:*

#### Lock

The 'lock byte' request rate (to lock a byte range)

**Test**

The 'test locked bytes' request rate (to test for byte range locks held on a specific byte range)

**Unlock**

The 'unlock bytes' request rate (to unlock a byte range)

**Waits**

The 'byte range lock wait' rate (lock request waiting before being awarded the requested lock)

**Retry**

The 'logical lock retries' rate (retry attempted to obtain the BFS requests logical lock)

**Retry Exceed**

The 'logical lock retries exceeded' rate (logical lock denied due to the lock retry count being exceeded)

**Lock Conflicts/s:****File**

The 'file logical lock conflicts' rate (request for a lock, unlock, or close on a file for serializing byte lock/unlock and file closes was denied or waited for because an implicit lock had already been created on that file, where the request for the lock was also from an implicit request)

**Direct**

The 'directory creation/deletion logical lock conflicts' rate (request for a lock on an object - file, directory, link or symbolic link - to be created or deleted was denied or waited for because an implicit lock was already held on the object in a byte file system)

**TokMgr**

The 'token manager logical lock conflicts' rate (token manager requested a WRITE VNODE lock but had to wait for the lock because the implicit lock was already held)

**NAMECAT**

The 'NAMECAT unallocated lock conflicts' rate (request for a lock on an unallocated NAMECAT row was denied because the implicit lock was already held on the row)

**Global**

The 'global storage logical lock conflicts' rate (request for a lock on the object was denied or waited for because an implicit lock was already held on the object)

*Token Callback:* Tokens control shared resources and data caching for byte file system clients/users, much like locks. When a byte file request requires a token that is held by another user, the requestor must wait until the client machine that holds the token returns it (responds to a callback of that token).

**Rate/sec**

The rate of callbacks due to token conflicts

**ms/Callbk**

The average callback wait time: The time spent waiting for callbacks of tokens

**T-Out Ret/s**

The 'token callback timeout retry' rate (for retries of callbacks because of a delay of the holding client machine to respond to the token callback request)

**Request Retry/s**

The 'token callback requestor retry' rate (Requestor retries because of extended delays in call back response. This occurs when it is necessary to give up waiting for a normal callback completion because of exceeding the retry limit for callback retries.

The string 'No data for server' will be inserted if no new monitor records were processed for a server (logged off or monitoring disabled).

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX193, System Selection Menu Screen – SYSMENU

**Performance Toolkit Subcommand:** “SYSMENU” on page 313.

The System Selection Menu Screen (FCX193):

FCX193	CPU nnnn	SER nnnnn	(Intermediate Server)		VMSYSA
Systems Accessible using APPC/VM					
System	Resource	System	Resource	System	Resource
VMSYSA	FCXRES0A	VMSYSB	VMSYSB	TESTVM1	RESRC1
TESTVM2	RESRC2	TESTVM3	RESRC3	TESTVM4	RESRC4
TESTVM5	RESRC5	TESTVM6	RESRC6	TESTVM7	RESRC7
TESTVM8	RESRC8	TESTVM9	RESRC9	TESTVM10	RESRC10
Select the system to be monitored					
Command ==> _					
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return					

Figure 92. Layout of System Selection Menu Screen (FCX193)

See the "Connecting to an S&F Server for Remote Connections" section in the [z/VM: Performance Toolkit Guide](#) for more information on this screen.

## FCX194, Variable Correlation Screen – CORREL

**Performance Toolkit Subcommand:** “CORREL” on page 149.

The Variable Correlation Screen (FCX194) provides a quick overview showing how well any of the many performance variables saved in the HISTLOG and HISTSUM history files correlates with a selected base variable.

FCX194		Input File: 'yyyymmdd HISTLOG *'		Perf. Monitor	
<b>Correl. Coefficients for Base Variable CPU</b>		(% total CPU		, 165 Samp)	
Selection: From (not set)		Days : All days			
To (not set)		Hours: All hours			
Corre- lation	Variable Name	Variable Description	Corre- lation	Variable Name	Variable Description
<b>1.000</b>	%LogLd	% logical CPU	<b>1.000</b>	%US	% user CPU
<b>.997</b>	%EM	% emulation CPU	<b>.957</b>	VMStl	VMDBK steal rate
<b>.953</b>	%CP	% supervisor CPU	<b>.950</b>	XAl/s	XSTORE alloc/s
<b>.948</b>	XSStl/s	XSTORE steals/s	<b>.946</b>	XSDel/s	XSTORE delete/s
<b>.940</b>	%SY	% system CPU	<b>.938</b>	XDeal/s	XSTORE dealloc/s
<b>.928</b>	DIAG	DIAG instr./s	<b>.926</b>	Allo/s	Page alloc. rate
<b>.920</b>	SIE	SIE instr. rate	<b>.920</b>	PRIV	Inst. simul./s
<b>.911</b>	SIEIC	SIE intercept/s	<b>.907</b>	Dsptch	Users dispatched
<b>.899</b>	VMCF>C	VMCF rate to CP	<b>.899</b>	VMCF>V	VMCF rate to VM
<b>.899</b>	VMCF_T	Total VMCF rate	<b>.898</b>	VRecv	User VMCF recv/s
<b>.898</b>	VSend	User VMCF send/s	<b>.896</b>	NT-Q	Non-triv. queue
<b>.896</b>	PGOUT/s	PAGOUT rate/s	<b>.895</b>	Scan1	Scan 1 rate
<b>.895</b>	MDUS	MDC insert users	<b>.892</b>	TR-Q	Trivial queue
<b>.892</b>	ExtInt	Ext. interrupt/s	<b>.892</b>	XPg/s	XSTORE page rate
Select new base variable to recompute					
Command ==>					
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return					

Figure 93. Layout of Variable Correlation Screen (FCX194)

A good correlation with the base variable is indicated by highlighting the entry. The correlation coefficient is shown in:

### Yellow

For correlation coefficients of 0.9 to 1.0.

### White

For correlation coefficients of 0.8 to 0.9.

### Green

For all other values.

The sample shows a screen with correlations for base variable CPU (total CPU load), the default. You can select any other base variable by performing one of these actions"

1. Specifying the variable name with the CORREL subcommand when selecting the screen.
2. Moving the cursor to any variable on the correlation display and then pressing ENTER.

### Field Descriptions:

*General Header Lines:*

#### Correl. Coefficients for Base Variable CPU

This line shows which variable was used as a "base" variable against which the correlation coefficients for all other variables were calculated.

#### Selection

Shows the criteria applied when selecting performance data from the history file. Record selection can be based on an overall period (From and To specifications) and can be restricted to a smaller subset by selecting only specific days and/or hours.

**From**

Shows the start date or time of a selected period, or (not set) when no start date and time have been specified.

**To**

Shows the end date or time of a selected period, or (not set) when no end date and time have been specified.

**Days**

Shows for which weekdays history log records have been processed.

**Hours**

Shows for which hours history log records have been processed.

**Correlation**

The correlation coefficient for the variable.

The correlation will be shown as a number between -1.0 and +1.0, with the following exceptions:

**\*ZERO\***

Indicates that all of the values processed were 0.

**\*CONS\***

Indicates that all of the values were the same, that is, the variable was a constant during the selected period.

Dots indicate that an insufficient number of valid measurements was found to allow calculation of a meaningful correlation.

**Variable Name**

The variable name, as described for the REDHIST display. (See [“REDHIST” on page 268.](#))

A question mark in front of the variable name indicates correlation factors of questionable validity because the number of samples for the variable was less than 1/4 of the number of samples for the base variable, as indicated in the header line.

**Variable Description**

A short description of the variable. See the corresponding paragraphs for the REDHIST display ([“REDHIST” on page 268](#)) for a more detailed description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.



## FCX195, Extended History Data Log Screen – REDHIST

**Performance Toolkit Subcommand:** “REDHIST” on page 268.

The Extended History Data Log Screen is displayed after entering the REDHIST subcommand for one of the extended HISTLOG or HISTSUM history data log files. It allows browsing through these files. The current day's file is shown by default.

FCX195		Input File: '20200615 HISTLOG *'								Perf. Monitor			
Date	Time		<----- CPU Load (%) ----->										
yyyy/mm/dd	hh:mm	#CPU	CPU	%US	%CP	%EM	%SY	%Spin	%WT	%LogLd	%VEC	%VEM	
2020/06/15	09:11	6.0	385	370	74.5	310	14.7	.2	215	384.8	..	..	
2020/06/15	09:16	6.0	366	348	74.4	292	17.6	.2	234	366.0	..	..	
2020/06/15	09:21	6.0	365	346	77.7	287	18.2	.3	235	364.8	..	..	
2020/06/15	09:26	6.0	360	342	77.7	282	17.9	.3	240	359.7	..	..	
2020/06/15	09:31	6.0	441	424	73.7	368	17.3	.2	159	441.5	..	..	
2020/06/15	09:36	6.0	314	299	66.6	248	15.2	.2	286	314.2	..	..	
2020/06/15	09:41	6.0	312	296	68.8	243	16.2	.2	288	312.1	..	..	
2020/06/15	09:46	6.0	326	309	73.7	252	16.3	.2	274	325.6	..	..	
2020/06/15	09:51	6.0	345	329	74.1	271	15.9	.2	255	345.4	..	..	
2020/06/15	09:56	6.0	294	278	64.8	229	15.2	.2	306	293.6	..	..	
2020/06/15	10:01	6.0	345	328	69.8	275	16.8	.2	255	344.8	..	..	
2020/06/15	10:06	6.0	412	394	68.4	343	17.5	.2	188	411.8	..	..	
2020/06/15	10:11	6.0	441	425	72.2	368	15.9	.2	159	440.8	..	..	
2020/06/15	10:16	6.0	395	377	84.2	311	18.3	.2	205	395.1	..	..	
2020/06/15	10:21	6.0	414	396	75.4	338	18.3	.2	186	413.9	..	..	
2020/06/15	10:26	6.0	445	426	81.0	364	18.8	.3	155	444.8	..	..	
2020/06/15	10:31	6.0	436	417	83.4	353	18.7	.3	164	436.1	..	..	
2020/06/15	10:36	6.0	349	333	67.4	282	16.2	.2	251	349.4	..	..	
2020/06/15	10:41	6.0	312	296	64.4	247	15.4	.2	288	311.9	..	..	
Select new base variable with cursor and press ENTER to recompute													
Command ==>													
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return													

Figure 94. Layout of Extended History Data Log Screen (FCX195)

When viewing detailed HISTLOG files, the numbers shown generally represent average values calculated for the monitor sample interval that ended at the time specified in the Time field, but in some cases they represent values as they were sampled at the end of the interval.

When viewing condensed HISTSUM files, all numbers are averages of the original sample interval values shown in the detailed HISTLOG files.

**Note:** This display is a simple browser for the extended history log files. It shows the performance variables in the sequence in which they were stored in these files. While the performance variables were originally grouped so as to keep variables pertaining to the same general area (DASD performance, for example) together, new fields added by future releases of Performance Toolkit will be appended to the end of the HISTLOG records so as not to disturb existing applications that rely on fixed offsets for locating specific variables.

### Field Descriptions:

#### Date

The date when a row's performance data were collected. A complete date is shown, in yyyy/mm/dd format, when viewing the leftmost part of the history data, while an abbreviated format (mm/dd) is shown after shifting to the right to retain more space for displaying the performance data fields.

#### Time

The time when a row's performance data were collected. The time is always shown in hh:mm:ss format, where the time stamp is:

- The end time of a data collection interval in the case of detailed history log files (file type HISTLOGn)
- The end time of the first data collect interval within a one hour period in the case of a summary history log file (file type HISTSUM) where one record contains the average values calculated for all data collection intervals within that hour.

CPU Load (%):

**#CPU**

The number of processors that contributed to the following overall CPU figures.

**CPU**

The percentage of total elapsed time that processors were busy (the sum of busy percentages for all processors shown).

**%US**

The percentage of total elapsed time that processor time was attributed to specific users (includes emulation time and CP time, the sum of all processors shown).

**%CP**

The percentage of total elapsed time that processors were busy for CP (the sum of all processors shown).

**%EM**

The percentage of total elapsed time that processors were operating in emulation mode, doing work for users (the sum of all processors shown).

**%SY**

The percentage of total elapsed time that processor time was charged to the system (CP time not attributed to specific users, the sum of all processors shown).

**%Spin**

The percentage of total elapsed time that processors spent in spin loops (the sum of all processors shown).

**%WT**

The percentage of total elapsed time that processors were idle (the sum of all processors shown).

**%LogLd**

The percentage of *non-suspended* elapsed time that processors were busy. This value is identical to the total CPU load CPU for systems running first-level on a single-image CPU. When running in an LPAR, or second-level under z/VM, the logical load is higher because it is based on the sum of processor busy time plus *voluntary* wait time, and this sum is shorter than the actual elapsed time by the time that the processor was suspended (not allowed to work) by the PR/SM microcode or by the first-level CP.

**%VEC**

This field is deprecated (dots are displayed).

**%VEM**

This field is deprecated (dots are displayed).

**VECLD**

This field is deprecated (dots are displayed).

**PRIV**

The number of instructions simulated per second.

**DIAG**

The number of diagnose instructions executed per second.

**SIGP**

The rate at which this CPU is the target of SIGP External Call instructions.

**IO/S**

The number of start subchannel and resume subchannel executions per second.

**%Empt**

The percentage of elapsed time that the dispatch vector had no VMDBKs. This is the average of the percentages measured for all processors.

**#Usrs**

The average number of users in a dispatch vector when it was not empty. This is the average for all processors.

**#Mast**

The average number of users in the master-only dispatch vector when it was not empty.

**VMStl**

The total rate at which VMDBKs were stolen from any processor's dispatch vector, for dispatching on another processor.

**ToMast**

The total rate at which VMDBKs were moved to the master processor for processing of master-only work.

**AvEmp**

The average number of seconds before the available list became empty.

**Dsptch**

The number of long paths through the dispatcher per second. This corresponds approximately to the rate at which users were dispatched for running.

**SIE**

The total Start Interpretive Execution (SIE) rate. It does not include instructions to run a virtual guest during guest SIE simulation and SIE instructions following a successful fastpath simulation.

**SIEIC**

The total SIE exits by interception rate (that is, SIE mode stopped because CP had to simulate an instruction). It does not include virtual guest SIE interceptions and interceptions handled by fastpath simulation.

**PgFlt**

The total host segment and guest page fault rate, not including any CP page faults.

**FPSim**

The number of fastpath instruction simulations per second.

**SpRds**

The number of spool reads per second.

**SpWrt**

The number of spool writes per second.

**CSCH**

The number of clear subchannel executions per second.

**HSCH**

The number of halt subchannel executions per second.

**ExtInt**

The number of external interrupts processed per second.

**SolInt**

The number of solicited interrupts received per second.

**UnSolI**

The number of unsolicited interrupts received per second.

**Users:****LOGN**

The number of users logged on.

**ACT**

The number of users found active during a monitor sample interval.

**VIO/s**

The total number of virtual I/O instructions executed per second.

**Real Storage and Paging:****PPAG**

The number of pageable pages in the dynamic page area.

**NonPag**

The number of non-pageable pages in the DPA.

**%ST**

The percentage of storage utilization, that is, the percentage of pageable storage occupied by the working sets of users in the dispatch and eligible lists.

**SHAR**

The number of page frames in use by shared segments.

**FPGS**

This field is deprecated (dots are displayed).

**%FR**

Percent of subpool free storage in use.

**V=R**

This field is deprecated (dots are displayed).

**VRFREE**

This field is deprecated (dots are displayed).

**FEx/s**

The number of frames taken per second to satisfy the free storage extend condition.

**FrTakn**

The number of requests for a free storage extend page which have not yet been satisfied.

**FrSub**

The number of pages in use for free storage subpools (number of DWs provided, converted to pages).

**FrSav**

The number of pages in use for save areas.

**Resrvd**

The number of reserved pages.

**Locked**

The number of pages locked by CP LOCK command.

**Avail**

The number of pages on the available list.

**MS-Life**

The estimated life of a real storage page, calculated from Little's law (pageable storage / pageout rate).

**PRd/s**

The page read rate per second, including directory, nucleus and warm start I/O, but no spool I/O.

**PWr/s**

The page write rate per second, including directory, nucleus and warm start I/O, but no spool I/O.

**PG/s**

The total page rate per second (sum of reads and writes).

**NSS/s**

The number of shared segment pages read per second.

**Guest/s**

The number of single page reads per second for virtual machines.

**Syst/s**

The number of single page reads per second for the system (system pages are always read as single pages).

**SPag/s**

The total rate of single page reads (sum of guest reads and system reads).

**Alo/s**

The page allocation rate (number of pages obtained from the available list, per second).

**#TW**

The number of tasks waiting for a page, per second.

**FrWt**

The number of tasks waiting for a frame.

**Scan1**

The number of times the demand scan ended successfully with the first scan per second.

**Scan2**

The number of times the demand scan ended successfully with the second scan per second.

**E-Scan**

The number of times the demand scan ended successfully with the emergency scan per second.

**ScFail**

The number of times the demand scan failed per second.

**ESCN**

The percentage of demand scans which did not complete after the first pass.

**BKRD**

The average blocking factor when reading pages from DASD.

**BKMG**

The average blocking factor when migrating pages out of XSTORE.

**BKST**

The average blocking factor for page steals, that is, when writing pages from main storage to DASD.

*Expanded Storage Data:***XSTPgBk**

The number of expanded storage blocks available for paging (CP partition size minus the space used for the minidisk cache).

**%XS**

The percentage of expanded storage CP partition blocks in use.

**XAl/s**

The expanded storage page block allocation rate per second.

**XPg/s**

The total expanded storage paging rate (sum of PGIN and PGOUT operations per second).

**XDeal/s**

The expanded storage page block deallocation rate per second.

**MigVis**

The number of times a shared system was visited by migration per second.

**PGIN/s**

The total number of PGIN operations per second (to copy the contents of expanded storage blocks to main storage).

**%FP\_PGIN**

The percentage of PGIN operations handled by fastpath logic.

**PGOUT/s**

The number of PGOUT operations per second, to copy pages from main storage to expanded storage blocks.

**Migr/s**

The number of pages migrated from expanded storage to DASD, per second.

**XAGE**

The average age of expanded storage paging blocks, in seconds.

**MAGE**

The average age of the expanded storage paging blocks which were migrated to DASD, in seconds.

**XSTLife**

The average age of expanded storage paging blocks, calculated from Little's law (number of XSTORE paging blocks divided by PGOUT rate).

#### *Minidisk Cache Data:*

**MnIdeal**

The ideal number of main storage page frames in the minidisk cache (based on average age of a page in the DPA, determined by arbiter).

**MnAct**

The actual number of main storage page frames used for the minidisk cache.

**MnMin**

The minimum number of main storage page frames to be used for the minidisk cache.

**MnMax**

The maximum number of main storage page frames to be used for the minidisk cache.

**MnDel/s**

The rate at which main storage pages were deleted from the minidisk cache, per second.

**MnStl/s**

The rate at which the steal function was invoked to steal main storage pages, per second.

**MnBias**

The bias for minidisk cache usage of real storage. A number less than 1 is a bias against minidisk cache; a number greater than 1 is a bias for minidisk cache.

**XSTAV**

The total number of expanded storage blocks available to CP, including the blocks used for the minidisk cache.

**XSIdeal**

The ideal number of expanded storage blocks in the minidisk cache (based on average age of an XSTORE block, as determined by arbiter).

**XSAct**

The actual number of expanded storage blocks used for the minidisk cache.

**MDNE**

The total number of pages used for the minidisk cache in both expanded *and* main storage.

**XSMin**

The minimum number of expanded storage blocks to be used for the minidisk cache.

**XSMax**

The maximum number of expanded storage blocks to be used for the minidisk cache.

**FairShr**

The fair share limit for the minidisk cache, that is, the maximum number of blocks any single user is allowed to insert into the minidisk cache per fair share interval.

**XSDel/s**

The rate at which expanded storage pages were deleted from the minidisk cache, per second.

**XSStl/s**

The rate at which the steal function was invoked to steal expanded storage pages from the minidisk cache, per second.

**XSBias**

The bias for minidisk cache use of expanded storage. A number less than 1 is a bias against minidisk cache; a number greater than 1 is a bias for minidisk cache.

**MDR/S**

The number of read requests to the minidisk cache per second.

**MDH/s**

The number of read requests to the minidisk cache, per second, where all the requested blocks were found in the cache, thereby avoiding an I/O operation.

**%MDH**

The percentage of read requests where all requested blocks were found in the cache.

**%MDRPH**

This field is deprecated.

**%MDRMS**

This field is deprecated.

**BkRd/s**

This field is deprecated.

**%MDR**

This field is deprecated.

**%Bk**

This field is deprecated.

**MDPIN/s**

The rate at which minidisk cache pages were moved from expanded storage to main storage.

**MDPOUT/s**

The rate at which minidisk cache pages were moved from main storage to expanded storage.

**MDAvAge**

The estimated average age, in seconds, of a minidisk cache block, based on Little's law (cache size divided by PGOUT rate).

**FSEx/s**

The fair share exceeded rate, that is, the rate that blocks could not be moved into the minidisk cache because their user's fair share limit was exceeded.

**BkIns/s**

Insertions per second into the "in transit" waiting queue, which keeps track of users requesting blocks that are already on the way to the minidisk cache.

**InvR/s**

The invalidation request rate, that is, the number of requests per second to invalidate minidisk cache blocks due to an I/O operation to a virtual device via a non-cachable I/O interface.

**InvB/s**

The number of blocks per second which were invalidated following an invalidation request.

**MDUS**

The number of users which had blocks inserted into the minidisk cache in the last fair share interval.

*Auxiliary Storage Data:***PgSlots**

The total number of slots allocated to the page area.

**%PGSL**

The percentage of allocated page slots which are in use.

**SpSlots**

The total number of slots allocated to the spool area.

**%SPSL**

The percentage of allocated spool slots which are in use.

**DuSlots**

The total number of slots allocated to the dump area.

**%DUSL**

The percentage of allocated dump slots which are in use.

**SpCrea**

The number of spool files created per monitor sample interval.

**SpCre/s**

The rate, per second, at which spool files were created.

**SpPurg**

The number of spool files purged per monitor sample interval.

**SpPur/s**

The rate, per second, at which spool files were purged.

**PgMLOAD**

The average 'modified device load' value for paging devices, in milliseconds. This is the I/O device service time, multiplied by the number of queued requests for that same device, a value upon which CP bases its decisions when selecting a paging device for output.

**SpMLOAD**

The average 'modified device load' value for spooling devices, in milliseconds. This is the I/O device service time, multiplied by the number of queued requests for that same device, a value upon which CP bases its decisions when selecting a spooling device for output.

*Scheduler Queue Data:***InQue**

The total number of VMDBKs found in the dispatch and eligible lists.

**Q0**

The number of class 0 (quick dispatch) users' VMDBKs found in the dispatch list.

**Q1**

The number of class 1 (interactive) users' VMDBKs found in the dispatch list.

**Qx**

The number of non-interactive users' VMDBKs found in the dispatch list (classes 0, 2 and 3).

**Q2**

The number of class 2 users' VMDBKs found in the dispatch list.

**Q3**

The number of class 3 users' VMDBKs found in the dispatch list.

**Q0L**

The number of loading class 0 (quick dispatch) users' VMDBKs found in the dispatch list.

**Q1L**

The number of loading class 1 (interactive) users' VMDBKs found in the dispatch list.

**QxL**

The number of loading non-interactive users' VMDBKs found in the dispatch list (classes 0, 2 and 3).

**Q2L**

The number of loading class 2 users' VMDBKs found in the dispatch list.

**Q3L**

The number of loading class 3 users' VMDBKs found in the dispatch list.

**LimLst**

The number of VMDBKs in the limit list (not allowed to run due to their maximum share settings).

**E1**

The number of class 1 (interactive) users' VMDBKs found in the eligible list.

**Ex**

The number of non-interactive users' VMDBKs found in the eligible list (classes 2 and 3).

**E2**

The number of class 2 users' VMDBKs found in the eligible list.

**E3**

The number of class 3 users' VMDBKs found in the eligible list.

**E1L**

The number of loading class 1 (interactive) users' VMDBKs found in the eligible list.

**ExL**

The number of loading non-interactive users' VMDBKs found in the eligible list (classes 2 and 3).

**E2L**

The number of loading class 2 users' VMDBKs found in the eligible list.



**E3L**

The number of loading class 3 users' VMDBKs found in the eligible list.

**C1ES**

The length of the class 1 elapsed time slice, in seconds. This is the elapsed time which a class 1 user is allowed to spend in the dispatch list, and which should allow 85% of all transactions to complete.

**AbsShr**

The sum of the absolute shares of all VMDBKs which were in the dispatch list.

**RelShr**

The sum of the relative shares of all VMDBKs which were in the dispatch list.

**Consid**

The total number of main storage frames considered when selecting users for dispatching.

**Q0WSS**

The sum of pages in the working sets of class 0 (quick dispatch) users in the dispatch list.

**Q1WSS**

The sum of pages in the working sets of class 1 (interactive) users in the dispatch list.

**Q2WSS**

The sum of pages in the working sets of class 2 users in the dispatch list.

**Q3WSS**

The sum of pages in the working sets of class 3 users in the dispatch list.

*Scheduler Data:***SC\_Pg/s**

The smoothed page read rate, per second, as used by scheduling algorithms.

**SC\_ms/P**

The average elapsed time per page read, as used by scheduling algorithms.

**SC\_PSV**

The page steal value, corresponds to the STEAL parameter in the 'INDICATE LOAD' output.

**SC\_SWgt**

The resource weight for storage. Used by the scheduler to determine how much of a bottleneck the storage resources are.

**SC\_PWgt**

The resource weight for paging. Used by the scheduler to determine how much of a bottleneck the paging resources are.

**SC\_C1P/s**

The average paging rate of class 1 users at the end of their Q1 dispatch list stay. Used by the scheduler to assign a projected paging rate to E1 users on entry to the eligible list.

**SC\_C1WSS**

The average working set size of class 1 users at the end of their Q1 dispatch list stay. Used by the scheduler to assign a projected WSS to E1 users on entry to the eligible list.

**SC\_AIP/s**

The average paging rate of all users in the dispatch and eligible lists.

**SC\_AIWSS**

The average working set size of all users in the dispatch and eligible lists.

**Cl1Exp**

The expansion factor for class 1 users, i.e. the ratio of time it is taking for a class 1 transaction to complete compared with the amount of time it would take if it were running alone on the system.

**Cl2Exp**

The expansion factor for class 2 users, i.e. the ratio of time it is taking for a class 2 transaction to complete compared with the amount of time it would take if it were running alone on the system.

**Cl3Exp**

The expansion factor for class 3 users, i.e. the ratio of time it is taking for a class 3 transaction to complete compared with the amount of time it would take if it were running alone on the system.

*Transaction Data:***Tr\_UP\_T**

The response time, in seconds, for trivial UP transactions.

**Tr\_MP\_T**

The response time, in seconds, for trivial MP transactions.

**TR-T**

The average response time, in seconds, for all trivial transactions (UP and MP).

**Nt\_UP\_T**

The response time, in seconds, for non-trivial UP transactions.

**Nt\_MP\_T**

The response time, in seconds, for non-trivial MP transactions.

**NT-T**

The average response time, in seconds, for all non-trivial transactions (UP and MP).

**UP\_Resp**

The average response time, in seconds, for all UP transactions (trivial and non-trivial).

**QD\_UP\_T**

The response time, in seconds, for all UP transactions of quick dispatch users.

**MeanResp**

The average response time, in seconds, for all transactions on the system (UP and MP, trivial and non-trivial).

**Tr\_UP/s**

The transaction rate per second for trivial UP transactions.

**Tr\_MP/s**

The transaction rate per second for trivial MP transactions.

**TR/S**

The average transaction rate per second for all trivial transactions (UP and MP).

**Nt\_UP/s**

The transaction rate per second for non-trivial UP transactions.

**Nt\_MP/s**

The transaction rate per second for non-trivial MP transactions.

**NT/S**

The average transaction rate per second for all non-trivial transactions (UP and MP).

**UP\_Tr/s**

The average transaction rate per second for all UP transactions (trivial and non-trivial).

**Qd\_Tr/s**

The transaction rate per second for quick dispatch transactions.

**Trans/s**

The average transaction rate per second for all transactions on the system (UP and MP, trivial and non-trivial).

**TR-Q**

The average number of concurrently active users with trivial transactions (UP and MP).

**NT-Q**

The average number of concurrently active users with non-trivial transactions (UP and MP).

**%Tr\_UP**

The percentage of all transactions which were trivial UP transactions.

**%Tr\_MP**

The percentage of all transactions which were trivial MP transactions.

**%Nt\_UP**

The percentage of all transactions which were non-trivial UP transactions.

**%Nt\_MP**

The percentage of all transactions which were non-trivial MP transactions.

**%Q\_Dsp**

The percentage of all transactions which were quick dispatch transactions.

**#Tr\_UP**

The number of trivial UP transactions per sample interval.

**#Nt\_UP**

The number of non-trivial UP transactions per sample interval.

**#Tr\_MP**

The number of trivial MP transactions per sample interval.

**#Nt\_MP**

The number of non-trivial MP transactions per sample interval.

**#Trans**

The total number of transactions per sample interval.

**#AllUP**

The total number of UP transactions per sample interval.

**%TrUP/UP**

The percentage of UP transactions which were trivial transactions.

**%NtUP/UP**

The percentage of UP transactions which were non-trivial transactions.

**ITR**

The 'system internal throughput ratio', i.e. the number of transactions executed on the system per second of total CPU time.

*User Wait States:* The following average user wait state percentages are based on active sample counts where a user in the dormant list is considered to be still active if he is also in SVM wait, i.e. waiting for the completion of a communication with a service virtual machine.

Users' wait states are tested for in the following sequence:

- I/O wait
- Console function wait
- Instruction simulation wait
- Page wait
- CPU wait
- Running
- SVM wait and in the eligible list
- Loading
- Dormant
- Dormant and in SVM wait
- I/O active
- Test idle wait
- SVM wait and test idle wait
- Page active wait
- Other

Only the first matching wait state found is accounted for.

**%Act**

The percentage of samples during which an average user was found to be active, that is, in the dispatch or eligible list, or in dormant state but waiting for some service from a service virtual machine.

**%Run**

The average percentage of active samples during which users were found to be running on a logical processor.

**%CPU**

The average percentage of active samples during which users were found to be waiting to run on the logical processor.

**%LDG**

The average percentage of active samples during which users were found to be in loading state, i.e. with a very high page read rate.

**%PGW**

The average percentage of active samples during which users were found in page wait state.

**%IOW**

The average percentage of active samples during which users were found in I/O wait state.

**%Sim**

The average percentage of active samples during which users were found in instruction simulation wait state.

**%TIW**

The average percentage of active samples during which users were found in test idle wait state, i.e. waiting to be dropped from the dispatch list.

**%CFW**

The average percentage of active samples during which users were found in console function wait, i.e. waiting for CP to execute one of the console function commands which allow altering the state and configuration of the virtual machine (e.g. LINK, ATTACH, and DETACH command).

**%SV+TI**

The average percentage of active samples during which users were found in both SVM wait and test idle wait. SVM wait state indicates that users were waiting for the completion of a communication with a service virtual machine.

**%SV+EL**

The average percentage of active samples during which users were found in both SVM wait and in the eligible list. SVM wait state indicates that users were waiting for the completion of a communication with a service virtual machine.

**%SV+DM**

The average percentage of active samples during which users were found in both SVM wait and in the dormant list. SVM wait state indicates that users were waiting for the completion of a communication with a service virtual machine, which is also the reason why this state is counted as active even though users were also in the dormant list.

**%IOA**

The average percentage of active samples during which users were found to be waiting for something other than the above while an I/O operation started by the user had not yet completed.

**%LIM**

The average percentage of active samples during which users were found to be in the limit list, i.e. prevented from running by a maximum share setting.

**%PGA**

The average percentage of active samples during which users were found in page active wait, i.e. had loaded a wait PSW while waiting for the completion of a page read operation. This wait state is to be expected for guest systems running with PAGEX ON.

**%OTH**

The percentage of active samples during which users were found in other states.

**%Q0**

The average percentage of active samples during which users were found in Q0 (in the dispatch list as a class 0 user).

**%Q1**

The average percentage of active samples during which users were found in Q1 (in the dispatch list as a class 1 user).

**%Q2**

The average percentage of active samples during which users were found in Q2 (in the dispatch list as a class 2 user).

**%Q3**

The average percentage of active samples during which users were found in Q3 (in the dispatch list as a class 3 user).

**%E0-E3**

The average percentage of active samples during which users were found in the eligible list.

*System Wait States:*

The following wait states are calculated from the same source data as the average user wait states above, but they are different in that they are calculated based only on the normal non-dormant state (without considering dormant plus SVM wait a non-dormant state too), i.e. we only look at wait states of VMDBKs in the dispatch list.

**%PQ**

The average percentage of in-queue samples where users were found in page wait state.

**%IQ**

The average percentage of in-queue samples where users were found in page I/O state.

**%LD**

The average percentage of in-queue samples where users were found in loading state, i.e. with a very high page read rate.

**%EL**

The average percentage of in-queue samples where users were found in the eligible list, waiting for a scarce resource to become available.

*Average DASD Load:***Links**

The average number of minidisk links to disks.

**Paths**

The average number of available channel paths to disks.

**DASD/s**

The average number of SSCH and RSCH instructions per disk executed per second from this system.

**Avoid/s**

The average number of real I/O operations per second and per disk which could be avoided because the blocks were found in the minidisk cache.

**Pend**

The average function pending time, in milliseconds. This is the average time an I/O request remained pending in the I/O subsystem due to path busy conditions.

**Disc**

The average device disconnected time, in milliseconds. This is the average time the device remained disconnected from the channel while executing I/O requests, e.g. time spent for executing SEEK and SET SECTOR orders and also any reconnect misses.

**Conn**

The average device connected time, in milliseconds. This is the average time the device was connected to the channel while executing I/O requests (e.g. receiving commands, SEARCH time or actual data transfer).

**Serv**

The average DASD service time, in milliseconds. This is the sum of the function pending, disconnected and connected times.

**Resp**

The average DASD response time, in milliseconds. This is the device service time plus the time an I/O request was waiting to be started. The queue delay time is calculated based on the average I/O request queue length and the I/O rate.

**CUWt**

The average control unit queuing time, i.e. the additional wait time caused if the target device was already busy with an I/O started from another system (in milliseconds).

**Qued**

The average I/O queue length for all disks. The value is taken from CP monitor high-frequency sampling data, it is an indicator for I/O contention on the device level.

Be aware of the restrictions as to the validity of this field in the case of CP owned disks and disks attached to guest systems, as described under [“DEVICE” on page 169](#) for field Req. Qued.

**Busy**

The percentage of elapsed time during which an average disk was busy.

**%Read**

The average percentage of I/O requests which were for READ operations only.

**Cyls**

The average number of cylinders skipped per SEEK order.

**Recov**

The number of SSCHs (*not* the rate) executed for error recovery purposes.

**Dly/s**

The average rate at which I/O requests were delayed due to throttling.

*Total System Services Activity:***IUVM\_T**

The total IUCV and VMCF activity per second.

**IUCV\_T**

The total IUCV activity per second.

**VMCF\_T**

The total VMCF activity per second.

**\*TOT\_T**

The total IUCV activity, per second, by all CP IUCV services.

**\*MSA\_T**

The total activity, per second, for the \*MSGALL IUCV service.

**\*MSG\_T**

The total activity, per second, for the \*MSG IUCV service.

**\*RPI\_T**

The total activity, per second, for the \*RPI IUCV service.

**\*BLK\_T**

The total activity, per second, for the \*BLOCKIO IUCV service.

**\*MON\_T**

The total activity, per second, for the \*MONITOR IUCV service.

**\*SIG\_T**

The total activity, per second, for the \*SIGNAL IUCV service.

**\*CCS\_T**

The total activity, per second, for the \*CCS IUCV service.

**\*SPL\_T**

The total activity, per second, for the \*SPL IUCV service.

**\*SYM\_T**

The total activity, per second, for the \*SYMPTOM IUCV service.

**\*ACC\_T**

The total activity, per second, for the \*ACCOUNT IUCV service.

**\*LOG\_T**

The total activity, per second, for the \*LOGREC IUCV service.

**\*CRM\_T**

The total activity, per second, for the \*CRM IUCV service.

**\*IDE\_T**

The total activity, per second, for the \*IDENT IUCV service.

**\*CNF\_T**

The total activity, per second, for the \*CONFIG IUCV service.

*System Services Activity from CP to Virtual Machines:*

**IUVM>V**

The sum of the IUCV and VMCF communication rates from CP to virtual machines, per second.

**IUCV>V**

The total IUCV communication rate from CP to virtual machines, per second.

**VMCF>V**

The total VMCF communication rate from CP to virtual machines, per second.

**\*TOT>V**

The total IUCV communication rate, per second, by all CP IUCV services from CP to virtual machines.

**\*MSA>V**

The communication rate from CP to virtual machines, per second, for the \*MSGALL IUCV service.

**\*MSG>V**

The communication rate from CP to virtual machines, per second, for the \*MSG IUCV service.

**\*RPI>V**

The communication rate from CP to virtual machines, per second, for the \*RPI IUCV service.

**\*BLK>V**

The communication rate from CP to virtual machines, per second, for the \*BLOCKIO IUCV service.

**\*MON>V**

The communication rate from CP to virtual machines, per second, for the \*MONITOR IUCV service.

**\*SIG>V**

The communication rate from CP to virtual machines, per second, for the \*SIGNAL IUCV service.

**\*CCS>V**

The communication rate from CP to virtual machines, per second, for the \*CCS IUCV service.

**\*SPL>V**

The communication rate from CP to virtual machines, per second, for the \*SPL IUCV service.

**\*SYM>V**

The communication rate from CP to virtual machines, per second, for the \*SYMPTOM IUCV service.

**\*ACC>V**

The communication rate from CP to virtual machines, per second, for the \*ACCOUNT IUCV service.

**\*LOG>V**

The communication rate from CP to virtual machines, per second, for the \*LOGREC IUCV service.

**\*CRM>V**

The communication rate from CP to virtual machines, per second, for the \*CRM IUCV service.

**\*IDE>V**

The communication rate from CP to virtual machines, per second, for the \*IDENT IUCV service.

**\*CNF>V**

The communication rate from CP to virtual machines, per second, for the \*CONFIG IUCV service.

*System Services Activity from Virtual Machines to CP:*

**IUVM>C**

The sum of the IUCV and VMCF communication rates from virtual machines to CP, per second.

**IUCV>C**

The total IUCV communication rate from virtual machines to CP, per second.

**VMCF>C**

The total VMCF communication rate from virtual machines to CP, per second.

**\*TOT>C**

The total IUCV communication rate, per second, by all CP IUCV services from virtual machines to CP.

**\*MSA>C**

The communication rate from virtual machines to CP, per second, for the \*MSGALL IUCV service.

**\*MSG>C**

The communication rate from virtual machines to CP, per second, for the \*MSG IUCV service.

**\*RPI>C**

The communication rate from virtual machines to CP, per second, for the \*RPI IUCV service.

**\*BLK>C**

The communication rate from virtual machines to CP, per second, for the \*BLOCKIO IUCV service.

**\*MON>C**

The communication rate from virtual machines to CP, per second, for the \*MONITOR IUCV service.

**\*SIG>C**

The communication rate from virtual machines to CP, per second, for the \*SIGNAL IUCV service.

**\*CCS>C**

The communication rate from virtual machines to CP, per second, for the \*CCS IUCV service.

**\*SPL>C**

The communication rate from virtual machines to CP, per second, for the \*SPL IUCV service.

**\*SYM>C**

The communication rate from virtual machines to CP, per second, for the \*SYMPTOM IUCV service.

**\*ACC>C**

The communication rate from virtual machines to CP, per second, for the \*ACCOUNT IUCV service.

**\*LOG>C**

The communication rate from virtual machines to CP, per second, for the \*LOGREC IUCV service.

**\*CRM>C**

The communication rate from virtual machines to CP, per second, for the \*CRM IUCV service.

**\*IDE>C**

The communication rate from virtual machines to CP, per second, for the \*IDENT IUCV service.

**\*CNF>C**

The communication rate from virtual machines to CP, per second, for the \*CONFIG IUCV service.

*System Services Communication Failures:*

**IUVM\_F**

The total failure rate, per second, for all IUCV and VMCF communications.

**IUCV\_F**

The total failure rate, per second, for all IUCV communications.

**VMCF\_F**

The total failure rate, per second, for all VMCF communications.

**\*TOT\_F**

The total failure rate, per second, for all CP IUCV services.

**\*MSA\_F**

The failure rate, per second, for the \*MSGALL IUCV service.



**\*MSG\_F**

The failure rate, per second, for the \*MSG IUCV service.

**\*RPI\_F**

The failure rate, per second, for the \*RPI IUCV service.

**\*BLK\_F**

The failure rate, per second, for the \*BLOCKIO IUCV service.

**\*MON\_F**

The failure rate, per second, for the \*MONITOR IUCV service.

**\*SIG\_F**

The failure rate, per second, for the \*SIGNAL IUCV service.

**\*CCS\_F**

The failure rate, per second, for the \*CCS IUCV service.

**\*SPL\_F**

The failure rate, per second, for the \*SPL IUCV service.

**\*SYM\_F**

The failure rate, per second, for the \*SYMPTOM IUCV service.

**\*ACC\_F**

The failure rate, per second, for the \*ACCOUNT IUCV service.

**\*LOG\_F**

The failure rate, per second, for the \*LOGREC IUCV service.

**\*CRM\_F**

The failure rate, per second, for the \*CRM IUCV service.

**\*IDE\_F**

The failure rate, per second, for the \*IDENT IUCV service.

**\*CNF\_F**

The failure rate, per second, for the \*CONFIG IUCV service.

*System Facilities Data:***Done**

The number of CCWs successfully translated by the fast CCW translation function, per second.

**Abort**

The number of attempted fast CCW translations which were aborted, per second.

**Notelig**

The number of CCWs which were found to be not eligible for fast CCW translation, per second.

**TotCCW**

The total number of CCWs handled by the fast CCW translation function, per second.

**Exit(us)**

The total time spent processing user exit code per sample interval, in microseconds.

*User Communication Activity:***Msg/s**

The total IUCV and VMCF activity of an average user, per second.

**ISend**

The average IUCV send activity per second and user.

**IRecv**

The average IUCV receive activity per second and user.

**IFail**

The average number of IUCV data transfer failures per second and user.

**VSnd**

The average number of VMCF data transfers to other users, per second and user.

**VRecv**

The average number of VMCF data transfers from other users, per second and user.

**VFail**

The average number of VMCF data transfer failures per second and user.

**SendQ**

The average number of IUCV send requests queued for users.

**RecvQ**

The average number of IUCV requests in a user's receive queue.

**ReplQ**

The average number of IUCV requests in a user's reply queue.

**VQueue**

The average number of VMCF messages queued per users.

*User Extremes:***Max. CPU**

The percentage of total CPU time consumed by the user with the highest CPU load.

**Max. VEC**

This field is deprecated (dots are displayed).

**Max. I/O**

The I/O rate per second of the user with the highest I/O activity.

**Max. Pag**

The page rate per second of the user with the highest paging activity.

**Max. Res**

The number of resident pages of the user with the highest number of pages resident in main storage.

**Max. MDC**

The rate of minidisk cache inserts per second of the user with the highest minidisk cache insert rate.

**Max. XST**

The number of expanded storage blocks occupied by the user with the highest expanded storage usage.

*Fields added for Storage:***Mn>2GB**

Number of main storage frames above the 2GB line.

**Below2G**

Pages moved below 2GB for translation, per second.

*System Log Data: Fast CCW Translation for Network Devices***N\_Done**

The number of network CCWs successfully translated by the 'fast CCW translation' function, per second.

**N\_Abort**

The number of attempted fast network CCW translations which were aborted, per second.

**N\_NotEl**

The number of network CCWs which were found to be not eligible for fast CCW translation, per second.

**N\_TotCC**

The total number of network CCWs handled by the fast CCW translation function, per second.

*I/O Processor Data:***SSCH/s**

The number of times the I/O processor(s) began a start subchannel, per second.

**Int/s**

The number of I/O interrupts generated by the I/O processor(s), per second.

**%Busy**

The average percentage of high frequency samples during which I/O processors were found to be busy.

**Channel**

The average percentage of channel busy conditions found per 'begin subchannel'.

**Switch**

The average percentage of switch busy conditions found per 'begin subchannel'.

**CtlUnit**

The average percentage of control unit busy conditions found per 'begin subchannel'.

**Device**

The average percentage of device unit busy conditions found per "begin subchannel".

**Usage Notes:**

1. The REDHIST display shows mainly data which could also be viewed in other "by time" logs; it is not expected to be used very frequently. Since the HISTLOG and HISTSUM files also contain a large amount of data, some compromises were made to avoid excessive storage requirements, but which lead also to some restrictions.
  - Only one copy of the REDHIST display exists at any point in time, similar to the detailed user and I/O device displays. If multiple users view and manipulate REDHIST displays concurrently, they will interfere with each other's actions.
  - The operation of the locate function has been changed for this display:
    - a. Only the date and time prefix area and the variable names specified in the second header line can be specified as strings to be located. The actual performance numbers will *not* be scanned.
    - b. The program bases its decision as to whether a forward/backward locate operation or a left/right shift is attempted on the contents of the string specified: If it contains only valid decimal digits and/or the characters / (forward-slash), : (colon), and blanks, a search for a string in the date/time area is assumed; otherwise a left/right shift is attempted.
    - c. When locating a performance variable name in a left/right shift, the directional indicator -/ (hyphen forward-slash) for a backward search will be ignored. The search will always be done from left to right, starting with the first field in the record, *not* the current leftmost field.
    - d. The following special search arguments have been included as mnemonics for locating specific areas within the data:

**CPU**

for general CPU and system load data

**USErs**

for the logged on and active user counts

**STOrutil**

for main storage utilization and paging activity

**XSTore**

for expanded storage utilization and activity

**MDCache**

for minidisk cache data

**AUXstor**

for auxiliary storage data

**SCHqueue**

for scheduler queue lengths

**SCHData**

for scheduler timer pop data

**TRANsact**

for user transaction data

**USTlog**

for average user wait state percentages

**SYSStat**

for some wait states of in-queue users only

**DASD**

for average DASD performance information

**CPSall**

for CP services data (all activity)

**CPSTovm**

for CP services activity from CP to virtual machines

**CPSTOCp**

for CP services activity from virtual machines to CP

**CPSFails**

for CP services failures

**SYSLog**

for fast CCW translations and user exit time

**UCOMmlog**

for user communication data

**MAXusers**

for the highest loads imposed by single users

- The PRINT command has not been implemented for REDHIST displays of HISTSUM and HISTLOG files. Although a print command will be accepted, only a small subset of the data will actually be printed.
- Data retrieval of the REDHIST display based on the old FCONRMT logic (data retrieval via RSCS) has not been implemented. Although a retrieval command will be accepted, the initial screen cannot be manipulated with scroll and shift commands to view more output.

## FCX198, System Load Overview Screen – FCONRMT, SYSLOAD

**Performance Toolkit Subcommand:** "FCONRMT" on page 106, "SYSLOAD" on page 311.

### The System Load Overview Screen (FCX198):

[illegible]

Figure 95. Layout of System Load Overview Screen (FCX198)

See the "Using Remote Performance Monitoring" section in the [z/VM: Performance Toolkit Guide](#) for more information on this screen.

# FCX199, Remote History Data Screen

**Performance Toolkit Subcommand:** “FCONRMT” on page 106.

The Remote History Data Screen (FCX199):

FCX199		Performance Toolkit Remote History Data Display			Node-ID
Sel	Date	Interval	Type of Display		
-	2002/08/24	10:14:28 - 10:15:28	:	CPU	
-	2002/08/24	10:14:28 - 10:15:28	:	USER	
-	2002/08/24	10:15:28 - 10:16:29	:	DEV	
-	2002/08/24	10:15:28 - 10:16:29	:	DEV S	
-	2002/08/24	10:16:29 - 10:17:30	:	UPAGE	
-	2002/08/24	14:31:05 - 14:32:06	:	CPU	
-	2002/08/24	14:31:05 - 14:32:06	:	DEV	
-	2002/08/24	14:32:06 - 14:33:08	:	USER	
-	2002/08/24	13:30:15 - 14:33:08	:	UST	
-	2002/08/24	Initial	14:35:23	:	DEV 260
-	2002/08/24	14:35:23 - 14:37:17	:	DEV 260	
-	2002/08/24	Initial	14:40:52	:	USER PROF CAL
-	2002/08/24	14:40:52 - 14:42:25	:	USER PROF CAL	
-	2002/08/24	14:44:29 - 14:45:32	:	RED	
-	2002/08/24	14:37:17 - 14:51:45	:	DEV 260	
-	2002/08/24	14:58:28 - 14:59:30	:	USER	
-	2002/08/24	14:42:25 - 15:02:38	:	USER PROF CAL	
17 perf. displays saved; insert 's' to display or 'd' to delete					
Command ==>					
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return					

Figure 96. Layout of Remote History Data Screen (FCX199)

See the "Performance History Data Selection Display" section in the [z/VM: Performance Toolkit Guide](#) for more information on this screen.

# FCX201, FCONXREP Summary History Files Screen - FCONXREP

**Performance Toolkit Command:** FCONXREP.

The FCONXREP Summary History Files Screen (FCX201):

FCX201

Perf. Toolkit Numerical Performance Analysis

Perf. Report

Period: Date ==> 2003/01/01 to 2003/02/28 (YYYY/MM/DD)

Time ==> 08 to 17 (all hours in period)

Sel. hours ==> -- -- -- -- -- (only selected hours)

Sel. days ==> M-F\_ \_ \_ \_ \_ (only selected days)

Desired detail: HOURS ==> TOT (ALL or TOT)

DAYS ==> ALL (ALL or TOT)

MONTHS ==> ALL (ALL or TOT)

Variable selection ==> CPU Total CPU load (all processors)

==> PG/S System page rate (reads & writes)

==> C1ES 85% class 1 elapsed time slice

==> TR/S Trivial transactions executed per sec.

==> NT/S Non-triv. transact. executed per sec.

==> -----

==> -----

==> -----

Processing file 'ACUM HISTSUM A1' (new SUM format)

F1=Help F2=Browse F3=Quit F5=XEDIT F10=Execute F12=Return

Figure 97. Layout of FCONXREP Summary History Files Screen (FCX201)

See [Appendix B, “Numerical Analysis of Performance History Data,” on page 971](#) for more information on this screen.

## FCX202, LPAR Load Log – LPARLOG

**Performance Toolkit Subcommand:** “LPARLOG” on page 222.

The LPAR Load Log report (FCX202):

```
FCX202      CPU 3906  SER 646E7  Interval 06:15:34 - 07:03:34  Perf. Monitor

Interval <Partition->
End Time Name      Nr.    Upid #Core Weight Wait-C Cap %Load  <-- Core --> %Busy %vhd  --- --- --- Type TypeCap MT GrpCapNm GrpCap
>>Mean>> BCCC      1      9      4      10      NO NO    ...    .0    .0 --- --- --- CP      ... .. ...
>>Mean>> BCT1      2      2      4      10      NO NO    ...    .0    .0 --- --- --- CP      ... .. ...
>>Mean>> BCT2      3      3      4      10      NO NO    .0    .1    .1 --- --- --- CP      ... .. ...
>>Mean>> BCT3      4     13      4      10      NO NO    .0    .0    .0 --- --- --- CP      ... .. ...
>>Mean>> BDEV1     5      8      3      10      NO NO    .0    .0    .0 --- --- --- CP      ... .. ...
>>Mean>> BKVMD1    13     25      2      10      NO NO    ...    .0    .0 --- --- --- IFL     ... .. ...
>>Mean>> BKVMD2    14     27      4      10      NO NO    ...    .0    .0 --- --- --- IFL     ... .. ...
>>Mean>> BKVMD3    15     28      4      10      NO NO    ...    .0    .0 --- --- --- IFL     ... .. ...
>>Mean>> BKVMD4    16     20      1      10      NO NO    ...    .0    .0 --- --- --- IFL     ... .. ...
>>Mean>> BLX       6      7      4      10      NO NO    ...    .0    .0 --- --- --- CP      ... .. ...
>>Mean>> BPRF1     7     10      4      10      NO NO    ...    .0    .0 --- --- --- CP      ... .. ...
>>Mean>> BPRF3     8     11      4      10      NO NO    .7    23.5  .0 --- --- --- IFL     ... .. ...
>>Mean>> BST1      9      1      4      10      NO NO    .0    .2    .1 --- --- --- CP      ... .. ...
>>Mean>> BST2     10      4      4      10      NO NO    ...    .0    .0 --- --- --- CP      ... .. ...
>>Mean>> BVM1     11      5      4      10      NO NO    ...    .0    .0 --- --- --- CP      ... .. ...
>>Mean>> BVM2     12      6      4      10      NO NO    ...    .0    .2 --- --- --- CP      ... .. ...
>>Mean>> EPAT     17     34      4      10      NO NO    .0    .1    .0 --- --- --- CP      ... .. ...
>>Mean>> EVIC     18     41      2      10      NO NO    ...    .0    .0 --- --- --- CP      ... .. ...
>>Mean>> FCFT     19     38     13     999      NO NO    .0    .5    .0 --- --- --- MIX     ... 2 ...
>>Mean>> PHOS     20     39      4      10      NO NO    .0    .8    .1 --- --- --- CP      ... .. ...
>>Mean>> VMSYSONE  21      ..      0      0      NO NO    ...    ...    ... --- --- --- ..      ... .. ...
>>Mean>> Total     ..      ..     140    1189      .. ..    .8    1.3   .0 --- --- --- ..      ... .. ...

06:20:34 BCCC      1      9      4      10      NO NO    ...    .0    .0 --- --- --- CP      ... .. ...
06:20:34 BCT1      2      2      4      10      NO NO    ...    .0    .0 --- --- --- CP      ... .. ...
06:20:34 BCT2      3      3      4      10      NO NO    .0    .0    .0 --- --- --- CP      ... .. ...
06:20:34 BCT3      4     13      4      10      NO NO    .0    .0    .0 --- --- --- CP      ... .. ...
06:20:34 BDEV1     5      8      3      10      NO NO    .0    .0    .0 --- --- --- CP      ... .. ...
06:20:34 BKVMD1    13     25      2      10      NO NO    ...    .0    .0 --- --- --- IFL     ... .. ...
06:20:34 BKVMD2    14     27      4      10      NO NO    ...    .0    .0 --- --- --- IFL     ... .. ...
06:20:34 BKVMD3    15     28      4      10      NO NO    ...    .0    .0 --- --- --- IFL     ... .. ...
06:20:34 BKVMD4    16     20      1      10      NO NO    ...    .0    .0 --- --- --- IFL     ... .. ...
06:20:34 BLX       6      7      4      10      NO NO    ...    .0    .0 --- --- --- CP      ... .. ...
06:20:34 BPRF1     7     10      4      10      NO NO    ...    .0    .0 --- --- --- CP      ... .. ...
06:20:34 BPRF3     8     11      4      10      NO NO    .2    5.5   .0 --- --- --- IFL     ... .. ...
06:20:34 BST1      9      1      4      10      NO NO    .0    .2    .1 --- --- --- CP      ... .. ...
FCXCMD140E PF-key undefined for current mode
Command ==>
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return
```

Figure 98. Layout of LPAR Load Log report (FCX202)

The information shown is based on CP monitor SYSTEM domain SAMPLE data.

The display shows the activity of all logical partitions, by time, and also some totals and averages for the overall load. The leftmost column contains the end time (in *hh:mm:ss* format) of the monitor interval for which the values were calculated.

The first entries are reserved for the overall averages since the last RESET command or since data collection was activated. They are designated by the string `>>Mean>>` in the `Interval End Time` column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the `>>Mean>>` entries, you can find the relevant interval in the header line.

#### Name

The name of the logical partition to which the following figures apply. A `Total` line is also shown, which contains totals for the overall load and averages for the logical processor load.

#### Nr.

The number of the logical partition. Note that this is an internal number presented by CP monitor data which does *not* correspond to the externally visible partition number.

#### Upid

The user partition ID number for the logical partition.



**#Core**

The number of logical cores defined for the partition.

The total number of real cores on the system is inserted for the Total line.

**Weight**

The weight defined for the partition. Weight values range from 1 to 999, with a default of 10. The string DED is inserted when the partition is dedicated.

For a mixed-engine partition, the weight for the partition's CPs will be shown.

The sum of the individual partitions' weights is inserted for the Total line.

**Wait-C**

The wait completion definition for the logical processor. It can be:

**YES**

if the wait completion flag has been set on, indicating the logical processor will keep running on a real core until its time slice is up even though it is waiting

**NO**

if the wait completion flag is off

**Cap**

The capping definition for the partition. It can be:

**YES**

if capping has been set on for the partition, indicating that CPU cycles are to be allocated only up to the specified maximum

**NO**

if capping is not active

Dots are displayed for systems where the relevant data are not available.

**%Load**

The relative load of the partition on the entire processor complex, expressed as the relation of the total time that logical cores were dispatched for this partition, compared to the totally-available processor time (the number of cores x the elapsed time), expressed as a percentage. If the enhanced LPAR monitor support for secondary processors such as ICFs, IFLs, and ZIIPs is not available, a load value is inserted for only those partitions that have CP cores. Dots are displayed for the relative load of partitions that do not have any CP cores. The total load shown in the Total line is calculated as the total of all partitions and includes LPAR management time where available (see the %Ovhd description). If the enhanced LPAR monitor support for secondary processors is not available, the total load includes only the load of partitions that have CP cores.

**Note:** When the system has enhanced LPAR monitor support, overall physical load figures for CP cores are always calculated correctly. When the system does not have enhanced LPAR monitor support, CP physical load figures will be incorrect if the CEC contains ICF or IFL processors. This is because ICF and IFL cores are not counted in the number of physical cores, but they are treated like CP cores when PR/SM provides z/VM with processor usage data. Because the overall physical load is calculated based on the assumption that there is no data from ICF or IFL cores, the resulting load figure for CPs will be too high.

*Load per Core:*

The following fields show average load figures per logical core, including for the Total line (except for the %Ovhd value). Only the %Busy and %Ovhd fields can be calculated for all partitions. Monitor data for the remaining fields are available only for the logical partition in which the collecting VM system is running.

**Core %Busy**

Average percentage of time logical cores were busy, defined as the percentage of elapsed time during which real cores were assigned to logical cores. Note that the value will include also wait time in the system that is running in the partition if wait completion has been enabled for the core.

**Core %Ovhd**

Average percentage of elapsed time the logical cores spent for LPAR management.

Dots are displayed for systems where the relevant data are not available.

**Exception:** The *Total* line shows *General LPAR mgmt. overhead*. This is general LPAR management time, expressed as a percentage of the available CPU time (elapsed time times the number of physical cores); it does not include LPAR management overhead for specific logical cores which is already shown in column %Ovhd.

This information is available only on systems with the LPAR management time facility.

**Note:** The %Susp, %VMld, and %Logld columns are excluded from this report. All information about Processor/Thread activity is available in [“FCX304, Processor Log Screen – PRCLOG”](#) on page 825.

### Type

The CPU type of the logical cores defined for the partition. Possible values are:

#### CP

Central Processor

#### ICF

Internal Coupling Facility

#### IPL

Integrated Facility for Linux

#### MIX

A mixture of different processor types within one partition

#### ZIIP

IBM Z Integrated Information Processor

### Note:

1. For the ZIIP processor type, the monitor data displays IIP. Performance Toolkit translates this value as ZIIP.
2. While you might also see the first five types in the LPAR Load report (see [“FCX126, LPAR Load Screen – LPAR”](#) on page 472), the MIX type is only valid in this display.

### TypeCap

The maximum amount of CPU power of this type the partition is permitted to use, where a value of 100 represents one whole physical core's worth of CPU power.

For a mixed-engine partition, the value for the partition's CPs will be shown.

The sum of the individual partitions' type caps is inserted for the *Total* line.

If the partition does *not* have a type cap, you will see dots in this field.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

### MT

The multithreading depth for the specified logical core of this partition.

If the partition does *not* have a type cap, you will see dots in this field.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

### GrpCapNm

The Group Capping Name assigned to this partition.

If the partition does *not* have a type cap, you will see dots in this field.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

### GrpCap

The Group Capping value assigned to this partition.

If the partition does *not* have a type cap, you will see dots in this field.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

### Usage note:

The term *core* describes a unit that contains one or more processors. All processors of a core have the same processor type. For machine types that do not have the multithreading (MT) facility, the machine has physical processors and the partitions have logical processors. For machine types that have the MT facility, the machine has physical cores, each of which contains two physical processors. A partition that has not activated the MT facility continues to have logical processors. A partition that has activated the MT facility has logical cores, each of which contains one or more logical processors.

**Effect of AVERAGE and INTERIM subcommands:**

None.

# FCX203, TCP/IP Performance Data Menu Screen – TCPMENU

**Performance Toolkit Subcommand:** “TCPMENU” on page 326.

The TCP/IP Performance Data Menu Screen (FCX203) shows a list of all the available TCP/IP data displays, for all the servers that currently provide TCP/IP monitor records.

FCX203	CPU 9672	SER 65993	TCP/IP Displays	Perf. Monitor
-----	Log File			
S	Userid	Name	Description	
. System	HIPSOCK		Hipersocket channel activity	
. System	QDIO		QDIO device activity	
. System	VNIC		Virtual Network Device activity	
. System	VSWITCH		Virtual Switches activity	
. TCPIP	TCPACTLG		General TCP/IP activity log	
. TCPIP	TCPBPLOG		TCP/IP buffer pools log	
. TCPIP	TCPCONF		TCP/IP server configuration	
. TCPIP	TCPDATLG		General TCP/IP data transfer log	
. TCPIP	TCPDOSLG		TCP/IP denial of service log	
. TCPIP	TCPICMP		TCP/IP ICMP messages log	
. TCPIP	TCPIOLOG		TCP/IP I/O activity log	
. TCPIP	TCPLINKS		TCP/IP links activity log	
. TCPIP	TCPSESS		TCP/IP TCP and UDP sessions log	
. TCPIP	TCPUSERS		TCP/IP users activity log	
. TCPIPX	TCPACTLG		General TCP/IP activity log	
. TCPIPX	TCPBPLOG		TCP/IP buffer pools log	
. TCPIPX	TCPCONF		TCP/IP server configuration	
. TCPIPX	TCPDATLG		General TCP/IP data transfer log	
. TCPIPX	TCPDOSLG		TCP/IP denial of service log	
. TCPIPX	TCPICMP		TCP/IP ICMP messages log	
. TCPIPX	TCPIOLOG		TCP/IP I/O activity log	
. TCPIPX	TCPLINKS		TCP/IP links activity log	
Select performance screen with cursor and press ENTER				
Command ==>				
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return				

Figure 99. Layout of TCP/IP Performance Data Menu Screen (FCX203)

The example is for a system with two TCP/IP servers, TCPIP and TCPIPX.

All TCP/IP logs are based on APPLDATA monitor records created by the server machines: They can be built and displayed only when:

- APPLDATA sample and event data collection has been enabled for the server(s), and when
- The servers themselves have been set up to create APPLDATA monitor records.

Where available, i.e. if hipersockets are defined, the top entry will allow selecting the Hipersocket channel activity display.

Similarly, if QDIO and/or Virtual Switches activity information is available, a QDIO and/or a VSWITCH entry will allow selecting the QDIO or Virtual Switches activity display. A 'Server Userid' of System will be inserted for the Hipersocket, QDIO and Virtual Switches selections.

Both the HIPSOCK and VSWITCH entries will be shown with a 'Server Userid' of System.

The logs are initially sorted to Server Userid; you can change the sequence to have them sorted to the log name instead.

## Field Description:

### S

Cursor selection field.

### Server Userid

Userid of the TCP/IP server machine whose configuration is shown, or activity logged.

### Log File Name

The log file name, equal to the subcommand with which the log data can be selected for display.

**Description**

Description of the performance display.

## FCX204, General TCP/IP Activity Log Screen – TCPACTLG

**Performance Toolkit Subcommand:** “TCPACTLG” on page 318.

The General TCP/IP Activity Log Screen (FCX204):

FCX204	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor						
TCP/IP Activity Log for Server TCPIP											
	<--- Connections/s --->				<----- TCP Segments/s ----->					<---- ARP/s	
Interval	<---- Opens ----->					Trans	Re-	Recvd	X-mit		Reply
End Time	Init	Accept	Fails	Reset	Recvd	mit	X-mit	Error	Reset	Recvd	X-mit
>>Mean>>	.004	.004	.001	.000	.224	.238	.001	.001	.000	4.718	.003
08:34:08	.003	.003	.003	.000	.047	.063	.003	.007	.000	4.477	.003
08:39:08	.000	.000	.000	.000	.007	.013	.003	.000	.000	5.837	.003
08:44:08	.000	.000	.000	.000	.003	.007	.000	.000	.000	5.027	.000
08:49:08	.000	.000	.000	.000	.003	.007	.003	.000	.000	4.540	.003
08:54:08	.000	.000	.000	.000	.007	.013	.000	.000	.000	4.700	.000
08:59:08	.000	.000	.000	.000	.007	.007	.003	.000	.000	3.760	.003
09:04:08	.000	.000	.000	.000	.007	.013	.000	.000	.000	5.867	.000
09:09:08	.000	.000	.000	.000	.000	.000	.000	.000	.000	4.327	.003
09:14:08	.000	.000	.000	.000	.013	.020	.000	.007	.000	3.790	.003
09:19:08	.000	.000	.000	.000	.170	.173	.000	.000	.000	6.310	.003
09:24:08	.000	.000	.000	.000	.447	.467	.000	.000	.000	3.333	.003
09:29:08	.000	.000	.000	.000	1.203	1.243	.000	.000	.000	4.553	.003
09:34:08	.030	.030	.000	.000	.367	.440	.000	.000	.000	6.190	.007
Command ==> _											
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F10=Left	F11=Right	F12=Return				

Figure 100. Layout of General TCP/IP Activity Log Screen (FCX204)

The information shown is based on CP monitor APPLDATA domain SAMPLE data, provided by the TCP/IP server machine.

The display shows general TCP/IP activity by time. The leftmost column contains the end time:

```
(format hh:mm:ss)
```

of the monitor interval for which the values were calculated. Note that only part of the performance variables can be shown on the 80 columns of an ordinary screen: use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window to the left or to the right for viewing the remaining data columns.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### Connections:

##### Opens Init

The rate per second at which TCP connection open requests were initiated.

##### Opens Accpt

The rate per second at which TCP connection open requests were accepted.

##### Opens Fails

The rate per second for TCP connection open failures.

##### Reset

The reset rate per second for TCP connections.

*TCP Segments/s:*

**Recv**

The TCP segments received rate per second.

**Transmit**

The TCP segments transmit rate per second.

**Re-X-mit**

The rate at which TCP segments were retransmitted, per second.

**Recv Error**

The rate at which TCP segments were received that had errors, per second.

**X-mit Reset**

The rate at which TCP segments were transmitted that included a reset, per second.

*ARP/s:* Shows Address Resolution Protocol activity.

**Recv**

The ARP requests received rate per second.

**Reply X-mit**

The rate at which ARP replies were transmitted, per second.

**Req. X-mit**

The rate at which ARP requests were transmitted, per second.

*Denial of Service Packet Discards/s:* These fields are available with TCP/IP function level 420 and later only; dots will be inserted otherwise.

**Smurf**

Smurf denial of service packet discards, per second.

**Fraggle**

Fraggle denial of service packet discards, per second.

**PoDeath**

Ping-o-Death denial of service packet discards, per second.

Dots can be shown instead of TCP/IP activity values, and either of the strings:

Not updated  
Logged off

inserted, when no new monitor records have been received for an interval, or when the server was found to have been logged off.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX205, General TCP/IP Data Transfer Log Screen – TCPDATLG

**Performance Toolkit Subcommand:** “TCPDATLG” on page 321.

The General TCP/IP Data Transfer Log Screen (FCX205):

FCX205	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor
TCP/IP Data Transfer Log for Server TCPIP				
<div> <div> <div>&lt;- IP Packets --&gt;</div> <div>&lt;- Received/s --&gt;</div> </div> <div> <div>&lt;----- Incoming/s -----&gt;</div> <div>&lt;----- Outgoing/s -----&gt;</div> </div> <div> <div>IP Datagrams -----</div> <div>&lt;Fragm. D</div> </div> </div>				
Interval	Total	Headr	Addr	For- Unknwn Dis- <User Prot> Dis- No Recei D
End Time			warded	Protcl card DelTo SupBy card Route ved g
>>Mean>>	2.384	.000	.602	.000 .000 .00 1.782 .245 .000 .000 .000
08:39:08	2.080	.000	.740	.000 .000 .00 1.340 .020 .000 .000 .000
08:44:08	2.160	.000	.580	.000 .000 .00 1.580 .010 .000 .000 .000
08:49:08	1.870	.000	.550	.000 .000 .00 1.320 .013 .000 .000 .000
08:54:08	2.890	.000	.573	.000 .000 .00 2.317 .017 .000 .000 .000
08:59:08	1.510	.000	.413	.000 .000 .00 1.097 .013 .000 .000 .000
09:04:08	2.343	.000	.707	.000 .000 .00 1.637 .017 .000 .000 .000
09:09:08	1.990	.000	.490	.000 .000 .00 1.500 .003 .000 .000 .000
09:14:08	2.610	.000	.410	.000 .000 .00 2.200 .023 .000 .000 .000
09:19:08	2.690	.000	.820	.000 .000 .00 1.870 .177 .000 .000 .000
09:24:08	2.273	.000	.370	.000 .000 .00 1.903 .470 .000 .000 .000
09:29:08	3.577	.000	.533	.000 .000 .00 3.043 1.247 .000 .000 .000
09:34:08	2.447	.000	.750	.000 .000 .00 1.697 .457 .000 .000 .000
Command ==> F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return				

Figure 101. Layout of General TCP/IP Data Transfer Log Screen (FCX205)

The information shown is based on CP monitor APPLDATA domain SAMPLE data, provided by the TCP/IP server machine.

The display shows general TCP/IP data transfer activity by time. The leftmost column contains the end time:

(format hh:mm:ss)

of the monitor interval for which the values were calculated. Note that only part of the performance variables can be shown on the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window to the left or to the right for viewing the remaining data columns.

The bytes received or bytes transmitted fields could wrap during the monitor interval on systems with a lot of activity so the numbers may not be accurate.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval or BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

*IP Packets Received/s:*

#### Total

Total IP packets received rate per second.

#### Errors Headr

Received rate per second for IP packets that had header errors.

#### Errors Addr

Received rate per second for IP packets that had addressing errors.



*IP Datagrams:*

*Incoming/s:*

**Forwarded**

The number of IP datagrams forwarded per second.

**Unknown Protcl**

The number of incoming IP datagrams that specified an unknown protocol, per second.

**Discard**

The number of incoming IP datagrams discarded per second.

**User Prot DelTo**

The number of incoming IP datagrams delivered to IP user protocols per second.

*Outgoing/s:*

**User Prot SupBy**

The number of outgoing IP datagrams supplied by IP user protocols for delivery, per second.

**Discard**

The number of outgoing IP datagrams discarded before delivery, per second.

**No Route**

The number of outgoing IP datagrams that had no route to their destination, per second.

*Fragm. Reass/s:*

**Received**

The number of IP fragments received per second that required reassembly.

**Datagrams**

The number of IP datagrams reassembled per second.

**Reass Error**

The number of IP datagram reassembly errors per second.

*Fragmented/s:*

**Datagrams**

The number of IP datagrams fragmented per second.

**Failures**

The number of IP datagram fragmentation failures per second.

**Created**

The number of IP datagram fragments created per second.

*UDP Datagrams/s.:*

**Received**

The number of UDP datagrams received per second.

**No Listener**

The number of UDP datagrams received per second for ports that had no listener.

**Receive Error**

The number of UDP datagrams received per second that had errors.

**Transmit**

The number of UDP datagrams transmitted per second.

Dots can be shown instead of data transfer values, and either of the strings:

Not updated  
Logged off

inserted, when no new monitor records have been received for an interval, or when the server was found to have been logged off.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX206, TCP/IP ICMP Messages Log Screen – TCPICMP

**Performance Toolkit Subcommand:** “TCPICMP” on page 323.

The TCP/IP ICMP Messages Log Screen (FCX206):

FCX206	CPU nnnn	SER nnnnn	Interval	HH:MM:SS	- HH:MM:SS	Perf. Monitor
TCP/IP ICMP Messages Log for Server TCPIP						
<----- ICMP Messages Received/s -----> <----->						
Interval	Total	Error	Un-	Time	Param	Source
End Time			reach	Exceed	Probl	Quench
>>Mean>>	.004	.000	.000	.000	.000	.000
08:34:08	.003	.000	.000	.000	.000	.000
08:39:08	.003	.000	.000	.000	.000	.000
08:44:08	.003	.000	.000	.000	.000	.000
08:49:08	.003	.000	.000	.000	.000	.000
08:54:08	.003	.000	.000	.000	.000	.000
08:59:08	.003	.000	.000	.000	.000	.000
09:04:08	.003	.000	.000	.000	.000	.000
09:09:08	.003	.000	.000	.000	.000	.000
09:14:08	.003	.000	.000	.000	.000	.000
09:19:08	.003	.000	.000	.000	.000	.000
09:24:08	.003	.000	.000	.000	.000	.000
09:29:08	.003	.000	.000	.000	.000	.000
09:34:08	.003	.000	.000	.000	.000	.000
Command ==> _						
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F10=Left	F11=Right
F12=Return						

Figure 102. Layout of TCP/IP ICMP Messages Log Screen (FCX206)

The information shown is based on CP monitor APPLDATA domain SAMPLE data, provided by the TCP/IP server machine.

The display shows TCP/IP ICMP messages activity by time. ICMP reports problems and relays network specific information, such as error status from some network device.

The leftmost column contains the end time:

```
(format hh:mm:ss)
```

of the monitor interval for which the values were calculated. Note that only part of the performance variables can be shown on the 80 columns of an ordinary screen: use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window to the left or to the right for viewing the remaining data columns.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

**Field Description:**

**Interval End Time**

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

*ICMP Messages Received/s:*

**Total**

Total ICMP messages received per second.

**Error**

Number of ICMP messages received per second that had errors.

**Unreach**

Number of ICMP 'Destination Unreachable' messages received per second.

**Time Exceed**

Number of ICMP 'Time Exceeded' messages received per second.

**Param Probl**

Number of ICMP 'Parameter Problem' messages received per second.

**Source Quench**

Number of ICMP 'Source Quench' messages received per second.

**Redirect**

Number of ICMP 'Redirect' messages received per second.

**Echo**

Number of ICMP 'Echo' and 'Echo Reply' messages received per second.

**Time Stamp**

Number of ICMP 'Timestamp' and 'Time Stamp Reply' messages received per second.

**Addr Mask**

Number of ICMP 'Address Mask' and 'Address Mask Reply' messages received per second.

*ICMP Messages Sent/s:*

**Total**

Total ICMP messages sent per second.

**Error**

Number of ICMP message transmission errors per second.

**Unreach**

Number of ICMP 'Destination Unreachable' messages sent per second.

**Time Exceed**

Number of ICMP 'Time Exceeded' messages sent per second.

**Param Probl**

Number of ICMP 'Parameter Problem' messages sent per second.

**Source Quench**

Number of ICMP 'Source Quench' messages sent per second.

**Redirect**

Number of ICMP 'Redirect' messages sent per second.

**Echo**

Number of ICMP 'Echo' and 'Echo Reply' messages sent per second.

**Time Stamp**

Number of ICMP 'Timestamp' and 'Time Stamp Reply' messages sent per second.

**Addr Mask**

Number of ICMP 'Address Mask' and 'Address Mask Reply' messages sent per second.

Dots can be shown instead of ICMP message activity, and either of the strings:

Not updated  
Logged off

inserted, when no new monitor records have been received for an interval, or when the server was found to have been logged off.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX207, TCP/IP TCP and UDP Sessions Log Screen – TCPSESS

**Performance Toolkit Subcommand:** “TCPSESS” on page 327.

The TCP/IP TCP and UDP Sessions Log Screen (FCX207):

FCX207	CPU	nnnn	SER	nnnnn	Interval	HH:MM:SS	-	HH:MM:SS	Perf.	Monitor
TCP/IP TCP & UDP Sessions Log (open/close) for Server TCPIP										
Time	Acti- vity	Local Port	Remote IP Address (for TCP only)	Client Name	<-- Bytes Receive	---> Sent	<Roundtrip> <Time (s)-> Smooth Var.			
09:31:21	Open	90	9.164.193.181:1080	VMWEBSRV	---	---	---	---	---	---
09:31:21	Close	90	9.164.193.181:1080	VMWEBSRV	250	13539	.010	.010	---	---
09:37:51	Open	80	9.164.178.140:2523	WEBSHAR5	---	---	---	---	---	---
09:37:54	Open	3441	9.164.178.140:0113	WEBSHAR5	---	---	---	---	---	---
09:37:54	Close	3441	9.164.178.140:0113	WEBSHAR5	0	1	.000	1.47	---	---
09:37:55	Close	80	9.164.178.140:2523	WEBSHAR5	201	3291	.133	.074	---	---
09:37:59	Open	80	9.164.178.140:2522	WEBSHAR5	---	---	---	---	---	---
09:37:59	Open	80	9.164.178.140:2521	WEBSHAR4	---	---	---	---	---	---
09:37:59	Open	3442	9.164.178.140:0113	WEBSHAR5	---	---	---	---	---	---
09:37:59	Close	3442	9.164.178.140:0113	WEBSHAR5	0	1	.000	1.47	---	---
09:37:59	Open	3443	9.164.178.140:0113	WEBSHAR4	---	---	---	---	---	---
09:37:59	Close	3443	9.164.178.140:0113	WEBSHAR4	0	1	.000	1.47	---	---
09:38:01	Close	80	9.164.178.140:2522	WEBSHAR5	243	20752	.187	.095	---	---
09:38:02	Close	80	9.164.178.140:2521	WEBSHAR4	251	9101	.062	.043	---	---
Command ===> _										
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return										

Figure 103. Layout of TCP/IP TCP and UDP Sessions Log Screen (FCX207)

The information shown is based on CP monitor APPLDATA domain EVENT data, provided by the TCP/IP server machine.

The display shows TCP and UDP sessions (TCB and UCB open/close activity) by time. See “TCPUSERS” on page 328 for summary data on users' sessions. The leftmost column contains the time:

(format hh:mm:ss)

of the open or close activity. For 'close' events only, information on the activity and performance of the session will be included.

Note that only part of the data can be shown on the 80 columns of an ordinary screen: use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window to the left or to the right for viewing the remaining data columns.

### Field Description:

#### Time

The time of the open or close activity.

#### Activity

Is either Open if the TCB or UCB for the session was opened, or Close if it was closed.

#### Local Port

The number of the local port used.

#### Remote IP Address

The remote IP address (for TCB open/close only), or the string (UDP) (for UCB open/close).

#### Client Name

The client user identifier.

#### Bytes Received

The number of bytes received.

#### Bytes Sent

The number of bytes sent.

*Roundtrip Time (s):* The 'round trip time' is defined as the time interval between sending a segment and receiving an acknowledgment for it.

**Smooth**

The smoothed round trip time in seconds.

**Var.**

The smoothed round trip time variance in seconds.

*Segments:*

**Total**

The total number of segments.

**Unacknowledged**

The maximum number of unacknowledged segments.

*Max. Size:*

**Segment**

The maximum segment size.

**Window Recv**

The maximum receive window size.

**Window Send**

The maximum send window size.

*Max. Buffer Queue Size:*

**Input**

The maximum input buffer queue size.

**Output**

The maximum output buffer queue size.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX208, TCP/IP Links Activity Log Screen – TCPLINKS

**Performance Toolkit Subcommand:** “TCPLINKS” on page 325.

The TCP/IP Links Activity Log Screen (FCX208):

FCX208	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor
TCP/IP Links Activity Log for Server TCPIP				
		<----- Received/s ----->		
		<----- Packets ----->		
		<----- Tra		
Interval	Link Name	Bytes	Uni- cast	Non- Dis- card Error
End Time				Unknown
>>Mean>>	LOSAD00	652.6	.234	6.868 .00 .000
>>Mean>>	IUCVX1	.000	.000	.00 .000
09:14:08	LOSAD00	651.2	.023	6.377 .00 .000
09:14:08	IUCVX1	.000	.000	.00 .000
09:19:08	LOSAD00	766.6	.180	8.820 .00 .000
09:19:08	IUCVX1	.000	.000	.00 .000
09:24:08	LOSAD00	562.3	.457	5.150 .00 .000
09:24:08	IUCVX1	.000	.000	.00 .000
09:29:08	LOSAD00	720.6	1.213	6.917 .00 .000
09:29:08	IUCVX1	.000	.000	.00 .000
09:34:08	LOSAD00	743.3	.387	8.250 .00 .000
09:34:08	IUCVX1	.000	.000	.00 .000
Command ==>				
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return				

Figure 104. Layout of TCP/IP Links Activity Log Screen (FCX208)

The information shown is based on CP monitor APPLDATA domain SAMPLE data, provided by the TCP/IP server machine.

The display shows TCP/IP links activity by time. The leftmost column contains the end time:

(format hh:mm:ss)

of the monitor interval for which the values were calculated. Note that only part of the performance variables can be shown on the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window to the left or to the right for viewing the remaining data columns.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

## Field Description:

### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

### Link Name

The name of the link whose activity is displayed on the line.

### Received/s:

#### Bytes

The number of bytes received per second on this link.

#### Unicast

The number of unicast packets received per second.

**Non-Unicast**

The number of non-unicast packets received per second.

**Discard**

The number of incoming packets discarded per second.

**Error**

The number of incoming packets that had errors, per second.

**Unknown Protocol**

The number of incoming packets that had unknown protocols, per second.

*Transmitted/s:*

**Bytes**

The number of bytes transmitted per second on this link.

**Unicast**

The number of unicast packets transmitted per second.

**Non-Unicast**

The number of non-unicast packets transmitted per second.

**Discard**

The number of outgoing packets discarded per second.

**Error**

The number of outgoing packets that had errors, per second.

Dots can be shown instead of link activity values, and either of the strings:

```
Not updated
Logged off
```

inserted, when no new monitor records have been received for an interval, or when the server was found to have been logged off.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX210, TCP/IP Buffer Pools Log Screen – TCPBPLOG

Performance Toolkit Subcommand: “TCPBPLOG” on page 319.

The TCP/IP Buffer Pools Log Screen (FCX210):

FCX210	CPU nnnn	SER nnnnn	Interval	HH:MM:SS	-	HH:MM:SS	Perf. Monitor						
TCP/IP Buffer Pool Management Log for Server TCPIP													
<----- Buffer Pool Level ----->													
Interval	Acti	Cli-	<- Enve- ->	Raw	Sock	BSD	TCP	UDP	<-Data	Buffer-->	Segm		
End Time	vity	ent	<- lope -->	IP	et	Sock	Ctl	Ctl	Reglr	Small	Tiny	Ackn	CB1k
	CB1k	CB1k	Reglr Large	CB1k	CB1k	CB1k	Blok	Blok					
>>Mean>>	977	60	749	48	50	1927	221	1975	82	120	0	0	...
>>Min.>>	913	57	727	48	50	1851	212	1910	81	114	0	0	...
08:39:08	977	60	749	48	50	1927	221	1975	82	120	0	0	...
08:44:08	977	60	749	48	50	1927	221	1975	82	120	0	0	...
08:49:08	977	60	749	48	50	1927	221	1975	82	120	0	0	...
08:54:08	977	60	749	48	50	1927	221	1975	82	120	0	0	...
08:59:08	977	60	749	48	50	1927	221	1975	82	120	0	0	...
09:04:08	977	60	749	48	50	1927	221	1975	82	120	0	0	...
09:09:08	977	60	749	48	50	1927	221	1975	82	120	0	0	...
09:14:08	977	60	749	48	50	1927	221	1975	82	120	0	0	...
09:19:08	977	60	749	48	50	1927	221	1975	82	120	0	0	...
09:24:08	977	60	749	48	50	1927	221	1975	82	120	0	0	...
09:29:08	977	60	749	48	50	1927	221	1975	82	120	0	0	...
09:34:08	977	56	749	48	50	1924	221	1968	82	120	0	0	...
Command ==>													
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F10=Left	F11=Right	F12=Return						

Figure 105. Layout of TCP/IP Buffer Pools Log Screen (FCX210)

The information shown is based on CP monitor APPLDATA domain SAMPLE data, provided by the TCP/IP server machine.

The display shows TCP/IP buffer pool levels by time. The leftmost column contains the end time:

(format hh:mm:ss)

of the monitor interval for which the values were obtained.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

The second entry is reserved for the minimum buffer pool level values encountered. It is designated by the string >>Min.>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval or BYTIME interval for which the values were obtained.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

*Buffer Pool Level:* The number of free blocks in the pool at the end of the sample interval.

#### Activity CB1k

Activity control block pool level.

#### Client CB1k

Client control block pool level.

#### Envelope Reglr

Regular envelope pool level.

#### Envelope Large

Large envelope pool level.



**Raw IP CBlk**

Raw IP control block pool level.

**Socket CBlk**

Socket control block pool level.

**BSD Sock CBlk**

BSD-style socket control block pool level.

**TCP Ctl Blok**

TCP control block pool level.

**UDP Ctl Blok**

UDP control block pool level.

**Data Buffer Reglr**

Regular data buffer pool level.

**Data Buffer Small**

Small data buffer pool level.

**Data Buffer Tiny**

Tiny data buffer pool level.

**Segm Ackn CBlk**

Segment acknowledgement control block pool level.

**Fixed Page**

Fixed Page Storage Pool fields are:

**Stor Level**

Fixed page storage pool level (full pages).

**Locked Avail**

Fixed page storage pool locked pages that are available (not in use).

No minimums are provided by the TCP/IP server, so dots are inserted into the >>Min.>> line for both the 'Locked Avail' and the 'Locked Used' columns.

**Locked Used**

Fixed page storage pool locked pages that are currently allocated by users of the storage pool.

Dots can be shown instead of buffer pool level values, and either of the strings:

Not updated  
Logged off

inserted, when no new monitor records have been received for an interval, or when the server was found to have been logged off.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX212, TCP/IP Server Configuration Screen – TCPCONF

**Performance Toolkit Subcommand:** “TCPCONF” on page 320.

The TCP/IP Server Configuration Screen (FCX212) shows the buffer pool and TCP/IP links configuration of the selected TCP/IP server machine for the latest interval.

FCX212	CPU nnnn	SER nnnnn	Status	HH:MM:SS	Perf. Monitor	
TCP/IP Configuration for Server TCPIP						
Buffer Pools	Total Blocks	---Allocation Limit---		Free Blocks	Min. Depth	Elem. Size
Activity Control Block	1024	51	102	1007	995	137
Client Control Block	152	7	15	104	104	296
Data Buffer	9000	450	900	8972	8966	16384
Small Data Buffer	500	25	50	498	498	2048
Tiny Data Buffer	10	1	1	3	3	256
Envelope	750	37	75	750	737	2048
Large Envelope	150	7	15	149	149	9216
Host Pool	0	0	0	0	0	0
Raw IP Control Block	52	2	5	50	50	232
Socket Control Block	1006	50	100	912	912	228
BSD-Style Socket Ctl. Block	1001	50	100	955	955	913
TCP Control Block	3004	150	300	2975	2975	667
UDP Control Block	105	5	10	79	78	268
Address Translation	1511	0	5	1504	1504	165
IP Route	301	0	6	282	281	129
Segment Acknowledgement	60080	3004	6008	60080	60078	24
Fixed Page Storage Pool	167	217798	195852	16	0	4096
Overall Storage Utilization:				Virtual Machine Storage Size	1024MB	
				Total Available for Allocation	843MB	
				Largest Block Available	824MB	
Configuration of link LOSAD00 (number X'00000001')						
IBM LCS						
Device number : X'00000000'		Network number : X'00000000'				
Device type : LCS		Network type : TokenRing				
Device name : DOSAD00		Network name : LOSAD00				
Interface number : X'00000001'		Max. transm. unit: 2048 Bytes				
Interface type : iso88025-tokenRing		Interface speed : 4000k bits/s				
Command ==>						
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return						

Figure 106. Layout of TCP/IP Server Configuration Screen (FCX212)

The information shown is based on configuration records provided by the TCP/IP server machine, and passed on by CP monitor as APPLDATA EVENT records. Buffer pool information and information for active links is then continuously refreshed from APPLDATA SAMPLE records to reflect the current status.

### Field Description:

*Buffer Pool Configuration:*

#### Buffer Pools

The description of the control blocks the pool is used for.

#### Total Blocks

The total number of blocks in the pool.

#### Unrestricted

The unrestricted allocation limit.

#### Restricted

The restricted allocation limit.

#### Free Blocks

The number of available blocks.

**Min. Depth**

The minimum depth reached.

**Elem. Size**

The pool element size.

The buffer pool information that is displayed depends on the originating TCP/IP level.

*Overall Storage Utilization:* Shows the sizes of the machine's total virtual storage, the amount of free storage available for allocation, and the size of the largest block of available storage.

**Link Configuration:**

The header lines for each link show the link name, the link number, and the link's description. The following additional information is included for all links:

**Device number**

The device number.

**Device type**

The device type. Possible device types are

- LCS
- ELANS
- ILANS
- DDN1822
- X25ICA
- PVMIUCV
- SNAIUCV
- X25NPSI
- CTC
- HCH
- HIPPI
- CLAW
- SNALU62
- Virtual (VIPA)
- ATM
- OSA Direct Express
- HiperSockets

**Device name**

The device name.

**Interface number**

The interface number.

**Interface type**

The interface type. Possible interface types are

- other
- regular1822
- hdh1822
- ddn-x25
- rfc877-x25
- ethernet-csmacd
- iso88023-csmacd
- iso88024-tokenBus

- iso88025-tokenRing
- iso88026-man
- starLan
- proteon-10Mbit
- proteon-80Mbit
- hyperchannel
- fddi
- lapb
- sdlc
- ds1
- e1
- basicISDN
- primaryISDN
- propPointToPointSerial
- ppp
- softwareLoopback
- eon
- ethernet-3Mbit
- nsip
- slip
- ultra
- ds3
- sip
- frame-relay

(Can be abbreviated to fit into the available space.)

#### **Network number**

The network number.

#### **Network type**

The network type. Possible network types are

- InternalLoopback
- ProNet
- Ethernet
- Ether0r802.3
- Only802.3
- TokenRing
- Util
- IUCV
- CTC
- DDN1822
- X25ICA
- A220
- HIPPI
- FDDI

- CLAWip
  - ControlTask
  - OffloadLink1
  - OffloadApiBroadcastMedia
  - OffloadApiPointToPoint
  - OffloadApiOtherKinds
  - Virt. Device VIPA
  - OSA ATM native mode
  - QDIO Ethernet mode
  - QDIO ATM mode
  - QDIO Token Ring
  - HiperSockets
- (Can be abbreviated to fit into the available space.)
- VSwitch-IUCV

**Network name**

The network name.

**Max. transm. unit**

The size of the maximum transmission unit.

**Interface speed**

The speed of the interface, in bits/s.

Be aware that TCP/IP returns a 'speed' that is just a constant based on the device type. The value may differ from the actual interface's capabilities.

The above fields can be shown for each link. For PVMIUCV, SNAIUCV, IPIIUCV and SNALU62 devices the following additional fields can be included:

**Local node name**

The local node name.

**Local user ident.**

The local user identifier.

**Remote node name**

The remote node name.

**Rem. user ident.**

The remote user identifier.

For CLAW devices the following additional fields can be included:

**Host name**

The host name.

**Adapter name**

The claw adapter name.

**Control task name**

The control task name.

**Read buffers**

The number and size of read buffers.

**Write buffers**

The number and size of write buffers.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX213, User-Defined Displays Selection Menu Screen – UDEFMENU

**Performance Toolkit Subcommand:** [“UDEFMENU” on page 334.](#)

The User-Defined Displays Selection Menu Screen (FCX213):

FCX213	CPU nnnn	SER nnnnn	User_Def Displays	Perf. Monitor
User Defined Screens Selection Menu				
Screen	Data	Description		
S Name	Type			
. SYSTSUM	Interval	Sample DEFSCRN - lines from CPU/DEVICE/CHANNEL/USER		
. GENERAL	Interval	Sample DEFSCRN - lines from USER/DEVICE/CPU		
. MYLOG	By Time	Sample DEFLOG - fields from the SYSTEM screen		
. MYLOG2	By Time	Sample DEFLOG - fields from STORAGE/RESPALL/USER/DEV		

Select performance screen with cursor and hit ENTER  
Command ==>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 107. Layout of User-Defined Displays Selection Menu Screen (FCX213)

The menu shows a listing of all of the currently-active user-defined performance screens with interval data, and of user-defined performance data logs with "by time" data. Displays can be selected by moving the cursor to the desired line and pressing the Enter key, or by entering the screen name (MYLOG, for example) on the command line.

**Field Descriptions:**

**Screen Name**

The name of the user-defined screen or performance data log, as defined with the FC DEFSCRN or FC DEFLOG subcommands. See [“FCONTROL DEFLOG” on page 21](#) and [“FCONTROL DEFSCRN” on page 24](#) for more information about these subcommands.

**Data Type**

The type of performance data that is shown in the display:

**Interval**

For displays that show data for a single interval, as defined by FC DEFSCRN subcommands.

**By Time**

For data logs that show data for multiple intervals, with one output line per interval, as defined by FC DEFLOG subcommands.

**Description**

An optional description of the display, set using the DESC argument of the FC DEFSCRN or FC DEFLOG subcommands.

## FCX214, TCP/IP TCP and UDP Sessions Log Screen – TCPUSERS

**Performance Toolkit Subcommand:** “TCPUSERS” on page 328.

The TCP/IP Users Activity Log Screen (FCX214):

FCX214		CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS				Perf. Monitor			
TCP/IP Users Activity Log for Server TCPIP											
Time	User Name	Completed Sessions		Time /Sess (sec)	<-Total Bytes X-ferred-> <-- TCP -->		<-- UDP -->		<Roundtrip> <Time (s)-> Smooth Var.		<- To
		TCP	UDP		Recvd	Sent	Recvd	Sent			
>>Mean>>	WEBSHAR3	28	0	.9	4794	201k	0	0	.011	.741	
>>Mean>>	WEBSHAR5	2	0	.5	321	9102	0	0	.003	.734	
>>Mean>>	WEBSHAR1	2	0	.5	310	26595	0	0	.004	.734	
>>Mean>>	WEBSHARE	2	0	.5	324	6666	0	0	.003	.735	
>>Mean>>	Total...	34	0	.8	5749	244k	0	0	.009	.740	
07:49:28	WEBSHAR3	22	0	1.0	3921	149k	0	0	.012	.742	
07:49:28	WEBSHAR5	0	0	...	...	...	...	...	...	..	
07:49:28	WEBSHAR1	0	0	...	...	...	...	...	...	..	
07:49:28	WEBSHARE	0	0	...	...	...	...	...	...	..	
07:49:28	Total...	22	0	1.0	3921	149k	0	0	.012	.742	
07:54:28	WEBSHAR3	0	0	...	...	...	...	...	...	..	
07:54:28	WEBSHAR5	0	0	...	...	...	...	...	...	..	
07:54:28	WEBSHAR1	0	0	...	...	...	...	...	...	..	
07:54:28	WEBSHARE	0	0	...	...	...	...	...	...	..	
07:54:28	Total...	0	0	...	...	...	...	...	...	..	
Command ==>											
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return											

Figure 108. Layout of TCP/IP Users Activity Log Screen (FCX214)

The information shown is based on CP monitor APPLDATA domain EVENT data, provided in the form of connection OPEN/CLOSE records by the TCP/IP server machine.

The display shows summary TCP and UDP activity by user and by time; see “TCPSESS” on page 327 for session details. The leftmost column contains the end time:

(format hh:mm:ss)

of the sample interval for which the OPEN/CLOSE information was accumulated.

Note that only part of the data can be shown on the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window to the left or to the right for viewing the remaining data columns.

### Field Description:

#### Time

The end time of the interval for which activity data was accumulated.

#### User Name

The user identification of the virtual machine that handled our side of the sessions.

**Completed Sessions:** Shows the number of sessions which were completed during the sample interval, i.e. for which a CLOSE event record was processed.

#### TCP

The number of completed TCP sessions.

#### UDP

The number of completed UDP sessions.

#### Time/Sess (sec)

The average session duration (elapsed time between OPEN and CLOSE).

*Total Bytes X-ferred:*

**TCP Recvd**

The number of bytes received during TCP sessions.

**TCP Sent**

The number of bytes sent during TCP sessions.

**UDP Recvd**

The number of bytes received during UDP sessions.

**UDP Sent**

The number of bytes sent during UDP sessions.

*Roundtrip Time (s):* The 'round trip time' is defined as the time interval between sending a segment and receiving an acknowledgment for it.

**Smooth**

The smoothed round trip time, in seconds.

**Var.**

The smoothed round trip time variance, in seconds.

*Segments:*

**Total**

The total number of segments.

**Unacknowledged**

The maximum number of unacknowledged segments.

*Max. Buffer Queue Size:*

**Input**

The maximum input buffer queue size.

**Output**

The maximum output buffer queue size.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.



## FCX215, FICON Channel Load Screen – FCHANNEL

**Performance Toolkit Subcommand:** “FCHANNEL” on page 179.

The FICON Channel Load Screen (FCX215):

FCX215 CPU nnnn SER nnnnn Interval HH:MM:SS - HH:MM:SS Perf. Monitor													
-----<----- FICON Channel Utilization % ----->													
<--- Total for System ---> <-Own Partition-->													
<--- Transfer rate, byte/s --->													
<---Total Data--> <---Maximum-->													
Channel Path ID	Shrd	Bus Cycle T_BCy	Work Units T_WUn	<Data> Write T_DUW	Read T_DUR	Work Units L_WUn	<Data> Write L_DUW	Read L_DUR	Write/s	Read/s	M_Write	M_Read	
06	Yes	0	0	0	0	0	0	0	0	0	763M	763M	
0F	Yes	0	0	0	0	0	0	0	6394	0	381M	381M	
12	Yes	0	0	0	0	0	0	0	0	0	250M	250M	
17	Yes	0	0	0	0	0	0	0	0	0	250M	250M	
18	Yes	0	0	0	0	0	0	0	0	0	763M	763M	
19	Yes	0	0	0	0	0	0	0	0	0	763M	763M	
1A	Yes	0	0	0	0	0	0	0	0	0	250M	250M	
1B	Yes	0	0	0	0	0	0	0	0	0	250M	250M	
1C	Yes	0	0	0	0	0	0	0	6598	6598	381M	381M	
1D	Yes	0	0	0	0	0	0	0	0	0	1600M	1600M	
1E	Yes	0	0	0	0	0	0	0	0	0	240M	240M	
1F	Yes	0	0	0	0	0	0	0	0	0	240M	240M	
20	Yes	0	0	0	0	0	0	0	0	0	1600M	1600M	
21	Yes	0	0	0	0	0	0	0	0	0	1600M	1600M	
22	Yes	0	0	0	0	0	0	0	0	0	1600M	1600M	
23	Yes	0	0	0	0	0	0	0	0	0	1600M	1600M	
24	Yes	0	0	0	0	0	0	0	0	0	763M	763M	
25	Yes	0	0	0	0	0	0	0	0	0	763M	763M	
26	Yes	0	0	0	0	0	0	0	0	0	763M	763M	
27	Yes	0	0	0	0	0	0	0	0	0	763M	763M	
28	Yes	0	0	0	0	0	0	0	0	0	240M	240M	
29	Yes	0	0	0	0	0	0	0	0	0	240M	240M	
2A	Yes	2	0	0	0	0	0	0	0	0	800M	800M	
2B	Yes	2	0	0	0	0	0	0	0	0	800M	800M	
30	Yes	0	0	0	0	0	0	0	0	0	240M	240M	

Enter 'CHannel' or 'LCHannel' for channel busy overviews  
 Command ===>  
 F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 109. Layout of FICON Channel Load Screen (FCX215)

FICON channel load analysis is based on data from the Extended Channel Path Measurement Facility, available from the CP monitor SYSTEM domain. The screen shows a list of all channels for which valid information was found.

The sample shows a screen with channels sorted in ascending channel path ID sequence (indicated by the underscore characters above the 'ID' header), the initial sequence. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

### Field Description for FICON Channel Load:

#### Channel Path ID

The hexadecimal 'channel path identifier'

#### Shrd

'Shared' indicator for the channel. It can be either

#### YES

For channels shared with other logical partitions, or

#### NO

For channels dedicated to this logical partition

**FICON Channel Utilization %:** The following fields show the relative utilization of resources compared to the maximum value, in percent.

*Total for System:*

**Bus Cycle T\_BCy**

The bus cycles utilization for the whole system.

**Work Units T\_WUn**

The work units utilization for the whole system.

**Data Units Write T\_DUW**

The data units write utilization for the whole system.

**Data Units Read T\_DUR**

The data units read utilization for the whole system.

*Own Partition:*

**Work Units L\_WUn**

The work units utilization for the own logical partition.

**Data Units Write L\_DUW**

The data units write utilization for the own logical partition.

**Data Units Read L\_DUR**

The data units read utilization for the own logical partition.

*Transfer Rate (Bytes/s):*

**Total Data Write/s**

The total number of bytes written per second, for the whole system.

**Total Data Read/s**

The total number of bytes read per second, for the whole system.

**Maximum M\_Write**

The CHPID hardware's write speed, in bytes/sec.

**Maximum M\_Read**

The CHPID hardware's read speed, in bytes/sec.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all load fields.

# FCX216, Reusable Server Kernel Performance Data Menu Screen – RSKMENU

**Performance Toolkit Subcommand:** “RSKMENU” on page 278.

The Reusable Server Kernel Performance Data Menu Screen (FCX216) shows a list of all the available RSK data displays, for all the servers that currently provide RSK monitor records:

FCX216	CPU nnnn	SER nnnnn	RSK Displays	Perf. Monitor
----- Server	Display			
S Userid	Name	Description		
. BPIDX1	RSKDET	Reusable Server Kernel details		
. BPIDX1	RSKDISK	RSK disk performance log		
. BPIDX1	RSKLOG	General RSK performance log		
. BPIDX1	RSKSTOR	RSK storage utilization log		
. BPSERV	RSKDET	Reusable Server Kernel details		
. BPSERV	RSKDISK	RSK disk performance log		
. BPSERV	RSKLOG	General RSK performance log		
. BPSERV	RSKSTOR	RSK storage utilization log		
. BPSRV1	RSKDET	Reusable Server Kernel details		
. BPSRV1	RSKDISK	RSK disk performance log		
. BPSRV1	RSKLOG	General RSK performance log		
. BPSRV1	RSKSTOR	RSK storage utilization log		
. MPTEST	RSKDET	Reusable Server Kernel details		
. MPTEST	RSKDISK	RSK disk performance log		
. MPTEST	RSKLOG	General RSK performance log		
. MPTEST	RSKSTOR	RSK storage utilization log		
Select performance screen with cursor and press ENTER				
Command ==> _				
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return				

Figure 110. Layout of Reusable Server Kernel Performance Data Menu Screen (FCX216)

The example is for a system with four servers: BPIDX1, BPSERV, BPSRV1 and MPTEST.

All Reusable Server Kernel performance displays are based on APPLDATA monitor records created by the server machines. They can be built and displayed only when both of the following are true:

- APPLDATA sample data collection has been enabled for the server(s)
- The server(s) have been set up to create APPLDATA monitor records.

The logs are initially sorted to Server Userid. You can change the sequence to have them sorted to the log name instead.

**Field Description:**

**S**

Cursor selection field.

**Server Userid**

Userid of the RSK server machine whose performance data are shown.

**Display Name**

The display name, equal to the subcommand with which the performance screen can be selected for display.

**Description**

Short description of the performance display.

## FCX217, Reusable Server Kernel Performance Details Screen – RSKDET

**Performance Toolkit Subcommand:** “RSKDET” on page 275.

The Reusable Server Kernel Performance Details Screen (FCX217) shows detailed performance data for the selected RSK server, with activity information for all the monitor rows that provided data.

```

FCX217      CPU nnnn  SER nnnnn      Status  HH:MM:SS      Perf. Monitor

Reusable Server Kernel Details: BPSERV

<----- Services Activity ----->
Service  Trans <Bytes/s>  Service  Trans <Bytes/s>  Service  Trans <Bytes/s>
Name      /s  Recv  Sent  Name      /s  Recv  Sent  Name      /s  Recv  Sent
>Total>>  7.178  347  501k  APPC      .000    0    0    AUTH      .000    0    0
CACHE     .003    0    1    CMS       .000    0    0    CONFIG    .000    0    0
CONSOLE   .000    0    0    CP        .000    0    0    ECHO      .000    0    0
ENROLL    .013    0    2    HTTP     7.161  347  501k  IUCV      .000    0    0
...

<----- Line Driver Activity ----->
Line      Trans <Bytes/s>  Line      Trans <Bytes/s>  Line      Trans <Bytes/s>
Driver     /s  Recv  Sent  Driver     /s  Recv  Sent  Driver     /s  Recv  Sent
>Total>>  7.178  347  501k  CONSOLE    .017    0    3    MSG        .000    0    0
SP00L     .000    0    0    SUBCOM     .000    0    0    TCP        7.161  347  501k

<----- Enrollment Set Activity ----->  <----- Worker API ----->
Name      Records  Bytes  Ins/s  Remov/s  Retrv/s  ClassName  TotlConn  ActConn
>Total>>   212  71232  .003   .000   28.55  BPINDEX    1         1
HTTPCNFG    3   1008  .003   .000   21.49  .....    ....    ...
HTTPMIME   209  70224  .000   .000   7.064  .....    ....    ...

<----- Trie Indexing Data ----->
Trie      <Data Space>  Records  <-Activity/s-->
Name      MBytes %Used  Indexed  Nodes  Lookups  Records
>Total>>   4096  53  731292  27M   .000   .000
D00000002  2048  34  421008  8697k .000   .000
D00000003  2048  71  310284  18M   .000   .000

<----- File Cache ----->  <-Opens-->  Discrd  <Authorize/s->  <Row Retrvl>
Name      Bytes %Used  Files  /s %Hit  /s  Permit  Inquire  Rows/s  %Hit
HTTPFILE  128M  63  777  6.911  77  .000  .000   .000   .000   ..

<----- Storage Group Activity ----->
Storage  I/O      <--- Page Reads --->  <--- Page Writes --->
Group    Technique  IO/s  Pages/s  msec/Pg  IO/s  Pages/s  msec/Pg
>Total>>  -----  .996  238.2  ...  .996  238.2  ...
SGDS     DataSpace .286  68.30  ...  .286  68.30  ...
SG250    Diag 0250 .710  169.9  ...  .710  169.9  ...

<----- Memory Subpool Activity ----->
Subpool   <-Storage->  <----- ssMemory Activity ----->  <--- CMSSTOR Activity --->
Name      Total Used  <- Allocate ->  <- Release -->  <-Extended->  <-Depleted->
Bytes     %  Rate/s  Bytes Rate/s  Bytes Times  Bytes Times  Bytes
>Total>>   132M  98  114.9  584M  111.5  584M  2609  105M  1302  13584k
BKWSBPAU  102400  100  .000   0  .000   0  0  0  0  0
BKWSBPCA  774208  83  5.540  245808  2.503  122544  938  124384  0  0
...
DMSSBPS0  28776  100  .016  28776  .000  28848  2  28848  0  28848
D00000000  0  ..  .000   0  .000   72  2  72  0  72

Command ===>
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return

```

Figure 111. Layout of Reusable Server Kernel Performance Details Screen (FCX217)

The information shown is based on APPLDATA sample domain records created by the RSK server machine.

The display shows all of the usage and activity information provided by the standard data rows contained in reusable server kernel monitor data. It will not show any application specific data, even if the RSK application does provide such additional information.

Names of rows that show any activity for the selected interval will be highlighted. Where no rows were found for a specific component the corresponding lines will be omitted from the display.

#### **Field Descriptions:**

*Services Activity:* Shows activity data obtained from Service data rows. Entries are sorted alphabetically for service names.

##### **Service Name**

The name of the service.

A '>Total>>' entry with overall totals for all services is inserted first if data for multiple services can be shown.

##### **Trans/s**

The transaction rate per second.

##### **Bytes/s**

The number of bytes transferred per second.

##### **Recv**

Bytes received per second.

##### **Sent**

Bytes sent per second.

*Line Driver Activity:* Shows activity data obtained from Line Driver data rows. Entries are sorted alphabetically for line driver names.

##### **Line Driver**

The line driver name.

A '>Total>>' entry with overall totals for all line drivers is inserted first if data for multiple line drivers can be shown.

##### **Trans/s**

The transaction rate per second.

##### **Bytes/s**

The number of bytes transferred per second.

##### **Recv**

Bytes received per second.

##### **Sent**

Bytes sent per second.

*Enrollment Set Activity:* Shows utilization and activity data obtained from Enrollment data rows. Entries are sorted alphabetically for enrollment set names.

##### **Name**

The enrollment set name.

A '>Total>>' entry with overall totals for all enrollment sets is inserted first if data for multiple enrollment sets can be shown.

##### **Records**

The number of records in the enrollment set.

##### **Bytes**

The number of bytes in use holding records.

##### **Ins/s**

The insertion rate per second.

##### **Remov/s**

The removal rate per second.

**Retr/s**

The retrieval rate per second.

*Worker API:* Shows data from the Worker row.

**ClassName**

The worker class name.

**TotConn**

The total number of connections to the class since the server started.

**ActConn**

The current number of connections to the class.

Data for up to three of the most active classes can be shown.

*Trie Indexing Data:* Shows utilization and activity data obtained from Trie data rows. Entries are sorted alphabetically for trie names.

**Trie Name**

The name of the trie.

A '>Total>>' entry with overall totals for all tries is inserted first if data for multiple tries can be shown.

*Data Space:*

**MBytes**

The data space size in MB.

**%Used**

The percentage of storage in use.

**Records Indexed**

The number of records indexed by the trie.

**Nodes**

The number of nodes in the trie.

*Activity/s:*

**Lookups**

The number of lookups per second.

**Records**

The number of records returned per second.

*File Cache:* Shows utilization and activity data obtained from Cache data rows. Entries are sorted alphabetically for file cache names.

**Name**

The cache name.

A '>Total>>' entry with overall totals for all caches is inserted first if data for multiple caches can be shown.

**Bytes**

The cache size in bytes.

**%Used**

The percentage of cache storage in use.

**Cached Files**

The number of files in the cache.

**Opens /s**

The file opens rate per second (ssCacheFileOpen calls).

**Opens %Hit**

The percentage of cache hits, out of all opens.

**Discrd /s**

The files discarded rate per second.

*Authorization:* Shows activity data obtained from the Authorization data row.

**Authorize/s**

The authorization activity

**Permit**

The permit rate per second.

**Inquire**

The inquire rate per second.

**Row Retrieval**

The row retrieval activity.

**Rows/s**

The number of rows retrieved per second.

**%Hit**

The percentage of row cache hits, out of all row retrievals.

*Storage Group Activity:* Shows activity data obtained from Storage Group data rows. Entries are sorted alphabetically for storage group names.

**Name**

The storage group name.

A '>Total>>' entry with overall totals for all storage groups is inserted first if data for multiple storage groups can be shown.

**I/O Technique**

The technique used for I/O execution. One of the following strings will be inserted:

**Diag 00A4**

Indicates use of Diagnose X'A4'.

**Diag 0250**

Indicates use of Diagnose X'250'.

**DataSpace**

Indicates use of VM Data Spaces.

*Page Reads:*

**IO/s**

The page read I/O rate per second.

**Pages/s**

The number of pages read per second.

**msec/Pg**

The average time needed to read a page, in msec. Page read time can be shown only if Diagnose X'A4' was used for executing I/Os; the value is not available for other I/O techniques.

*Page Writes:*

**IO/s**

The page write I/O rate per second.

**Pages/s**

The number of pages written per second.

**msec/Pg**

The average time needed to write a page, in msec. Page write time can be shown only if Diagnose X'A4' was used for executing I/Os; the value is not available for other I/O techniques.

*Memory Subpool Activity:* Shows utilization and activity data obtained from Memory data rows. Entries are sorted alphabetically for memory subpool names.

**Subpool Name**

The memory subpool name.

A '>Total>>' entry with overall totals for all memory subpools is inserted first if data for multiple subpools can be shown.

**Storage Total**

The total storage used by the subpool, i.e. the sum of free storage in server kernel cache plus allocated storage.

**Storage Used %**

The percentage of total subpool storage that has been allocated.

*ssMemory Activity:*

**Allocate Rate/s**

The ssMemoryAllocate rate per second.

**Allocate Bytes**

The number of bytes allocated by ssMemoryAllocate during the whole interval.

**Release Rate/s**

The ssMemoryRelease rate per second.

**Release Bytes**

The number of bytes released by ssMemoryRelease during the whole interval.

*CMSSTOR Activity:*

**Extended Times**

The number of times storage was extended through CMSSTOR during the interval.

**Extended Bytes**

The number of bytes obtained through CMSSTOR during the interval.

**Depleted Times**

The number of times storage was depleted through CMSSTOR during the interval.

**Depleted Bytes**

The number of bytes returned through CMSSTOR during the interval.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all load fields, except for the 'Worker API' fields that are not affected by current/average settings.



# FCX218, Reusable Server Kernel Performance Log Screen – RSKLOG

**Performance Toolkit Subcommand:** “RSKLOG” on page 277.

The Reusable Server Kernel Performance Log Screen (FCX218):

FCX218	CPU	nnnn	SER	nnnnn	Interval	HH:MM:SS	-	HH:MM:SS	Perf.	Monitor	
Reusable Server Kernel Log: BPSERV											
Interval	<----- Services ----->				<--- Line-Drivers ---->				<----- Tri		
End Time	Act	Trans	<---Bytes--->		Act	Trans	<---Bytes--->		Act	Total	%
>>Mean>>	ive	/s	Recv/s	Sent/s	ive	/s	Recv/s	Sent/s	ive	MBytes	Used Re
11:36:05	23	6.806	336	369k	5	6.806	336	369k	2	4096	53 7
11:41:05	23	.000	0	0	5	.000	0	0	2	4096	53 7
11:46:04	23	10.43	516	651k	5	10.43	516	651k	2	4096	53 7
11:51:05	23	15.26	766	1020k	5	15.26	766	1020k	2	4096	53 7
11:56:05	23	16.44	794	613k	5	16.44	794	613k	2	4096	53 7
12:01:05	23	3.877	195	282k	5	3.877	195	282k	2	4096	53 7
12:06:05	23	.000	0	0	5	.000	0	0	2	4096	53 7
12:11:04	23	.000	0	0	5	.000	0	0	2	4096	53 7
12:16:04	23	.000	0	0	5	.000	0	0	2	4096	53 7
12:21:04	23	.000	0	0	5	.000	0	0	2	4096	53 7
12:26:05	23	.000	0	0	5	.000	0	0	2	4096	53 7
12:31:05	23	.000	0	0	5	.000	0	0	2	4096	53 7
12:36:05	23	.000	0	0	5	.000	0	0	2	4096	53 7
Command ==> _											
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F10=Left	F11=Right	F12=Return				

Figure 112. Layout of Reusable Server Kernel Performance Log Screen (FCX218)

The information shown is based on CP monitor APPLDATA domain SAMPLE data, provided by the RSK server machine.

The display shows activity for a number of Reusable Server Kernel components by time. The leftmost column contains the end time:

```
(format hh:mm:ss)
```

of the monitor interval or the BYTIME interval for which the values were obtained.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

Note that the information extends beyond the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window so as to view the remaining data.

**Field Description:**

**Interval End Time**

The end time of the monitor interval for which the values were obtained.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

*Services:* Shows total activity data obtained from Service data rows.

**Active**

The number of services for which data rows were obtained.

**Trans/s**

The total services transaction rate per second.

**Bytes Recv/s**

The total number of bytes received per second.

**Bytes Sent/s**

The total number of bytes sent per second.

*Line Drivers:* Shows total activity data obtained from Line Driver data rows.

**Active**

The number of line drivers for which data rows were obtained.

**Trans/s**

The total line drivers transaction rate per second.

**Bytes Recv/s**

The total number of bytes received per second.

**Bytes Sent/s**

The total number of bytes sent per second.

*Trie Indexing:* Shows activity data obtained from Trie data rows.

**Active**

The number of trie indices for which data rows were obtained.

**Total MBytes**

The storage reserved for trie indexing, in MB.

**% Used**

The percentage of storage in use.

**Total Records**

The total number of indexed records.

**Total Nodes**

The total number of nodes.

**Activity/s: Lookups**

The number of lookups per second.

**Activity/s: Records**

The number of records returned per second.

Dots can be shown instead of RSK activity data, and either of the strings:

Not updated  
Logged off

inserted, when no new monitor records have been received for an interval, or when the server was found to have been logged off. Dots will also be shown for specific components if no new performance data could be obtained for the component.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX219, Reusable Server Kernel Disk Performance Log Screen – RSKDISK

**Performance Toolkit Subcommand:** “RSKDISK” on page 276.

The Reusable Server Kernel Disk Performance Log Screen (FCX219):




Figure 113. Layout of Reusable Server Kernel Disk Performance Log Screen (FCX219)

The information shown is based on CP monitor APPLDATA domain SAMPLE data, provided by the RSK server machine.

The display shows activity for a number of disk related Reusable Server Kernel components by time. The leftmost column contains the end time:

```
(format hh:mm:ss)
```

of the monitor interval for which the values were obtained.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

Note that the information extends beyond the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window so as to view the remaining data.

## Field Description:

### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were obtained.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

*Authorize:* Shows activity data obtained from the Authorization data row.

### Permit/s

The permit rate per second.

### Inquire/s

The inquire rate per second.

### Rows/s

The number of rows retrieved per second.

**%Hit**

The row retrieval hit percentage.

*Enrollment Sets:* Shows total utilization and activity data obtained from Enrollment data rows.

**Active**

The number of enrollment sets for which data rows were obtained.

**Total Record**

The total number of records in all enrollment sets.

**Total Bytes**

The total number of bytes used by all enrollment sets.

**Activity/s Insrt**

The total insertion rate per second.

**Activity/s Remov**

The total removal rate per second.

**Activity/s Retrv**

The total retrieval rate per second.

*Cache:* Shows total utilization and activity data obtained from Cache data rows.

**Active**

The number of caches for which data rows were obtained.

**Total Bytes**

The total cache size in bytes.

**% Used**

The percentage of bytes used.

**Total Files**

The total number of files cached.

**Opens /s**

The total file opens rate per second.

**Opens %Hit**

The percentage of cache hits, out of all opens.

**Discard /s**

The total files discarded rate per second.

Dots can be shown instead of RSK activity data, and either of the strings:

```
Not updated
Logged off
```

inserted, when no new monitor records have been received for an interval, or when the server was found to have been logged off. Dots will also be shown for specific components if no new performance data could be obtained for the component.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX220, Reusable Server Kernel Storage Log Screen – RSKSTOR

**Performance Toolkit Subcommand:** “RSKSTOR” on page 279.

The Reusable Server Kernel Storage Log Screen (FCX220):

FCX220	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS				Perf. Monitor		
Reusable Server Kernel Storage Log: BPSERV									
	<----- Storage Groups ----->						<----- ss		
	<----- Paging Activity ----->						<---- Alloc		
Interval	Act-ive	Reads	Writes			Act-ive	Total Used	Rate/s	
End Time	Rows	I/O/s	Page/s	ms/Pg	I/O/s	Page/s	ms/Pg	%	
>>Mean>>	2	1.871	447.5	...	1.870	447.4	...	17	100M 90 54.06
11:41:05	3	9.524	2280	...	9.524	2280	...	17	138M 84 93.33
11:46:04	3	1.775	424.2	...	1.775	424.2	...	17	138M 97 114.3
11:51:05	3	1.716	410.6	...	1.716	410.6	...	17	138M 96 117.1
11:56:05	3	1.647	393.8	...	1.647	393.8	...	17	138M 98 33.92
12:01:05	3	1.642	393.4	...	1.642	393.4	...	17	138M 98 2.984
12:06:05	3	1.732	413.1	...	1.732	413.1	...	17	138M 98 3.136
12:11:04	3	1.738	416.4	...	1.738	416.4	...	17	138M 98 3.136
12:16:04	3	1.767	422.9	...	1.767	422.9	...	17	138M 98 3.220
12:21:04	3	1.766	421.9	...	1.766	421.9	...	17	138M 98 3.199
12:26:05	3	1.785	426.7	...	1.785	426.7	...	17	138M 98 3.233
12:31:05	3	1.799	431.1	...	1.802	432.0	...	17	138M 98 3.261
12:36:05	3	1.840	439.8	...	1.836	439.0	...	17	138M 98 3.332
Command ==> _									
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return									

Figure 114. Layout of Reusable Server Kernel Storage Log Screen (FCX220)

The information shown is based on CP monitor APPLDATA domain SAMPLE data, provided by the RSK server machine.

The display shows activity for Reusable Server Kernel storage groups and memory subpools. The leftmost column contains the end time:

```
(format hh:mm:ss)
```

of the monitor interval or the BYTIME interval for which the values were obtained.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

Note that the information extends beyond the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window so as to view the remaining data.

### Field Description:

#### Interval End Time

The end time of the monitor interval for which the values were obtained.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

**Storage Groups:** Shows total activity data for RSK storage groups.

#### Active Rows

The number of storage groups for which data rows were obtained.

**Paging activity, Reads:**

#### IO/s

The total page read I/O rate per second.

#### Page/s

The total number of pages read per second.

**ms/Pg**

The average time needed to read a page, in msec. Any page read time shown will pertain only to I/Os made using Diagnose X'A4'; the time is not available if another I/O technique is used.

*Paging activity, Writes:*

**IO/s**

The total page write I/O rate per second.

**Page/s**

The total number of pages written per second.

**ms/Pg**

The average time needed to write a page, in msec. Any page write time shown will pertain only to I/Os made using Diagnose X'A4'; the time is not available if another I/O technique is used.

*Memory Subpools:* Shows total utilization and activity data for RSK memory subpools.

**Active Rows**

The number of memory subpools for which data rows were obtained.

**Total Bytes**

The total storage used by all subpools, i.e. the sum of all free storage in server kernel cache plus allocated storage.

**Used %**

The percentage of total subpool storage that has been allocated.

*ssMemory Activity:*

**Allocate Rate/s**

The total ssMemoryAllocate rate per second.

**Allocate Bytes**

The total number of bytes allocated by ssMemoryAllocate during the whole interval.

**Release Rate/s**

The total ssMemoryRelease rate per second.

**Release Bytes**

The total number of bytes released by ssMemoryRelease during the whole interval.

*CMSSTOR Activity:*

**Extended Times**

The total number of times storage was extended through CMSSTOR during the interval.

**Extended Bytes**

The total number of bytes obtained through CMSSTOR during the interval.

**Depleted Times**

The total number of times storage was depleted through CMSSTOR during the interval.

**Depleted Bytes**

The total number of bytes returned through CMSSTOR during the interval.

Dots can be shown instead of RSK storage activity data, and either of the strings:

Not updated  
Logged off

inserted, when no new monitor records have been received for an interval, or when the server was found to have been logged off. Dots will also be shown for specific components if no new performance data could be obtained for the component.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX221, SFS and BFS Logs Selection Menu Screen – SFSMENU

**Performance Toolkit Subcommand:** “SFSMENU” on page 292.

The SFS and BFS Logs Selection Menu Screen (FCX221) shows a list of all the available Shared File System and Byte File System performance logs:

FCX221	CPU nnnn	SER nnnnn	SFS/BFS Displays	Perf. Monitor
<b>Shared File System Logs</b>				
S	Log Name	Description		
-	SFSIOLOG	Shared file system I/O activity log		
-	SFSLOG	General shared file system performance log		
-	SFSREQ	Shared file system requests log		
<b>Byte File System Logs</b>				
S	Log Name	Description		
-	BFSFILES	Byte file system file activity log		
-	BFSLOCKS	Byte file system locking activity log		
-	BFSPIPES	BFS pipes and linking activity log		
Select performance screen with cursor and press ENTER				
Command ==>				
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return				

Figure 115. Layout of SFS and BFS Logs Selection Menu Screen (FCX221)

All Shared File System logs are based on APPLDATA monitor records created by the server machines. They can be built and displayed only when APPLDATA sample data collection has been enabled for the server(s).

**Field Description:**

**S**

Cursor selection field. The field will be highlighted and made cursor-selectable (set as 'unprotected') only if the corresponding shared file system or byte file system performance log is available.

**Log Name**

The log file name, equal to the subcommand with which the log data can be selected for display. The field will be highlighted if the corresponding shared file system or byte file system performance log is available.

**Description**

Description of the performance display.

## FCX222, TCP/IP I/O Activity Log Screen – TCPIOLOG

**Performance Toolkit Subcommand:** “TCPIOLOG” on page 324.

The TCP/IP I/O Activity Log Screen (FCX222):

FCX222	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor
TCP/IP I/O Activity Log for Server TCPIP				
Interval	<---- I/O per sec. ---->			
End Time	<Requests-->	<-- Bytes -->	<Data Transfers>	<-- Polls-->
>>Mean>>	Read	Write	Receive	X-mit
07:11:35	75.14	76.97	795k	10284
07:16:35	69.38	69.18	768k	5011
07:21:34	88.17	88.77	993k	6567
07:26:34	79.38	80.08	906k	6668
07:31:35	19.16	18.84	164k	3019
07:36:35	89.74	90.93	988k	7024
07:41:35	90.28	92.51	1027k	9141
07:46:35	77.15	79.18	880k	7625
07:51:35	15.49	15.35	97910	3089
07:56:35	56.86	58.35	579k	6310
08:01:35	91.77	92.25	924k	8709
08:06:34	94.34	95.01	913k	8036
08:11:34	75.62	75.83	765k	6452
	34.43	34.67	337k	5005
Command ===>				
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return				

Figure 116. Layout of TCP/IP I/O Activity Log Screen (FCX222)

The information shown is based on CP monitor APPLDATA domain SAMPLE data, provided by the TCP/IP server machine.

The display shows TCP/IP I/O activity by time. The leftmost column contains the end time

(format hh:mm:ss)

of the monitor interval for which the values were calculated. Note that only part of the performance variables can be shown on the 80 columns of an ordinary screen: use the LEFT/RIGHT commands, or the corresponding PF-keys, for shifting the display window to the left or to the right for viewing the remaining data columns.

The bytes received or bytes transmitted fields could wrap during the monitor interval on systems with a lot of activity so the numbers may not be accurate.

The first entry is reserved for the overall averages since the last RESET command, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

I/O per sec.:

#### Requests read

The read requests rate per second.

#### Requests write

The write requests rate per second.

#### Bytes Receive

The number of bytes received per second.



**Bytes X-mit**

The number of bytes transmitted per second.

*QDIO activity per sec.:*

**Data Transfers Inbound**

The number of inbound QDIO data transfers per second.

**Data Transfers Outbound**

The number of outbound QDIO data transfers per second.

**PCI Inter**

The number of PCI interrupts per second. The PCI interrupt is used by the hardware adapter to request additional QDIO buffers for inbound data transfers. This event should only occur if TCP/IP enters a wait state or is not polling the queues at a sufficient rate.

**Polls Total**

The total poll rate per second.

**Polls Idle**

The number of 'idle' polls per second, i.e. of polling operations in which no QDIO data transfer has taken place.

*IUCV/s.:*

**Receive**

The IUCV requests receive rate per second.

**Reject**

The IUCV requests reject rate per second.

**Reply**

The IUCV requests reply rate per second.

**Send**

The IUCV requests send rate per second.

*VMCF/s.:*

**Success**

The rate per second of successful VMCF sends.

**Abnorm**

The rate per second of abnormal VMCF sends, and of send failures.

Dots can be shown instead of TCP/IP activity values, and either of the strings

Not updated  
Logged off

inserted, when no new monitor records have been received for an interval, or when the server was found to have been logged off.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX223, Linux Systems Selection Screen – LINUX

**Performance Toolkit Subcommand:** “LINUX” on page 215.

The Linux Systems Selection Screen (FCX223) will be shown if you enter the LINUX subcommand with no *linuxid*:

FCX223	CPU nnnn	SER nnnnn	Linux Systems		Perf. Monitor
Selectable Linux Systems					
GFREE18	LINUX1	LINUX2	LINUX3	LINUX4	LINUX5
LINUX6	LINUX7	LNXBENKE	LNXBENK1	W3PILOT1	W3VML
LX00001	LX00002	LX00003	LX00004	LX00005	LX00006
LX00007	LX00008	LX00009	LX00010	LX00011	LX00012
LX00013	LX00014	LX00015	LX00016	LX00017	LX00018
Select a system for Linux details					
Command ==> _					
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return					

Figure 117. Layout of Linux Systems Selection Screen (FCX223)

This is a menu that allows selecting different Linux systems for displaying detailed Linux internal performance data. The display simply lists all the systems for which an IP address has been defined in file FCONX LINUXUSR: actual performance data retrieval is done by sending XML requests to the RMF PM server interface on these Linux systems. Therefore, implementation and activation of that interface is an absolute prerequisite for reporting Linux performance with Performance Toolkit, and the corresponding TCP/IP interface of Performance Toolkit must also be activated (see the “Monitoring Linux Internal Performance Data” section in [z/VM: Performance Toolkit Guide](#) for implementation guidelines).

**Field Description:**

To select a Linux system for performance analysis either:

- Place the cursor on the name of the system you are interested in and press ENTER, or
- Select the system directly by entering the subcommand 'LINUX *linuxid*' on the command line.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX224, Linux Details Selection Screen – LINUX

**Performance Toolkit Subcommand:** “LINUX” on page 215.

The Linux Details Selection Screen (FCX224) will be shown if you specify a *linuxid* on the LINUX subcommand:

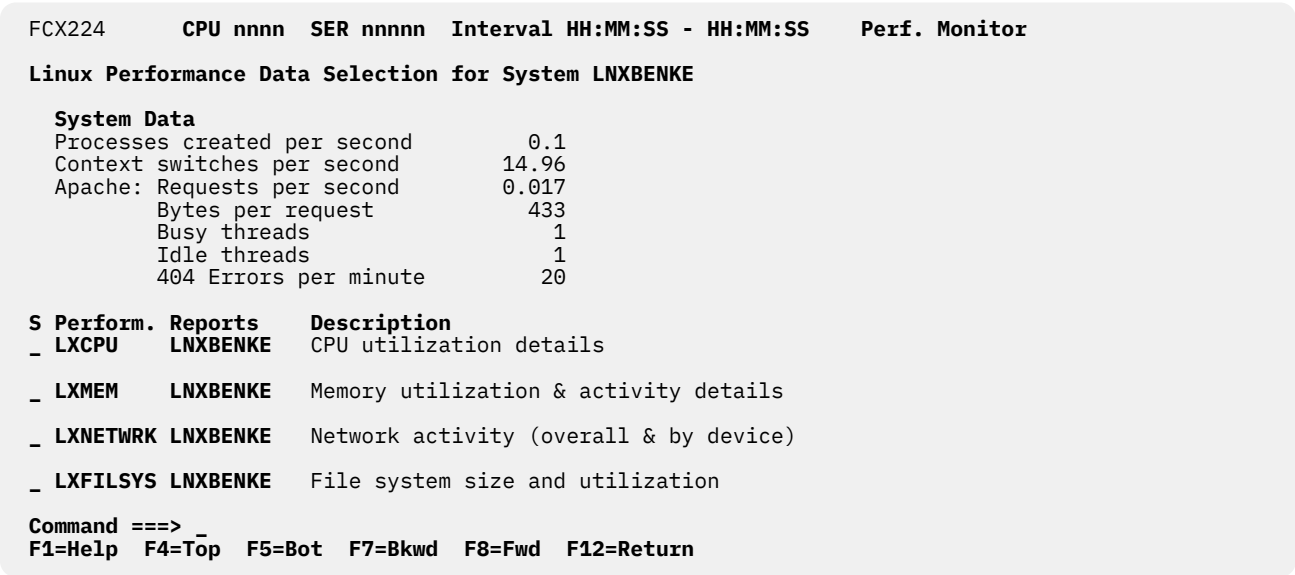


Figure 118. Layout of Linux Details Selection Screen (FCX224)

All of the Linux performance information is retrieved via TCP/IP connection, from the RMF PM DDS interface of the selected Linux system. No data can be shown if either this interface is not available or if the TCP/IP link is down.

The display shows some general 'system' performance values for the selected Linux system and allows selecting other displays with performance data. For more detailed information on CPU, memory, network or Linux filesystem utilization either:

- Place the cursor on the line of the performance data you are interested in and press ENTER, or
- Select the desired screen directly by entering the corresponding subcommand on the command line.

### Field Description for System Data:

#### Processes created per second

The number of processes created per second by this system.

#### Context switches per second

The number of context switches per second on this system.

#### Apache: Requests per second

The number of requests processed per second by the Apache HTTP server.

#### Apache: Bytes per request

The average number of bytes transferred per request to the Apache HTTP server.

#### Apache: Busy threads

The average number of busy threads in the Apache HTTP server.

#### Apache: Idle threads

The average number of idle threads in the Apache HTTP server.

#### Apache: 404 errors per minute

The number of 'file not found' errors (404 errors) per minute.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall period averages will be shown if possible. Linux performance data are collected in intervals that are independent of the CP monitor setup in VM, so it may not be possible to show averages for the same period.

## FCX225, System Summary Log Screen – SYSSUMLG

**Performance Toolkit Subcommand:** “SYSSUMLG” on page 315.

The System Summary Log Screen (FCX225):

FCX225	CPU nnnn SER nnnnn				Interval HH:MM:SS - HH:MM:SS				Perf. Monitor												
<----- CPU ----->																	P				
<--Ratio-->																	o				
Interval	Pct		On-		Log-		SSCH	DASD	Users	<Stg>	<-Paging-->	<Spl>	UP+MP Transactions				----->				
End Time	Busy	T/V	User	line	ged	Activ	+RSCH	Resp	in	PGIN+	Read+	Pages	Triv	Non-	Quick	Triv	Non-	Quick	UnPk	l	
>>Mean>>	.5	4.74	.4582	21.0	42	20	96.8	.4	.0	.0	.0	4.0	.036	.613	.187	4.01	1.37	2.45	7.0	V	
09:25:52	.5	4.27	.4763	21.0	42	20	96.6	.5	.0	.0	.0	3.6	.048	.513	.185	4.64	1.05	2.55	7.1	V	
09:30:52	.5	5.13	.4417	21.0	42	19	96.6	.4	.0	.0	.0	4.0	.026	.657	.199	3.62	1.58	2.35	7.0	V	
09:31:52	.5	4.59	.4822	21.0	42	22	98.0	.2	.0	.1	.0	5.0	.035	.589	.144	4.03	1.30	2.62	7.0	V	
Command ==>																					
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return																					

Figure 119. Layout of System Summary Log Screen (FCX225)

The information shown is based on CP monitor SYSTEM, I/O and USER domain SAMPLE data.

The display shows a summary of system performance by time. The leftmost column contains the end time:

(format hh:mm:ss)

of the monitor interval for which the values were calculated.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

**CPU:** Shows processor load summary.

#### Pct Busy

Average processor percent busy for all processors.

#### Ratio T/V

The average ratio of total to virtual CPU time for all processors.

#### Ratio User

The ratio of user CPU time to total CPU time.

#### Online

The average number of processors that were online during the interval.

**Users:** Shows user counts.

#### Logged

The number of users logged on.

#### Activ

The number of users that were active per monitor sample interval, i.e. that used some CPU.

**I/O:** Shows summary of I/O activity.

#### SSCH+RSCH/s

The total start subchannel and resume subchannel rate per second.

#### DASD Resp msec

The average DASD I/O response time, in milliseconds.

*Stg*: Resource contention.

**Users in Elist**

The number of VMDBKs found waiting in either E1, E2 or E3.

*Paging Rate/s*: Overall paging activity.

**PGIN+PGOUT**

The total expanded storage PGIN plus PGOUT rate per second.

**Read+Write**

The total DASD page read and write rate per second, excluding spool I/O.

*Spl*: Spooling activity.

**Pages/s**

The total rate of spool pages written to, and read from DASD, per second.

*UP+MP Transactions, Response Time*: User response time summary

**Triv**

The average response time for trivial transactions of UP and MP users.

**Non-Triv**

The average response time for non-trivial transactions of UP and MP users.

**Quick Disp**

The average response time for transactions of UP and MP users with the QUICKDSP option.

*UP+MP Transactions, Transactions/s*: User transactions summary

**Triv**

The average trivial UP and MP transactions rate/s

**Non-Triv**

The average non-trivial UP and MP transactions rate/s

**Quick Disp**

The average UP and MP transactions rate/s for users with the QUICKDSP option.

**UnPk**

The average number of processors that were in unparked state during the interval.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

**Polz**

The polarization mode at the end of the measured interval. Possible values are:

**H**

Horizontal mode

**V**

Vertical mode

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see a dash in this field.

**Effect of AVERAGE and INTERIM subcommands:**

None.

## FCX226, User Configuration – UCONF

**Performance Toolkit Subcommand:** “UCONF” on page 333.

The User Configuration report (FCX226) shows virtual machine configuration information for each user:

FCX226	CPU nnnn	SER nnnnn	Status	HH:MM:SS	Perf. Monitor										
-----	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----
Userid	SVM	Mach	Flg	Qck	No	Atta-	Stor	Reserved	<----	Virt.	CPU	-----	<- Share -->	<-- Max Share --->	% CPU
		Mode	ReO	DSP	MDC	ched	Size	Pages	Type	Aff	Def.	Ded.	Stop	SRel.	%
APCSCIF	No	ESA	Off	Off	No	0	64M	0	CP	On	1	0	0	100	...
CFT2ND	No	EME	Off	Off	No	0	256M	0	CP	On	1	0	0	100	...
CFT2NDA	No	EME	Off	Off	No	0	1024M	0	CP	On	1	0	0	100	...
CFT2NDB	No	EME	Off	Off	No	0	3072M	0	CP	On	25	0	8	100	...
									IFL	On	12	...	10	100	...
CHSCTSTN	No	XC	Off	On	No	0	16M	0	CP	On	1	0	0	100	...
DGA4	No	ESA	Off	Off	No	0	32M	0	CP	On	1	0	0	100	...
DGA41	No	ESA	Off	Off	No	0	25M	0	CP	On	1	0	0	100	...
DGA42	No	ESA	Off	Off	No	0	25M	0	CP	On	1	0	0	100	...
DGA43	No	ESA	Off	Off	No	0	22M	0	CP	On	1	0	0	100	...
DGA44	No	ESA	Off	Off	No	0	23M	0	CP	On	1	0	0	100	...
DTCVSW1	Yes	ESA	Off	On	No	0	32M	0	CP	On	1	0	0	100	...
DTCVSW2	Yes	ESA	Off	On	No	0	32M	0	CP	On	1	0	0	100	...
FTPSEVER	Yes	XC	Off	On	No	0	32M	0	CP	On	1	0	0	100	...
FTPSEVER50	No	ESA	Off	Off	No	0	32M	0	CP	On	1	0	0	100	...
GCS	No	ESA	Off	Off	No	0	16M	768	CP	On	1	0	0	100	...
HOWARDJS	No	ESA	Off	Off	No	0	300M	0	CP	On	1	0	0	100	...
MAKER001	No	XC	Off	Off	No	0	18M	0	CP	On	1	0	0	500	...
MCGLYNND	No	ZXC	Off	Off	No	0	32M	0	CP	On	1	0	0	100	...
MISCSEVER	No	ESA	Off	Off	No	0	64M	0	CP	On	1	0	0	100	...
MONWRITE	No	ESA	Off	Off	No	0	4M	0	CP	On	1	0	0	...	2
MPROUTE	Yes	EME	Off	On	No	0	32M	0	CP	On	1	0	0	3000	...
M8USER	No	ESA	Off	Off	No	0	256M	0	CP	On	1	0	0	100	...
OPERATOR	No	ESA	Off	Off	No	0	32M	1024	CP	On	1	0	0	100	...
OPERSYMP	No	ESA	Off	Off	No	0	32M	0	CP	On	1	0	0	100	...
PERFSVM	No	EME	Off	On	No	0	1024M	0	CP	On	1	0	0	...	2
PORTMAP	Yes	ESA	Off	On	No	0	32M	0	CP	On	1	0	0	100	...
PVM	No	ESA	Off	Off	No	0	8M	0	CP	On	1	0	0	3000	...
RELAY	No	ESA	Off	Off	No	0	100M	0	CP	On	1	0	0	100	...
RSCS	Yes	ESA	Off	Off	No	0	32M	768	CP	On	1	0	0	100	...
RSCSDNS	Yes	ESA	Off	Off	No	0	32M	0	CP	On	1	0	0	100	...
RSTL3B	No	Z	Off	On	No	0	4496M	0	CP	On	1	0	0	100	...
RSTL4B	No	EME	Off	On	No	0	1024M	0	CP	On	1	0	0	100	...
SENDIT	No	ESA	Off	Off	No	0	64M	0	CP	On	1	0	0	100	...
SFPURGER	No	ESA	Off	Off	No	0	32M	0	CP	On	1	0	0	100	...

Select a user for user details  
Command ==>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 120. Layout of User Configuration report (FCX226)

The information shown is originally based on CP monitor MONITOR domain configuration data. Where available, the fields are updated continually with the latest information from USER domain sample and event data, because configuration records are only available for users that were already logged on when data collection was started.

Figure 120 on page 679 shows a report with the users sorted in ascending user ID sequence, as indicated by the underscore characters above the Userid heading. This is the initial setting. Other possible sort sequences are indicated by dots above the corresponding columns. These sequences can be set using the SORT subcommand with the appropriate column heading as the sort argument (see “SORT” on page 295), or by positioning the cursor on the heading of the field that is to be sorted and then pressing Enter.

For more detailed information about a specific user, do one of the following:

1. Place the cursor on the Userid field of the virtual machine you are interested in and press Enter.
2. Enter this subcommand on the command line:

```
USER userid
```

### Field Descriptions:

#### Userid

The user identification. This field is highlighted for virtual machines that are currently logged on.

**SVM**

Indicates whether this is a service virtual machine as defined by the SVMSTAT option in the user's directory entry (Yes or No).

This field contains dots if the SVM status is not available for users that logged on after starting data collection.

**Mach Mode**

The architectural mode in which the virtual machine is running:

**EME**

The virtual machine architecture mode was previously ESA/390, but it is now z/Architecture.

**ESA**

The virtual machine architecture mode is ESA/390.

**XC**

The virtual machine architecture mode is ESA/XC.

**Z**

The virtual machine architecture mode is z/Architecture.

**ZXC**

The virtual machine architecture mode is z/XC.

**???**

The virtual machine architecture mode is unknown.

**Flg ReO**

The REORDER processing setting for this user (On or Off).

This field contains dashes when running Performance Toolkit against z/VM 6.1 or earlier data.

**Qck DSP**

Indicates whether the QUICKDSP option has been set for the user (On or Off).

**No MDC Fair**

Indicates whether the NOMDCFS option has been set for the user (Yes or No).

**Attached XSTOR**

The size, in megabytes, of expanded storage that is attached to the user. If this value is 0, no expanded storage is attached.

This field contains dots if Performance Toolkit processes a trend file that is prepared with data collected when running Performance Toolkit against with z/VM 6.1 or earlier data.

**Stor Size**

The virtual machine's storage size.

**Reserved Pages**

The number of pages that are reserved for the virtual machine.

*Virt. CPUs*

The next five fields contain information about virtual processors.

**Type**

The type of virtual processor group. Valid values are:

**CP**

Central Processor

**ICF**

Internal Coupling Facility

**IFL**

Integrated Facility for Linux

**ZIIP**

IBM Z Integrated Information Processor

Dots are displayed if the processor type is unknown.



**Aff**

Indicates the affinity of the virtual processor type. Valid values are:

**Off**

Indicates that CPU affinity has not been set on for this user. When CPU affinity is OFF, the virtual processor is dispatched on only logical processors of primary type.

**On**

Indicates that CPU affinity has been set on for this user with the SET CPUAFFINITY command and that this virtual processor is dispatched on a logical processor of the same type.

**Sup**

Indicates that CPU affinity is suppressed. Although CPU affinity has been set on for this user with the SET CPUAFFINITY command, there is currently no matching logical processor on which to run.

**Def.**

The number of virtual processors of the specified type in the virtual machine's configuration.

**Ded.**

The number of dedicated logical processors in running state (that is, not in stopped state). This value is applicable for the type of primary logical processor only. This field contains dots for any other processor type.

**Stop**

The number of virtual processors that are in a stopped state.

*Share:*

The next two fields contain information about share data.

**SRel.**

The virtual machine's relative share summary value for this CPU type group, if set. Dots are inserted if an absolute share has been set for the virtual machine.

**% SAbs.**

The virtual machine's absolute share summary value for this CPU type group, if set. Dots are inserted if a relative share has been set for the virtual machine.

*Max Share:*

The next three fields contain information about maximum share data.

**Limit**

The type of limit set for the maximum share: Soft or Hard. Dots are inserted if no limit has been set.

**MRel.**

The virtual machine's maximum share relative value, if set. Dots are inserted if no maximum relative share has been set.

**% MAbs.**

The virtual machine's maximum share absolute value, if set. Dots are inserted if no maximum absolute share has been set.

**CPU PoolName**

The CPU pool name that was assigned to the user at the end of the prior interval.

**Effect of AVERAGE and INTERIM subcommands:**

None. The last available status is always displayed.

# FCX227, Linux Network Activity Screen – LXNETWRK

**Performance Toolkit Subcommand:** “LXNETWRK” on page 232.

The Linux Network Activity Screen (FCX227) will be shown if you specify a *linuxid* on the LXNETWRK subcommand:

```
FCX227      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS      Perf. Monitor

Linux Network Activity for System LNXBENKE

-----
Network      <-----> Received/s -----> <-----> Transmitted/s ----->
Device      RcvPack  RcvByte  RcvError  SndPack  SndByte  SndError
>Total>      0.68      17        0        0.78      17        0
ctc0         0.5       0         0         0.6       0         0
lo           0.18      17        0         0.18      17        0

Command ==> _
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 121. Layout of Linux Network Activity Screen (FCX227)

All of the Linux performance information is retrieved via TCP/IP connection, from the RMF PM DDS interface of the selected Linux system. No data can be shown if either this interface is not available or if the TCP/IP link is down.

The display shows network activity, both as total activity and by network device.

The sample shows a screen with the 'by device' values sorted in ascending device name sequence (indicated by the underscore characters above the 'Network Device' header), the initial setting. Other possible sort sequences are indicated by dots above the corresponding columns; they can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

**Field Description:**

**Network Device**

The name of the network device to which the remaining values apply. The string >Total> is inserted for overall totals.

*Received/s*

**RcvPack**

The number of packets received per second. Though the size of a packet is not constant, this, together with the SndPack value, can give you an idea about the network processor utilization and the CPU utilization caused by network traffic. If you have many small network packets you do not have problems with bandwidth but with (network) processor speed.

**RcvByte**

The number of bytes received per second.

**RcvError**

The number of receive errors per second.

*Transmitted/s:*

**SndPack**

The number of packets transmitted per second. Though the size of a packet is not constant, this, together with the RcvPack value, can give you an idea about the network processor utilization and the CPU utilization caused by network traffic. If you have many small network packets you do not have problems with bandwidth but with (network) processor speed.

**SndByte**

The number of bytes transmitted per second.

**SndError**

The number of transmit errors per second.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall period averages will be shown if possible. Linux performance data are collected in intervals that are independent of the CP monitor setup in z/VM, so it may not be possible to show averages for the same period.

# FCX228, Linux Filesystem Usage Screen – LXFILSYS

**Performance Toolkit Subcommand:** “LXFILSYS” on page 228.

The Linux Filesystem Usage Screen (FCX228):

```
FCX228      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor

Linux Filesystem Usage for System LNXBENKE

DASD I/O Activity
I/O request rate per second      ...
I/O response time/request (msec)  ...
I/O response time/sector (msec)  ...

-----
Filesystem      <---- MBytes ---->  <-Percent->
Name            Size      Free  %Used %Free
>Total>         2274      856   60.2  39.7
/dev/dasdb1     2274      856   60.2  39.7

Command ==> _
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 122. Layout of Linux Filesystem Usage Screen (FCX228)

All of the Linux performance information is retrieved via TCP/IP connection, from the RMF PM DDS interface of the selected Linux system. No data can be shown if either this interface is not available or if the TCP/IP link is down.

The display shows overall DASD I/O activity and response time, and both overall and 'by filesystem' information on filesystem size and usage.

The sample shows a screen with the 'by filesystem' values sorted in descending 'filesystem name' sequence (indicated by the underscore characters above the 'Filesystem Name' header), the initial setting. Other possible sort sequences are indicated by dots above the corresponding columns; they can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

**Field Description:**

**DASD I/O Activity:** The following I/O related values are only available for Linux images running 2.4 kernel (or higher) on IBM S/390 or z/Series mainframes.

**I/O request rate per second**

Number of DASD I/O requests per second.

**I/O response time/request (msec)**

Image-wide average response time per DASD I/O request, in milliseconds.

**I/O response time/sector (msec)**

Image-wide average response time per sector, in milliseconds.

**Filesystem Usage:**

**Filesystem Name**

The mount point of the filesystem to which the remaining values apply. The string >Total> is inserted for overall totals.

**MBytes Size**

The filesystem size, in MBytes.

**MBytes Free**

The available space, in MBytes, at the end of a cycle time.

**Percent %Used**

The percentage of filesystem space that was used at the end of a cycle time.

**Percent %Free**

The percentage of filesystem space that was free at the end of a cycle time.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall period averages will be shown if possible. Linux performance data are collected in intervals that are independent of the CP monitor setup in z/VM, so it may not be possible to show averages for the same period.

# FCX229, Linux Memory Utilization and Activity Screen – LXMEM

**Performance Toolkit Subcommand:** “LXMEM” on page 229.

The Linux Memory Utilization and Activity Screen (FCX229) will be shown if you specify a *linuxid* on the LXMEM subcommand:

FCX229	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor																																																							
<b>Linux Memory Util. &amp; Activity Details for System LNXBENKE</b>																																																											
Total memory size	124MB	Swap space size	140MB																																																								
Total memory used	114MB	% Swap space used	0.51%																																																								
Used for buffer	67MB	Swap-in rate	0/s																																																								
Used for shared	37MB	Swap-out rate	0/s																																																								
Used for cache	10MB	Page-in rate	1.566/s																																																								
Total free memory	10MB	Page-out rate	1.566/s																																																								
<table> <tr> <th></th><th><b>Size</b></th><th><b>Page Fault</b></th><th><b>Rate/s</b></th><th></th></tr> <tr> <th></th><th><b>(Bytes)</b></th><th><b>Minor</b></th><th><b>Major</b></th><th><b>&lt;-Incl.Children-&gt;</b></th></tr> <tr> <th><b>Process Name</b></th><th><b>VirtSize</b></th><th><b>ResidSet</b></th><th><b>MinPgFlt</b></th><th><b>MajPgFlt</b></th></tr> <tr> <td>httpd.4545</td><td>12857300</td><td>9332</td><td>0</td><td>0</td></tr> <tr> <td>httpd.3964</td><td>12636200</td><td>9100</td><td>0</td><td>0</td></tr> <tr> <td>httpd.199</td><td>12558300</td><td>8932</td><td>0</td><td>0</td></tr> <tr> <td>httpd.3963</td><td>12546000</td><td>8912</td><td>0</td><td>0</td></tr> <tr> <td>gpmddsriv.10882</td><td>7258110</td><td>5452</td><td>24</td><td>4</td></tr> <tr> <td>gpmddsriv.10883</td><td>7258110</td><td>5452</td><td>0</td><td>0</td></tr> <tr> <td>gpmddsriv.30333</td><td>7258110</td><td>5452</td><td>0</td><td>0</td></tr> <tr> <td>gpmddsriv.4667</td><td>7258110</td><td>5452</td><td>0</td><td>0</td></tr> </table>						<b>Size</b>	<b>Page Fault</b>	<b>Rate/s</b>			<b>(Bytes)</b>	<b>Minor</b>	<b>Major</b>	<b>&lt;-Incl.Children-&gt;</b>	<b>Process Name</b>	<b>VirtSize</b>	<b>ResidSet</b>	<b>MinPgFlt</b>	<b>MajPgFlt</b>	httpd.4545	12857300	9332	0	0	httpd.3964	12636200	9100	0	0	httpd.199	12558300	8932	0	0	httpd.3963	12546000	8912	0	0	gpmddsriv.10882	7258110	5452	24	4	gpmddsriv.10883	7258110	5452	0	0	gpmddsriv.30333	7258110	5452	0	0	gpmddsriv.4667	7258110	5452	0	0
	<b>Size</b>	<b>Page Fault</b>	<b>Rate/s</b>																																																								
	<b>(Bytes)</b>	<b>Minor</b>	<b>Major</b>	<b>&lt;-Incl.Children-&gt;</b>																																																							
<b>Process Name</b>	<b>VirtSize</b>	<b>ResidSet</b>	<b>MinPgFlt</b>	<b>MajPgFlt</b>																																																							
httpd.4545	12857300	9332	0	0																																																							
httpd.3964	12636200	9100	0	0																																																							
httpd.199	12558300	8932	0	0																																																							
httpd.3963	12546000	8912	0	0																																																							
gpmddsriv.10882	7258110	5452	24	4																																																							
gpmddsriv.10883	7258110	5452	0	0																																																							
gpmddsriv.30333	7258110	5452	0	0																																																							
gpmddsriv.4667	7258110	5452	0	0																																																							
<b>Command ==&gt;</b> <b>F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return</b>																																																											

Figure 123. Layout of Linux Memory Utilization and Activity Screen (FCX229)

All of the Linux performance information is retrieved via TCP/IP connection, from the RMF PM DDS interface of the selected Linux system. No data can be shown if either this interface is not available or if the TCP/IP link is down.

The display shows overall Linux memory utilization and activity, as perceived by the Linux system, and also utilization and activity details by process.

The sample shows a screen with the 'by process' values sorted in descending 'resident set' size sequence (indicated by the underscore characters above the 'ResidSet' header), the initial setting. Other possible sort sequences are indicated by dots above the corresponding columns; they can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

## Field Description:

### Overall Utilization and Activity:

#### Total memory size

Total size of memory available, in MBytes.

#### Total memory used

Memory used, in MBytes, at the end of a cycle time. This includes memory used for buffer caches.

#### Used for buffer

Memory used for buffers, in MBytes, at the end of a cycle time. The Linux buffer cache is a disk cache intended to relieve processes from having to wait for relatively slow disks to retrieve or store data. Linux automatically uses otherwise unused memory for disk buffers. When free memory becomes scarce, buffer frames are automatically released.

The buffer cache kernel thread can be tuned using */proc/sys/vm/bdflush*. Refer to the file *Documentation/proc.txt* in the Linux kernel source tree.

**Used for shared**

Memory, in MBytes, usable by more than one process. If any part of memory could be used by more than one process it is counted to be shared memory.

**Used for cache**

Memory used for cache, in MBytes, at the end of a cycle time.

**Total free memory**

The size, in MBytes, of memory that was not allocated at the end of a cycle time.

**Swap space size**

The size of the max. available swap space, in MBytes.

**% Swap space used**

Percentage of swap space used at the end of a cycle time.

**Swap-in rate**

The number of pages swapped in per second. Note that the Linux kernel uses swapping only if there is not enough memory available. Swapping is normally bad for performance and application responsiveness. Exploiting VM virtual disks in storage as swapping devices can help reduce the negative impact to performance or swapping.

**Swap-out rate**

The number of pages swapped out per second. Note that the Linux kernel uses swapping only if there is not enough memory available. Swapping is normally bad for performance and application responsiveness. Exploiting VM virtual disks in storage as swapping devices can help reduce the negative impact to performance or swapping.

**Page-in rate**

The number of pages paged in per second.

**Page-out rate**

The number of pages paged out of memory per second.

**Utilization and Activity by Process:****Process Name**

The name of the process.

*Size*

**VirtSize**

Virtual memory size at the end of a cycle time, in bytes. This is normally a large number, but most parts of the virtual size are often never used (not even paged in).

**ResidSet**

Resident set size, in kBytes. The resident set size, RSS, is the total size of all parts of a process (code, data, shared libraries) actually resident in memory. This can tell you how much memory the processes are 'really' using.

*Page Fault Rate/s*

**Minor MinPgFlt**

Number of minor page faults per second for the process. If there is (from the Kernel's perspective) no disk access involved in the page fault (like Copy on Write for a shared page), the page fault is called a *minor* page fault.

**Major MajPgFlt**

Number of major page faults per second for the process. A major page faults occurs if - from the Linux operating system's perspective - disk access is involved in handling the page fault.

**Incl.Children MinPfltC**

Number of minor page faults per second for the process, including all children processes. If there is (from the Kernel's perspective) no disk access involved in the page fault (like Copy on Write for a shared page), the page fault is called a *minor* page fault.

**Incl.Children MajPfltC**

Number of major page faults per second for the process, including all children processes. A major page fault occurs if - from the Linux operating system's perspective - disk access is involved in handling the page fault.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall period averages will be shown if possible. Linux performance data are collected in intervals that are independent of the CP monitor setup in VM, so it may not be possible to show averages for the same period.



## FCX230, Linux CPU Utilization Screen – LXCPU

**Performance Toolkit Subcommand:** “LXCPU” on page 226.

The Linux CPU Utilization Screen (FCX230) will be shown if you specify a *linuxid* on the LXCPU subcommand:

FCX230	CPU nnnn	SER nnnnn	Interval	HH:MM:SS	- HH:MM:SS	Perf. Monitor		
Linux CPU Utilization for System LNXBENKE								
Processor	<--- Percent		CPU Utilization			<- Accumulated (s) ->		
	Total	User	Kernel	Nice	Idle	TotTm	UserTm	KernTm
>>Mean>>	0.35	0.24	0.10	0	99.64	---	---	---
cpu0	0.68	0.46	0.21	0	99.31	---	---	---
cpu1	0.61	0.48	0.13	0	99.38	---	---	---
cpu2	0.04	0	0.04	0	99.95	---	---	---
cpu3	0.01	0	0.01	0	99.98	---	---	---
Process Name								
gpmddsrv.10882	0.95	0.8	0.15	0	---	308.4	257.6	50.79
gengat.4651	0.03	0	0.03	0	---	668.5	2.61	665.9
procgat.4662	0.03	0.01	0.01	0	---	828.8	66.23	762.6
httpd.3964	0.01	0	0.01	0	---	4.84	0.58	4.26
nscd.274	0.01	0	0.01	0	---	62.89	3.49	59.4
atd.210	0	0	0	0	---	0	0	0
cron.255	0	0	0	0	---	0.07	0	0.07
httpd.199	0	0	0	0	---	4.41	2.08	2.33
inetd.217	0	0	0	0	---	0.06	0.05	0.01
Command ==>								
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return			

Figure 124. Layout of Linux CPU Utilization Screen (FCX230)

All of the Linux performance information is retrieved via TCP/IP connection, from the RMF PM DDS interface of the selected Linux system. No data can be shown if either this interface is not available or if the TCP/IP link is down.

The display shows overall Linux CPU utilization, as it is perceived by the Linux system, and also utilization details by processor and by process.

The sample shows a screen with the 'by process' values sorted in descending total CPU utilization sequence (indicated by the underscore characters above the 'Total' header), the initial setting. Other possible sort sequences are indicated by dots above the corresponding columns; they can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER. The SORT subcommand does not affect the 'by processor' lines which will always be sorted in processor name sequence.

### Field Description:

#### Processor/Process Name

The processor designation (for CPU utilization by processor) or the name of the process (for CPU utilization by process). The string >>Mean>> is inserted for overall averages.

#### Percent CPU Utilization

##### Total

Percentage of CPU total active time.

##### User

Percentage of CPU time in user mode, i.e. executing ordinary user code. If a user program calls a system routine like *open()*, the execution time needed for the *open()* system routine is counted as kernel mode, while the normal processing of the user program is counted as user mode.

##### Kernel

Percentage of CPU time in kernel mode.

**Nice**

*For >>Mean>> and 'by processor' lines:* Percentage of CPU time in 'nice' mode. A process runs in 'nice' mode if the scheduling priority is lower than normal. If the superuser has increased the scheduling priority of some processes to values higher than normal, this is not considered 'nice' mode.

*For 'by process' lines:* The 'nice' value assigned to the process. It can be changed with the *nice* operator command to values ranging from -20 (highest priority) to 19 (lowest priority).

**Idle**

Percentage of CPU idle time. The value is inserted for the overall average and 'by processor' lines only.

*Accumulated (s):* Shows accumulated CPU time in seconds, by process.

**TotTm**

Accumulated CPU time, counted since process creation.

**UserTm**

Accumulated CPU time in user mode, counted since process creation.

**KernTm**

Accumulated CPU time in kernel mode, counted since process creation.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall period averages will be shown if possible. Linux performance data are collected in intervals that are independent of the CP monitor setup in VM, so it may not be possible to show averages for the same period.

## FCX231, Hipersocket Channel Activity Screen – HIPSOCK

**Performance Toolkit Subcommand:** “HIPSOCK” on page 194.

### The Hipersocket Channel Activity Screen (FCX231):

```

FCX231          CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS      Perf. Monitor

-----
<----- Hipersocket Activity/Sec ----->
<--- Total for System ---> <----- Own Partition ----->
<-Transferred--> Failed <-Transferred--> <--- Failed --->
ID      Shrd  T_Msgs  T_DUnits  T_NoBuff  L_Msgs  L_DUnits  L_NoBuff  L_Other
FB      No    .63    15033    1.94    .59    15033    1.94    1.94
FC      No    .00     ...     .00     .00     .00     .00     .00
FD      No    .00     .00     .00     .00     .00     .00     .00
FE      No    .00     .00     .00     .00     .00     .00     .00

Command ==> _
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return

```

Figure 125. Layout of Hipersocket Channel Activity Screen (FCX231)

Hipersocket channel load analysis is based on data from the Extended Channel Path Measurement Facility, available from the CP monitor SYSTEM domain. The screen shows a list of all hipersocket channels for which valid information was found.

The sample shows a screen with channels sorted in ascending channel path ID sequence (indicated by the underscore characters above the 'ID' header), the initial sequence. Other possible sort sequences are indicated by dots above the corresponding columns; they can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on [page 295](#)) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

### Field Description for Hipersockets Channel Load:

### Channel Path ID

The hexadecimal 'channel path identifier'

## Shrd

'Shared' indicator for the channel. It can be either

**YES**

For channels shared with other logical partitions, or

**NO**

For dedicated channels

*Hipersocket Activity/Sec.:* The remainder of the screen shows hipersockets activity per second during the selected interval.

*Total for System:*

## Transferred T\_Msgs

The number of messages sent per second for the whole system.

### Transferred T\_DUnits

The number of data units sent per second for the whole system.

### Failed T NoBuff

The number of sends per second that failed due to unavailable receive buffers, for the whole system.

*Own Partition:*

**Transferred L\_Msgs**

The number of messages sent per second by programs in the own partition.

**Transferred L\_DUnits**

The number of data units sent per second for the own partition.

**Failed L\_NoBuff**

The number of sends per second that failed due to unavailable receive buffers, for the own partition.

**Failed L\_Other**

The number of sends per second that failed due to other problems, for the own partition.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all load fields.

## FCX232, I/O Processor Log Screen – IOPROCLG

**Performance Toolkit Subcommand:** “IOPROCLG” on page 203.

The I/O Processor Log Screen (FCX232) displays I/O processor activity by time.

FCX232	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS				Perf. Monitor
Interval End Time	Proc Number	<-Activity/Sec--> Beg_SSCH I/O_Int		Proc %Busy	<- Busy conditions Channel Switch	per CU	SSCH -> Device
>>Mean>>	00	223.9	855.5	.291	.006	.000	.000
>>Mean>>	01	223.9	877.0	.243	.006	.000	.000
>>Mean>>	02	223.9	827.6	.241	.006	.000	.000
>>Mean>>	03	223.9	865.6	.258	.006	.000	.000
>>Mean>>	04	224.0	33.4	.072	.000	.000	.000
>>Mean>>	05	223.9	32.3	.075	.000	.000	.000
>>Mean>>	06	223.9	25.6	.075	.000	.000	.000
>>Mean>>	07	223.9	18.3	.072	.000	.000	.000
>>Mean>>	08	223.9	2.8	.069	.000	.000	.000
>>Mean>>	09	223.9	2.9	.069	.000	.000	.000
>>Mean>>	0A	223.9	8.7	.072	.000	.000	.000
>>Mean>>	0B	223.9	24.2	.072	.000	.000	.000
>>Mean>>	0C	223.9	6.3	.067	.000	.000	.000
>>Mean>>	0D	223.8	15.1	.074	.000	.000	.000
>>Mean>>	0E	223.8	6.1	.066	.000	.000	.000
>>Mean>>	0F	223.9	6.0	.076	.000	.000	.000
14:55:07	00	222.3	847.2	.352	.002	.000	.000
14:55:07	01	222.8	869.0	.269	.001	.000	.000
14:55:07	02	222.5	819.7	.242	.001	.000	.000
14:55:07	03	222.9	869.0	.255	.002	.000	.000
14:55:07	04	222.5	33.1	.076	.000	.000	.000
14:55:07	05	222.9	32.3	.062	.000	.000	.000
14:55:07	06	223.5	25.7	.131	.000	.000	.000
14:55:07	07	222.9	18.7	.055	.000	.000	.000
14:55:07	08	222.8	2.9	.076	.000	.000	.000
14:55:07	09	222.3	2.9	.076	.000	.000	.000
14:55:07	0A	223.2	8.6	.069	.000	.000	.000
14:55:07	0B	222.3	24.1	.083	.000	.000	.000
14:55:07	0C	222.5	6.4	.097	.000	.000	.000
14:55:07	0D	222.6	15.1	.097	.000	.000	.000
14:55:07	0E	222.4	6.7	.097	.000	.000	.000
14:55:07	0F	223.0	6.2	.083	.000	.000	.000
14:56:07	00	225.4	865.6	.359	.003	.000	.000
14:56:07	01	225.5	886.4	.242	.002	.000	.000
14:56:07	02	225.8	832.3	.311	.002	.000	.000
14:56:07	03	226.2	866.2	.359	.003	.000	.000
14:56:07	04	226.0	32.9	.055	.000	.000	.000
14:56:07	05	225.8	32.1	.069	.000	.000	.000
14:56:07	06	225.1	26.9	.097	.000	.000	.000
14:56:07	07	225.1	17.9	.083	.000	.000	.000
14:56:07	08	225.5	3.0	.062	.000	.000	.000
14:56:07	09	225.8	3.1	.062	.000	.000	.000
14:56:07	0A	225.6	9.2	.048	.000	.000	.000
14:56:07	0B	225.8	24.6	.028	.000	.000	.000
14:56:07	0C	225.7	6.4	.076	.000	.000	.000
14:56:07	0D	225.5	15.3	.083	.000	.000	.000
14:56:07	0E	225.8	5.9	.104	.000	.000	.000
14:56:07	0F	225.5	7.0	.062	.000	.000	.000
Command ==>							
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return		

Figure 126. Layout of I/O Processor Log Screen (FCX232)

The information shown is based on CP monitor PROCESSOR domain SAMPLE data.

On this screen, the leftmost column contains the end time (in *hh:mm:ss* format) of the interval for which the values were calculated.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

**Field Descriptions:****Interval End Time**

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

**Proc Number**

The I/O processor ID in hexadecimal format.

*Activity/Sec:*

**Beg\_SSCH**

The number of times per second that the I/O processor began a 'start subchannel'.

**I/O\_Int**

The number of times per second that the I/O processor generated an I/O interrupt.

**Proc %Busy**

The percentage of high frequency samples that found the I/O processor to be busy.

**Note:** The HF sampling counters upon which this value is based may overflow within approximately four minutes. The value can only be meaningful, therefore, when working with monitor sample intervals not longer than three minutes.

*Busy Conditions per SSCH:*

**Channel**

The average number of channel busy conditions encountered for each SSCH begun on this I/O processor.

**Switch**

The average number of switch busy conditions encountered for each SSCH begun on this I/O processor.

**CU**

The average number of control unit busy conditions encountered for each SSCH begun on this I/O processor.

**Device**

The average number of device busy conditions encountered for each SSCH begun on this I/O processor.

**Effect of AVERAGE and INTERIM subcommands:**

None.

## FCX233, TCP/IP Denial of Service Log Screen – TCPDOSLG

**Performance Toolkit Subcommand:** "TCPDOSLG" on page 322.

The TCP/IP Denial of Service Log Screen (FCX233):

FCX233	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS					Perf. Monitor	
TCP/IP Denial of Service Log for Server TCPIP									
Interval	<----- Denial of Service Packet Discards/Sec ----->								
End Time	Smurf	Fraggle	PoDeath	BLAT	Stream	R4P3D	KOD	KOX	SynFlood
>>Mean>>	.000	.000	.000	.000	.000	.000	.000	.000	.000
09:46:04	.000	.000	.000	.000	.000	.000	.000	.000	.000
09:47:03	.000	.000	.000	.000	.000	.000	.000	.000	.000
09:48:04	.000	.000	.000	.000	.000	.000	.000	.000	.000
09:49:03	.000	.000	.000	.000	.000	.000	.000	.000	.000
09:50:03	.000	.000	.000	.000	.000	.000	.000	.000	.000
09:51:03	.000	.000	.000	.000	.000	.000	.000	.000	.000
09:52:03	.000	.000	.000	.000	.000	.000	.000	.000	.000
09:53:03	.000	.000	.000	.000	.000	.000	.000	.000	.000
09:54:03	.000	.000	.000	.000	.000	.000	.000	.000	.000
09:55:03	.000	.000	.000	.000	.000	.000	.000	.000	.000
09:56:03	.000	.000	.000	.000	.000	.000	.000	.000	.000
09:57:03	.000	.000	.000	.000	.000	.000	.000	.000	.000
09:58:03	.000	.000	.000	.000	.000	.000	.000	.000	.000
09:59:03	.000	.000	.000	.000	.000	.000	.000	.000	.000
Command ==> _									
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return				

Figure 127. Layout of TCP/IP Denial of Service Log Screen (FCX233)

The information shown is based on CP monitor APPLDATA domain SAMPLE data, provided by the TCP/IP server machine.

The display shows packet discard rates resulting from the TCP/IP stack's handling of various forms of 'denial of service' attacks. The leftmost column contains the end time (format *hh:mm:ss*) of the interval for which the values were calculated.

The first entry is reserved for the overall averages since the last RESET command, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

**Denial of Service Packet Discards/Sec:** These fields are available with TCP/IP function level 420 and later only; dots will be inserted otherwise.

#### Smurf

Smurf denial of service packet discards, per second.

#### Fraggle

Fraggle denial of service packet discards, per second.

#### PoDeath

Ping-o-Death denial of service packet discards, per second.

The following fields are available with TCP/IP function level 430 and later only; dots will be inserted otherwise.

#### BLAT

BLAT denial of service packet discards, per second.

#### Stream

Stream denial of service packet discards, per second.

**R4P3D**

R4P3D denial of service packet discards, per second.

**KOD**

KOD denial of service packet discards, per second.

**KOX**

KOX denial of service packet discards, per second.

**SynFlood**

SynFlood denial of service packet discards, per second.

Dots can be shown instead of TCP/IP activity values, and either of the strings

Not updated Logged off

inserted, when no new monitor records have been received for an interval, or when the server was found to have been logged off.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.



# FCX234, Processor/Core Configuration Log Screen – PROCCONF

**Performance Toolkit Subcommand:** “PROCCONF” on page 258.

The Processor/Core Configuration Log Screen (FCX234):

```
FCX234      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor

Processor Configuration Log
Physical Core capability: 492.0000  Our level's config. name: .....
Physical Cores (total)   :      128  . higher levels, CP name: .....

Interval    <- Phys Cores Config ->  <----- Virtual Configuration ----->
End Time    Config StandBy Reserved  Count Config StandBy Reserved  CAF
>>Mean>>     57.0      .0      71.0      ..      ..      ..      ..      ..
03:19:23      57.0      .0      71.0      ..      ..      ..      ..      ..
03:20:23      57.0      .0      71.0      ..      ..      ..      ..      ..
03:21:23      57.0      .0      71.0      ..      ..      ..      ..      ..
03:22:23      57.0      .0      71.0      ..      ..      ..      ..      ..
03:23:23      57.0      .0      71.0      ..      ..      ..      ..      ..
03:24:23      57.0      .0      71.0      ..      ..      ..      ..      ..

Command ==> _
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 128. Layout of Processor/Core Configuration Log Screen (FCX234)

The information shown is based on CP monitor SYSTEM domain SAMPLE data.

The display shows physical core configuration by time. The leftmost column contains the end time (format *hh:mm:ss*) of the interval for which the values were calculated.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the Interval End Time column.

**Field Descriptions:**

**Physical Core capability**

The capability of one of the cores in the configuration. This value is used as an indication of the capability of a core relative to the capability of a core in other CPC types or models.

**Physical Cores (total)**

The total number of physical cores in the configuration. It includes all cores in the configured, standby and reserved state.

**Our level's config name**

The configuration name (the user identification in the higher level CP) of our configuration.

**n higher levels, CP name**

Identifies the virtual machine control program that provides the virtual configuration, where *n* indicates the number of higher CP levels.

*By Time Values:*

**Interval End Time**

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

*Phys Cores Config:*

The physical core counts are counts of only the general-purpose (CP) cores.

**Config**

The average number of physical cores found in the configured state.

**StandBy**

The average number of physical cores found in the standby state.

**Reserved**

The average number of physical cores found in the reserved state.

*Virtual Configuration:***Count**

The average number of virtual processors found in this virtual configuration.

**Config**

The average number of virtual processors found in the configured state.

**StandBy**

The average number of virtual processors found in the standby state.

**Reserved**

The average number of virtual processors found in the reserved state.

**CAF**

The average value of the processor capability adjustment factor. It specifies the amount of underlying level-1, level-2, or level-3 configuration capability that is allowed to be used for this level-3 configuration by the virtual machine control program. The maximum level is 1000; lower values indicate a correspondingly lower available processor capability.

**Usage note:**

The term *core* describes a unit that contains one or more processors. All processors of a core have the same processor type. For machine types that do not have the multithreading (MT) facility, the machine has physical processors and the partitions have logical processors. For machine types that have the MT facility, the machine has physical cores, each of which contains two physical processors. A partition that has not activated the MT facility continues to have logical processors. A partition that has activated the MT facility has logical cores, each of which contains one or more logical processors.

**Effect of AVERAGE and INTERIM subcommands:**

None.

## FCX235, LPAR Configuration Log Screen – LPARCONF

**Performance Toolkit Subcommand:** “LPARCONF” on page 221.

The LPAR Configuration Log Screen (FCX235):

FCX235	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor
LPAR Configuration Log for Partition PLB5 (Nr 5)				
Interval	<----- Logical Core Configuration ----->			
End Time	Count	Config	StandBy	Reserved Dedicated Shared CAF
>>Mean>>	32.0	20.0	.0	.0 .0 20.0 363
07:54:00	32.0	20.0	.0	.0 .0 20.0 363
07:55:00	32.0	20.0	.0	.0 .0 20.0 363
07:56:00	32.0	20.0	.0	.0 .0 20.0 363
07:57:00	32.0	20.0	.0	.0 .0 20.0 363
07:58:00	32.0	20.0	.0	.0 .0 20.0 363
07:59:00	32.0	20.0	.0	.0 .0 20.0 363
08:00:00	32.0	20.0	.0	.0 .0 20.0 363
Command ==> _				
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return				

Figure 129. Layout of LPAR Configuration Log Screen (FCX235)

The information shown is based on CP monitor SYSTEM domain SAMPLE data.

The display shows LPAR core configuration by time for the logical partition that is collecting the monitor data. The leftmost column contains the end time (format *hh:mm:ss*) of the interval for which the values were obtained.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the Interval End Time column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### Count

The average of the total number of logical cores of the primary type found, regardless of the state of the logical cores.

#### Config

The average number of logical cores of the primary type found in the configured state.

#### StandBy

The average number of logical cores of the primary type found in the standby state.

#### Reserved

The average number of logical cores of the primary type found in the reserved state.

#### Dedicated

The average number of logical cores of the primary type that were dedicated CPUs.

#### Shared

The average number of logical cores of the primary type that were shared cores.

#### CAF

The average value of the LPAR capability adjustment factor. It specifies the amount of underlying level-1 configuration capability that is allowed to be used for this level-2 configuration by the LPAR hypervisor. The maximum level is 1000; lower values indicate a correspondingly lower available CPU capability.

### Usage note:

The term *core* describes a unit that contains one or more processors. All processors of a core have the same processor type. For machine types that do not have the multithreading (MT) facility, the machine has physical processors and the partitions have logical processors. For machine types that have the MT facility, the machine has physical cores, each of which contains two physical processors. A partition that has not activated the MT facility continues to have logical processors. A partition that has activated the MT facility has logical cores, each of which contains one or more logical processors.

**Effect of AVERAGE and INTERIM subcommands:**

None.

# FCX236, Processor Load and Configuration Logs Menu – PROCMENU

**Performance Toolkit Subcommand:** “PROCMENU” on page 260.

The Processor Load and Configuration Logs Menu (FCX236) shows a list of all of the processor-specific "by time" logs.

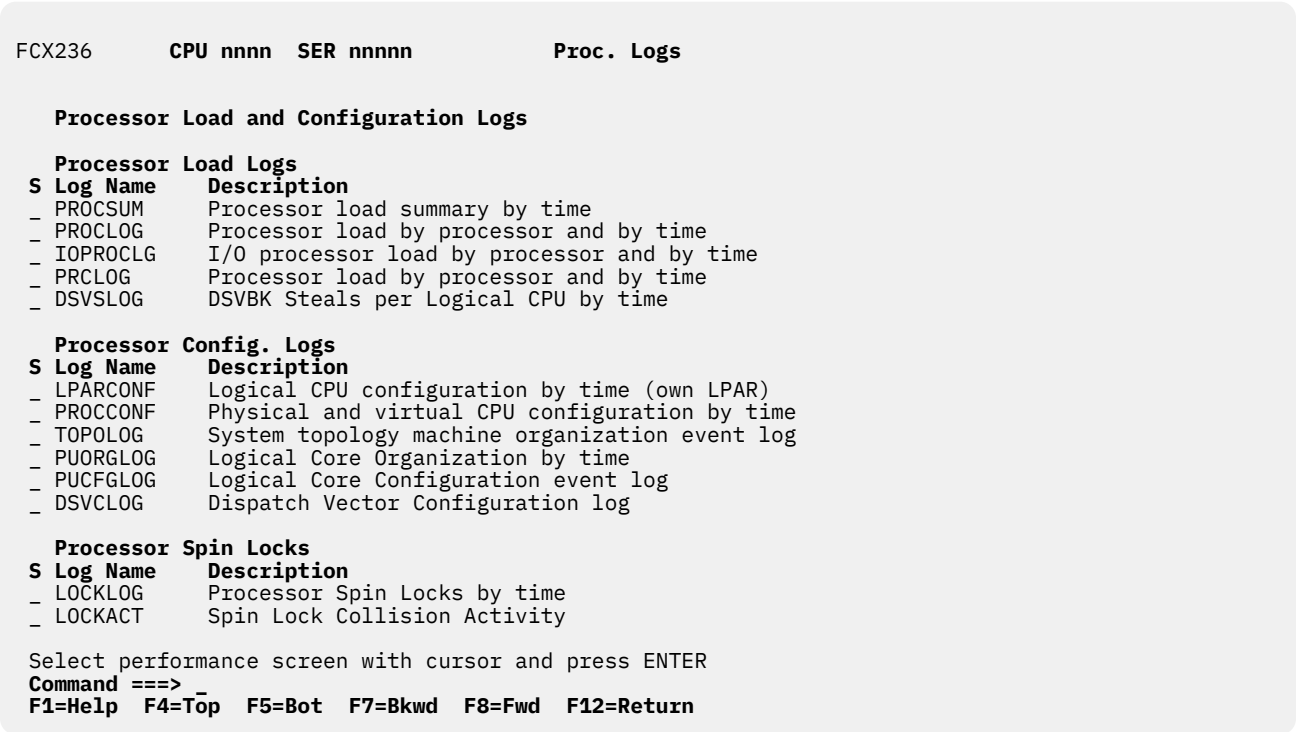


Figure 130. Layout of Processor Load and Configuration Logs Menu (FCX236)

The processor load and configuration logs are based on SYSTEM and PROCESSOR domain SAMPLE data.

## Field Descriptions:

### S

Indicates the cursor selection field. This field is highlighted and made cursor-selectable (set as unprotected) only if the corresponding processor load or configuration log is available.

### Log Name

The log file name, equal to the command with which the log data can be selected for display. The field is highlighted if the corresponding processor load or configuration log is available.

### Description

A description of the performance display.

## FCX237, I/O Priority Queueing Screen – IOPQUEUE

**Performance Toolkit Subcommand:** “IOPQUEUE” on page 202.

The I/O Priority Queueing Screen (FCX237) shows information on the use of the I/O priority queueing facility. User lines are shown highlighted on the screen if some of the I/Os got out-prioritized during the selected interval.

FCX237		CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS				Perf. Monitor
I/O Priority Queueing		HW active						
System minimum		9, maximum 15						
-----								
Userid	Store Subch	Range Type	<-Req. R_Low	Range-> R_High	<-Eff. E_Low	Range-> E_High	Outprior Rate/s	
ACKERK	No	Abs	2	12	5	10	.00	
CFT2NDA	Yes	Rel	0	0	5	5	.00	
IRD00001	No	Rel	10	12	5	5	.00	
IRD00002	No	Rel	8	10	5	5	.00	
IRD00003	No	Rel	10	12	5	5	.00	
IRD00004	No	Rel	10	12	5	5	.00	
IRD00005	No	Rel	10	12	5	5	.00	
IRD00006	No	Rel	10	12	5	5	.00	
IRD00007	No	Rel	9	11	5	5	.00	
IRD00008	No	Rel	10	12	5	5	.00	
IRD00009	No	Rel	10	12	5	5	.00	
IRD00010	No	Rel	10	12	5	5	.00	
IRD00011	No	Rel	10	12	5	5	.00	
Select a user for user details								
Command ==>								
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return								

Figure 131. Layout of I/O Priority Queueing Screen (FCX237)

The information shown is based on USER domain SAMPLE data. It is available only for z/VM 4.3 and later.

The sample shows a screen with the users sorted in ascending sequence of the user identifications, the initial setting (indicated by the underscore characters above the Userid header). Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the 'SORT' subcommand with the appropriate column heading as sort argument (see “SORT” on page 295), or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

For more detailed information on one specific user either:

- Place the cursor on the Userid field of the virtual machine you are interested in and press ENTER, or
- Enter the subcommand 'USER userid' on the command line

### Field Description:

The main header shows whether the I/O priority queueing hardware facility is active or not. Note that CP will simulate the facility to allow prioritizing I/O requests from different virtual machines even when the corresponding hardware facility is not active. System minimum *min*, maximum *max* indicates the minimum and maximum range limits set for the system.

The remaining lines describe the I/O priority queueing settings and activity for individual virtual machines:

#### Userid

User identification of the virtual machine for which the line contains I/O priority queueing information.

#### Store Subch

Awareness of subchannel details by virtual machine.

#### Yes

Guest requested subchannel details which includes I/O priority facility status. This usually indicates that the guest is I/O priority aware, and that it may set its own I/O priorities.

**No**

Guest not aware of I/O priority facility. This indicates that changes to the priority of the virtual machine's I/Os can only be made by CP command.

**Range Type**

The range type set: can be either

**Abs**

If the range is absolute, or

**Rel**

If the range is relative.

**Req. Range R\_Low**

The low value of the requested range.

**Req. Range R\_High**

The high value of the requested range.

**Eff. Range E\_Low**

The low value of the effective range.

**Eff. Range E\_High**

The high value of the effective range.

**Outprior Rate/s**

The number of I/O requests that were outprioritized per second for this user. Note that this value refers to CP's simulation of the I/O priority queueing facility when passing on I/O requests from virtual machines to the hardware. It does *not* reflect the activity of the hardware feature.

**Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages will be shown for all numerical fields.

# FCX238, Idle Users Screen – IDLEUSER

Performance Toolkit Subcommand: “IDLEUSER” on page 199.

The Idle Users Screen (FCX238):

FCX238	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor
-----	Min	-----	Min	-----	Min
-----	Min	-----	Min	-----	Min
Userid	Idle	Userid	Idle	Userid	Idle
AVS	9	BEACH1	1722	BUCKETS	1722
DOUG	1722	GCS	1722	GREGORJL	1722
IMAP3SRV	1722	INTTAPE1	1722	ISPVM	1722
MICHEL	1722	MULTISRV	1138	OPERACCT	1722
PJC	1722	PORTMAP	1722	PVM	2
SFSESA	1722	SNMPD	1722	SQLMACH	2
TPOPER	1722	VMSESVS	1105	VMSESVU	1722
				VSCS	1722
				DATAMOVE	2
				HAMILTJL	1722
				K4SERV	1722
				OPERATOR	1106
				RXAGENT1	1722
				TCPMAINB	1722
				YVETTE	9
				DIRMAINT	2
				HOLDER	1722
				MARINOKA	1722
				OPERSYMP	1722
				SFCM1	1722
				TOOLS	1722

Select a user for user details  
Command ===>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 132. Layout of Idle Users Screen (FCX238)

The display shows a list of all virtual machines that have been idle during the last monitor sample interval and the period during which they have been idle.

The sample shows a screen with the users sorted in ascending userid sequence, the initial setting (indicated by the underscore characters above the Userid headers). The sequence can be changed by means of the 'SORT' subcommand with the appropriate column heading as sort argument (see “SORT” on page 295), or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

For more detailed information on one specific user either:

- Place the cursor on the Userid field of the virtual machine you are interested in and press ENTER, or
- Enter the subcommand 'USER userid' on the command line.

Field Description:

Userid

User identification of idle users.

Min Idle

The period in minutes since the users' last activity or since program initialization, whichever is shorter.

Effect of 'AVERAGE' and 'INTERIM' subcommands: None.



## FCX239, Processor Summary Log Screen – PROCSUM

**Performance Toolkit Subcommand:** “PROCsum” on page 261.

The Processor Summary Log Screen (FCX239) displays a summary of processor activity by time.

FCX239	CPU nnnn SER nnnnn Interval HH:MM:SS - HH:MM:SS Perf. Monitor																					
<----- CPU -----> <----- Spin Lock Collision Activity ----->																						
<--Ratio--> <----- Total -----> <-- Scheduler --> <- Timer Request -> <- SIE/sec -> <-- Diag9C/sec ->																						
Interval	Pct				On-	Colls		AvSpin	Spin	Colls		AvSpin	Spin	Colls		AvSpin	Spin	In-Inter-		by by by		
End Time	Busy	T/V	User	line	/sec	usec	%Busy	/sec	usec	%Busy	/sec	usec	%Busy	/sec	usec	%Busy	struct	cept	DSP	SYN	HVR	UnPk
>>Mean>>	6.3	1.08	.9547	18.0	735.3	.899	.066	122.9	1.880	.023	6.8	.173	.000	5786	4883	.0	.0	.0	.0	.0	.0	17.0
15:02:07	6.3	1.08	.9553	18.0	752.3	.882	.066	125.1	1.849	.023	5.7	.159	.000	5860	4951	.0	.0	.0	.0	.0	.0	17.0
15:03:07	6.2	1.08	.9554	18.0	742.0	.905	.067	127.3	1.887	.024	6.4	.163	.000	5650	4806	.0	.0	.0	.0	.0	.0	17.0
15:04:07	6.2	1.08	.9552	18.0	774.5	.911	.071	136.5	1.938	.026	5.6	.164	.000	5570	4655	.0	.0	.0	.0	.0	.0	17.0
15:05:07	6.2	1.08	.9558	18.0	730.7	.897	.066	136.0	1.844	.025	6.0	.141	.000	5546	4742	.0	.0	.0	.0	.0	.0	17.0
Command ==>																						
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return																						

Figure 133. Layout of Processor Summary Log Screen (FCX239)

The information shown is based on CP monitor SYSTEM domain SAMPLE data.

On this screen, the leftmost column contains the end time (in *hh:mm:ss* format) of the interval for which the values were calculated.

The first entry is reserved for the overall averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the Interval End Time column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

**CPU:** Shows the processor load summary.

#### Pct Busy

The average processor percent busy for all processors.

#### Ratio T/V

The average ratio of total to virtual CPU time for all processors.

#### Ratio User

The ratio of user CPU time to total CPU time.

#### Online

The average number of processors that were online during the interval.

**Spin Lock Collision Activity:** Shows the overall, timer request, and scheduler spin lock activity. The Total columns account for all of the system's formal spin lock activity, whether or not the spun-upon lock is tabulated in LOCKLOG.

#### Total Colls/sec

The total spin lock collisions per second for all processors.

#### Total AvSpin usec

The average number of microseconds spent spinning per lock collision.

#### Total Spin %Busy

Percentage of elapsed time spent in spin lock collisions.

#### Scheduler Colls/sec

The scheduler spin lock collisions per second for all processors.

#### Scheduler AvSpin usec

The average number of microseconds spent spinning per scheduler lock collision.

#### Scheduler Spin %Busy

The percentage of elapsed time spent in scheduler spin lock collisions.

**Timer Request Colls/sec**

Timer request spin lock collisions per second for all processors.

**Timer Request AvSpin usec**

The average number of microseconds spent spinning per timer request collision.

**Timer Request Spin %Busy**

The percentage of elapsed time spent in timer request spin lock collisions.

*SIE/sec*: Shows SIE activity.

**Instruct**

The total number of SIE calls per second for all processors.

**Intercept**

The total number of SIE intercepts per second for all processors.

*Diag 9C/sec*: Shows Diagnose instruction activity.

**Diag 9C byDSP**

The rate of issuing Diagnose X'9C' by HCPDSP.

**Diag 9C bySYN**

The rate of issuing Diagnose X'9C' by HCPSYN.

**Diag 9C byHVR**

The rate of issuing Diagnose X'9C' by HCPHVR.

**UnPk**

The average number of processors that were in unparked state during the interval.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

**Effect of AVERAGE and INTERIM subcommands:**

None.

## FCX240, Virtual Switch Activity Screen – VSWITCH

**Performance Toolkit Subcommand:** “VSWITCH” on page 363.

In the Virtual Switch Activity Screen (FCX240), the lines for Virtual Switches which have been active during the last measuring interval are highlighted on the screen.

Note that besides issuing the VSWITCH subcommand, this screen can also be reached from the TCP/IP menu by selecting "Virtual Switch Activity", or by selecting any row in the General Virtual Switch Description Screen and pressing ENTER. (See “FCX266, General Virtual Switch Description Screen – GVSWITCH” on page 750.)

If you select any row in this screen and press ENTER, you will go to the Extended Virtual Switch Activity Screen for additional data. (See “FCX267, Extended Virtual Switch Activity Screen – EVSWITCH” on page 753.)

```

FCX240      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor
-----
      Q Time  <--- Outbound/s ---> <--- Inbound/s ---> <--- Signals ---> < Intpts > <- Bridge->
      S Out  Bytes <--Packets--> Bytes <--Packets--> <-- issued/s ---> <- Port ->
Addr  Name  Controlr V  Sec  T_Byte T_Pack T_Disc R_Byte R_Pack R_Disc Write Read Sync Rcv/s Pro/s BufUs BufMx
>> System ( 5) << 0 300 68 0 0 24 0 0 0 .3 .0 7.9 1.1 1.1 0 0
60D5 PEGIVL DTCVSW1 0 300 128 1 0 0 0 0 0 .7 .0 9.2 1.2 1.2 0 0
63D0 GLOB1 DTCVSW2 0 300 70 0 0 0 0 0 0 .3 .0 4.4 .6 .6 0 0
63D3 GLOB2 DTCVSW2 0 300 3 0 0 0 0 0 0 .0 .0 7.4 1.1 1.1 0 0
64C0 GLOB1 DTCVSW1 0 300 1 0 0 95 0 0 0 .0 .0 10.0 1.4 1.4 0 0
64C3 GLOB2 DTCVSW1 0 300 0 0 0 0 0 1 .0 .0 7.5 1.1 1.1 0 0

See EVSWITCH / GVSWITCH / GLONACT for other Vswitch reports
Command ==>
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

```

Figure 134. Layout of Virtual Switch Activity Screen (FCX240)

The information shown is based on CP monitor I/O domain SAMPLE data.

The sample shows a screen with Virtual Switches sorted in ascending device number sequence (indicated by the underscore characters above the 'Addr' header). Other possible sort sequences are indicated by dots above the corresponding columns; they can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER. Note that the information extends beyond the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window so as to view the remaining data.

### Field Description:

#### Addr

Real device number. The first line contains the string >> System << instead of a device number and owner identification. It shows the average load figures for all Virtual Switches for which performance data were collected. The string >> System << also includes the total number of Virtual Switches (in parentheses).

#### Name

The name of the Virtual Switch. A Virtual Switch is given a 16-byte identification, or LAN ID, when it is created. The VSWITCH LAN ID consists of two 8-byte character fields: OWNER||NAME. The OWNER is always SYSTEM. This column is the NAME component of the LAN ID.

#### Controlr

The user identification of the virtual machine the device is currently attached to.

#### QSV

The queuestorage value (a value between 1 and 8).

#### Timeout sec

The timeout value set for the Virtual Switch, in seconds.

*Outbound/s:*

**Bytes T\_Byte**

Number of bytes transmitted per second.

**Packets T\_Pack**

Number of packets transmitted per second.

**Packets T\_Disc**

Number of outbound packets discarded per second.

*Inbound/s:*

**Bytes R\_Byte**

Number of bytes received per second.

**Packets R\_Pack**

Number of packets received per second.

**Packets R\_Disc**

Number of inbound packets discarded per second.

*Signals issued/s:*

**Write**

Number of write signals issued per second.

**Read**

Number of read signals issued per second.

**Sync**

Number of sync signals issued per second.

**Int Rcv/s**

Number of interrupts, such as PCI and AI, received per second.

**Int Pro/s**

Number of interrupts, such as PCI and AI, that resulted in input queue processing, per second.

**BufUs**

Cardinal count of the number of configured buffers currently in use.

**BufMx**

Count of the maximum number of asynchronous buffers allowed.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all activity fields. The *QSV* and *Timeout* fields will always contain current information.

## FCX241, VM Resource Manager Screen – VMRM

**Performance Toolkit Subcommand:** “VMRM” on page 359.

In the VM Resource Manager Screen (FCX241), the names of workloads which have been active during the last measuring interval are highlighted on the screen.

FCX241	CPU nnnn	SER nnnnn	Interval	HH:MM:SS	-	HH:MM:SS	Perf. Monitor
-----	-----	-----	-----	-----	-----	-----	-----
VM Resource Manager	Server	Workload	Importance	<-- D-Goal	D-Act	<-- CPU C-Goal	C-Act
							Active Samples
VMRMSVM	WORK1		10	100	...	100	...
VMRMSVM	WORK2		5	50	...	50	...
VMRMSVM	WORK3		1	1	...	1	...
VMRMSVM	WORK4		10	100	100	100	87
VMRMSVM	WORK5		5	50	100	50	43
VMRMSVM	WORK6		1	1	100	1	7
VMRMSVM	WORK7		10	100	100	100	83
VMRMSVM	WORK8		5	50	100	50	41
VMRMSVM	WORK9		1	1	...	1	...

Command ==>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 135. Layout of VM Resource Manager Screen (FCX241)

The information shown is based on CP monitor application data domain SAMPLE data.

The sample shows a screen with the workload information sorted in ascending server ID and workload name sequence (indicated by the underscore characters above the 'Server' header). Other possible sort sequences are indicated by dots above the corresponding columns; they can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

**Note:** If the VMRM server is restarted or a new configuration file is used, the VMRM data will be reset and data will be reported starting at the reset.

### Field Description:

#### Server

User identification of the virtual machine where the VM Resource Manager is running.

#### Workload

The name of the workload for which information is shown. Workloads that have been active (i.e. are associated with virtual machines logged on) in the last monitor sample interval are highlighted.

#### Importance

The importance value assigned to the workload.

#### DASD D-Goal

The target DASD velocity value assigned.

#### DASD D-Act

The actual DASD velocity value calculated. Dots will be inserted if a 'null' value has been received (i.e. if the velocity value could not be calculated because this would have resulted in a division by zero).

#### CPU C-Goal

The target CPU velocity value assigned.

**CPU C-Act**

The actual CPU velocity value calculated. Dots will be inserted if a 'null' value has been received (i.e. if the velocity value could not be calculated because this would have resulted in a division by zero).

**Active Samples**

The number of samples processed for the workload. Monitor data will be generated for a workload only while that workload is considered 'active' with one or more virtual machines that are logged on.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages or totals will be shown for all fields, except in the case of the 'Active Samples' field which will show the total count of active samples in the interval.

# FCX242, Linux Data Selection Menu Screen – LINUXMEN

**Performance Toolkit Subcommand:** “LINUXMEN” on page 216.

The Linux Data Selection Menu Screen (FCX242) is the master menu for selecting Linux performance data displays:

```
FCX242      CPU nnnn  SER nnnnn      Linux Displays      Perf. Monitor

Linux screens selection
S Display   Description
LINUX      RMF PM system selection menu
LXCPU      Summary CPU activity display
LXMEM      Summary memory util. & activity display
LXNETWRK   Summary network activity display

Select performance screen with cursor and press ENTER
Command ==> _
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 136. Layout of Linux Data Selection Menu Screen (FCX242)

All Linux performance displays with summary data are based on APPLDATA monitor records created by the server machines: They can be built and displayed only when APPLDATA sample data collection has been enabled for the server(s), and when the Linux systems themselves have been set up to provide APPLDATA records.

The prerequisites for using RMF PM based Linux performance data are described in the “Monitoring Linux Internal Performance Data” section in [z/VM: Performance Toolkit Guide](#).

**Field Description:**

- S**  
Cursor selection field. The field will be highlighted and made cursor-selectable (set as 'unprotected') only if the corresponding display is available.
- Display**  
The name of the screen, equal to the command with which it can be selected for display. The field will be highlighted if the corresponding display is available.
- Description**  
Description of the performance display or selection menu.

# FCX243, General Linux CPU Utilization Screen – LXCPU

**Performance Toolkit Subcommand:** “LXCPU” on page 226.

The General Linux CPU Utilization Screen (FCX243) will be shown if you enter the LXCPU subcommand with no *linuxid*:

FCX243

CPU	nnnn	SER	nnnnn	Interval	HH:MM:SS	-	HH:MM:SS	Perf.	Monitor
-----									
Total CPU									
-----									
Linux	Virt	Utilization (%)							
-----	-----	-----							
Userid	CPUs	TotCPU	User	Kernel	Nice	IRQ	SoftIRQ	IOWait	Idle Stolen
>System<	1.0	.3	.1	.1	.0	.0	.1	.3	129.4
LXM00001	4	1.4	.0	.3	.0	.0	.9	.6	398.0
LX00001	1	...	...	...	...	...	...	...	...
LX00002	1	...	...	...	...	...	...	...	...
LX00003	1	...	...	...	...	...	...	...	...

Select a highlighted guest for RMF PM Linux details  
Command ==>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 137. Layout of General Linux CPU Utilization Screen (FCX243)

The display shows overall Linux CPU utilization and processes, as perceived by the Linux guest systems. Data can be collected and an output line built only:

- If the Linux system makes the data available to the CP monitor, and
- If the APPLDATA sample domain has been enabled for the Linux guest.

The sample shows a screen with the summary lines from different Linux guests sorted in ascending 'Userid' sequence (indicated by the underscore characters above the 'Linux Userid' header), the initial setting. Other possible sort sequences are indicated by dots above the corresponding columns; they can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

## Field Description

The following field descriptions apply to both the LXCPU and LXCPULOG *userid* displays, with the exception of the Linux Userid field which will be replaced by the field Interval End Time on the LXCPULOG display.

### Linux Userid

The user identification of the Linux guest.

### Virt CPUs

The number of virtual CPUs defined for the guest.

*Total CPU Utilization (%)*: Total CPU utilization, i.e. the **sum** of the utilization for all virtual processors, is shown to allow sorting Linux guests according to their overall CPU load. This logic also corresponds to the manner in which CPU load is indicated on the 'User Resource Usage Screen' FCX112.

See the “Usage Notes” in “LXCPU” on page 226 for information on the accuracy of CPU utilization percentages.

### TotCPU

Percent total CPU utilization.



**User**

Percent user mode CPU utilization.

**Kernel**

Percent kernel mode CPU utilization.

**Nice**

Percent 'nice' mode CPU utilization (modified priority).

**Idle**

Percent CPU idle.

**IRQ**

Percent time in interrupts.

**SoftIRQ**

Percent time in soft IRQs.

**IOWait**

Percent time in I/O Wait.

**Stolen**

Percent time hypervisor in control.

*Processes:***Current Runabl**

The number of runnable processes at sampling time.

**Current Waiting**

The number of processes waiting for I/O at sampling time.

**Current Total**

Total number of processes at sampling time.

**Average Running 1\_Min**

The average number of processes found running during the last minute.

**Average Running 5\_Min**

The average number of processes found running during the last 5 minutes.

**Average Running 15\_Min**

The average number of processes found running during the last 15 minutes.

**Nr of Users**

The number of Linux guests that contributed to the averages shown in this line.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all load fields for all users.

# FCX244, General Linux Memory Utilization Screen – LXMEM

**Performance Toolkit Subcommand:** “LXMEM” on page 229.

The General Linux Memory Utilization Screen (FCX244) will be shown if you enter the LXMEM subcommand with no *linuxid*:

FCX244	CPU nnnn	SER nnnnn	Interval	HH:MM:SS	- HH:MM:SS	Perf. Monitor
-----	Memory Allocation (MB)					-----
Linux	Main		High		Buffers	Cache
Userid	M_Total	%MUsed	H_Total	%HUsed	Shared	Used
>System<	237.0	45.1	.0	.0	.0	14.1
LINUX23	227.8	63.3	.0	.0	.0	9.0
LINUX24	246.3	26.8	.0	.0	.0	19.3

Select a highlighted guest for RMF PM Linux details

Command ==>

F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return

Figure 138. Layout of General Linux Memory Utilization Screen (FCX244)

The display shows overall Linux memory utilization and activity, as perceived by the Linux system. Data can be collected and an output line built only:

- If the Linux system makes the data available to the CP monitor, and
- If the APPLDATA sample domain has been enabled for the Linux guest.

The sample shows a screen with the summary lines from different Linux guests sorted in in ascending 'Userid' sequence (indicated by the underscore characters above the 'Linux Userid' header), the initial setting. Other possible sort sequences are indicated by dots above the corresponding columns; they can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

## Field Description

The following field descriptions apply to both the LXMEM and LXMEMLOG *userid* displays, with the exception of the Linux Userid field which will be replaced by the field Interval End Time on the LXMEMLOG display.

### Linux Userid

The user identification of the Linux guest.

*Memory Allocation (MB):*

### Main M\_Total

Total main memory size, in MBytes. This value does not include some space reserved by the kernel.

### Main %MUsed

Main memory utilization, in percent.

### High H\_Total

High memory size, in MBytes.

### High %HUsed

High memory utilization, in percent.

**Shared**

Memory, in MBytes, that is usable by more than one process. If any part of memory could be used by more than one process it is counted to be shared memory. Value currently set to zero by Linux kernel (2.4 and 2.6).

**Buffers/CaFree**

Memory reserved for buffers and free cache, in MBytes.

**Cache Used**

Memory used for buffers, in MBytes. The Linux buffer cache is a disk cache intended to relieve processes from having to wait for relatively slow disks to retrieve or store data. Linux automatically uses otherwise unused memory for disk buffers. When free memory becomes scarce, buffer frames are automatically released. The buffer cache kernel thread can be tuned using `/proc/sys/vm/bdflush`. Please refer to the file *Documentation/proc.txt* in the Linux kernel source tree.

*Swapping:***Space (MB) S\_Total**

The size of the max. available swap space, in MBytes.

**Space (MB) %SUsed**

Percentage of swap space used.

**Pgs/sec In**

Swap in rate, in 4kB pages per second. Please note that the Linux kernel uses swapping only if there is not enough memory available. If not backed by IBM VM on a S/390 or zSeries mainframe, swapping is normally very bad for performance and application responsiveness.

**Pgs/sec Out**

Swap out rate, in 4kB pages per second. Please note that the Linux kernel uses swapping only if there is not enough memory available. If not backed by IBM VM on a S/390 or zSeries mainframe, swapping is normally very bad for performance and application responsiveness.

*Paging:***Allocates**

Number of physical page allocations per second in the system.

**Faults Major**

Number of major page faults per second for the process. A major page fault occurs if -- from the Linux operating system's perspective -- disk access is involved in handling the page fault.

**Faults Minor**

Number of minor page faults per second for the process. If there is (from the Kernel's perspective) no disk access involved in the page fault (like Copy-on-Write for a shared page), the page fault is called a minor page fault.

*Block IO:***kB/sec Read**

Block I/O data read rate, in kB per second.

**kB/sec Write**

Block I/O data write rate, in kB per second.

**Nr of Users**

The number of Linux guests that contributed to the averages shown in this line.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Averages will be shown as requested.

Usage Note: The Linux internal data collection is not synchronized with CP's data collect activity. That is, The Linux internal data will usually have been obtained for a different period than that used for collecting other VM data. This can lead to discrepancies between the internal Linux values and the VM activity seen from the outside.

# FCX245, General Linux Networking Activity Screen – LXNETWRK

**Performance Toolkit Subcommand:** “LXNETWRK” on page 232.

The General Linux Networking Activity Screen (FCX245) will be shown if you enter the LXNETWRK subcommand with no *linuxid*:

```
FCX245      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor
-----
Linux      <--- Data Transfer/sec ---> <----- Errors/sec ----->
Userid     Inter <--Packets--> <---Bytes---> BadPac Packet <-No Space-> Transmit
>System<   faces P_Recv P_Xmit B_Recv B_Xmit  Recvd  X_Mit Buffer Linux Collisn
LINUX24    3.0  1.528  .332   239   45   .000  .000  .000  .000  .000
```

Select a highlighted guest for RMF PM Linux details

Command ===>

F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return

Figure 139. Layout of General Linux Networking Activity Screen (FCX245)

The display shows overall Linux networking activity, as perceived by the Linux guest systems. Data can be collected and an output line built only:

- If the Linux system makes the data available to the CP monitor, and
- If the APPLDATA sample domain has been enabled for the Linux guest.

The sample shows a screen with the summary lines from different Linux guests sorted in in ascending 'Userid' sequence (indicated by the underscore characters above the 'Linux Userid' header), the initial setting. Other possible sort sequences are indicated by dots above the corresponding columns; they can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

## Field Description

The following field descriptions apply to both the LXNETWRK and LXNETWRK *userid* displays, with the exception of the Linux Userid field which will be replaced by the field Interval End Time on the LXNETWRK display.

### Linux Userid

The user identification of the Linux guest.

### Interfaces

The number of networking interfaces defined.

*Data Transfer/sec:*

### Packets P\_Recv

The number of packets received per second.

### Packets P\_Xmit

The number of packets transmitted per second.

### Bytes B\_Recv

The number of Bytes received per second.

**Bytes B\_Xmit**

The number of Bytes transmitted per second.

*Errors/sec:*

**BadPac Recvd**

Bad packets received per second.

**Packet X\_Mit**

Packet transmit problems per second.

**No Space Buffer**

No space found in Linux buffers, per second.

**No Space Linux**

No space available in Linux, per second.

**Transmit Collisn**

Collisions while transmitting, per second.

**Nr of Users**

The number of Linux guests that contributed to the averages shown in this line.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all load fields.

## FCX246, Linux CPU Load Log Screen – LXCPULOG

**Performance Toolkit Subcommand:** [“LXCPULOG” on page 227.](#)

When requesting the Linux CPU Load Log Screen (FCX246) for a user, the *userid* must be that of a Linux guest machine which has previously been selected for benchmarking, and for which Linux APPLDATA records were collected. It shows a 'by time' log of the selected guest's internal view of CPU utilization.

FCX246	CPU nnnn	SER nnnnn	Interval	HH:MM:SS	- HH:MM:SS	Perf. Monitor						
Linux CPU Load Log for User LINUX24												
			<----- Total CPU ----->					<-----				
Interval	Virt	<----- Utilization (%) ----->					<---- Current					
End Time	CPUs	TotCPU	User	Kernel	Nice	IRQ	SoftIRQ	IOWait	Idle	Stolen	Runabl	Waiting
>>Mean>>	2	26.7	.6	8.1	.0	5.4	12.6	3.5	65.6	104.1	15.2	.0
09:13:19	2	2.3	.2	.7	.0	.6	.8	2.6	182.8	12.3	30	0
09:14:18	2	29.2	.8	9.0	.0	5.8	13.7	4.1	44.7	121.9	24	0
09:15:22	2	26.6	.7	8.2	.0	5.3	12.4	.2	42.1	131.0	12	0
09:16:22	2	29.2	.8	9.0	.0	5.7	13.6	4.0	45.3	121.5	19	0
09:17:22	2	28.6	.8	8.9	.0	5.6	13.3	4.2	41.5	125.7	21	0
09:18:22	2	28.3	.7	8.9	.0	5.7	12.9	6.1	45.8	119.7	31	0

**Command** ==> **F1=Help** **F4=Top** **F5=Bot** **F7=Bkwd** **F8=Fwd** **F12=Return**

Figure 140. Layout of Linux CPU Load Log Screen (FCX246)

The information shown is based on APPLDATA domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

## Interval End Time

The end time of the monitor interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'Linux Userid' field (and the 'Nr of users' field which is omitted from this screen), all data columns are identical to the ones of the 'General Linux CPU Utilization' (LXCPU) display. Please see [“FCX243, General Linux CPU Utilization Screen – LXCPU”](#) on page 712 for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX247, Linux Memory Utilization Log Screen – LXMEMLOG

**Performance Toolkit Subcommand:** “LXMEMLOG” on page 230.

When requesting the Linux Memory Utilization Log Screen (FCX247) for a user, the *userid* must be that of a Linux guest machine which has previously been selected for benchmarking, and for which Linux APPLDATA records were collected. It shows a 'by time' log of the selected guest's internal view of memory utilization.

FCX247	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor
Linux Memory Utilization Log for User LINUX24				
	Memory Allocation (MB)			
Interval	<--- Main --->	<--- High --->	Buffers	Cache <---Space (MB)--> <
End Time	M_Total %MUsed	H_Total %HUsed	Shared /CaFree	Used S_Total %SUsed
>>Mean>>	246.3 29.6	.0 .0	.0 32.9	14.5 256.0 .0
17:11:29	246.3 29.6	.0 .0	.0 32.9	14.5 256.0 .0
17:12:29	246.3 29.6	.0 .0	.0 32.9	14.5 256.0 .0
17:13:29	246.3 29.6	.0 .0	.0 32.9	14.5 256.0 .0
17:14:29	246.3 29.6	.0 .0	.0 32.9	14.5 256.0 .0
17:15:29	246.3 29.6	.0 .0	.0 32.9	14.5 256.0 .0
17:16:29	246.3 29.6	.0 .0	.0 32.9	14.5 256.0 .0
17:17:29	246.3 29.6	.0 .0	.0 32.9	14.5 256.0 .0
17:18:30	246.3 29.6	.0 .0	.0 32.9	14.5 256.0 .0
17:19:30	246.3 29.6	.0 .0	.0 32.9	14.5 256.0 .0
17:20:30	246.3 29.6	.0 .0	.0 32.9	14.5 256.0 .0
17:21:30	246.3 29.6	.0 .0	.0 32.9	14.5 256.0 .0
17:22:30	246.3 29.6	.0 .0	.0 32.9	14.5 256.0 .0
17:23:30	246.3 29.6	.0 .0	.0 32.9	14.5 256.0 .0
Command ===>				
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return				

Figure 141. Layout of Linux Memory Utilization Log Screen (FCX247)

The information shown is based on APPLDATA domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column. Note that the information extends beyond the 80 columns of an ordinary screen: use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window so as to view the remaining data.

### Field Description:

#### Interval End Time

The end time of the monitor interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'Linux Userid' field (and the 'Nr of users' field which is omitted from this screen), all data columns are identical to the ones of the 'General Linux Memory Utilization' (LXMEM) display. Please see “FCX244, General Linux Memory Utilization Screen – LXMEM” on page 714 for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX248, Linux Network Activity Log Screen – LXNETLOG

**Performance Toolkit Subcommand:** “LXNETLOG” on page 231.

When requesting a Linux Network Activity Log (FCX248) for a user, the *userid* must be that of a Linux guest machine which has previously been selected for benchmarking, and for which Linux APPLDATA records were collected. It shows a 'by time' log of the selected guest's internal view of network activity.

```
FCX248      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS      Perf. Monitor
Linux Network Activity Log for User LINUX24

Interval Inter <--- Data Transfer/sec ---> <----- Errors/sec ----->
End Time faces P_Recv P_Xmit B_Recv B_Xmit BadPac Packet <-No Space-> Transmit
>>Mean>> 3.0 1.077 .358 230 48 .000 .000 .000 .000 .000
16:56:26 3.0 1.198 .333 216 45 .000 .000 .000 .000 .000
16:57:26 3.0 .782 .366 154 49 .000 .000 .000 .000 .000
16:58:27 3.0 .982 .333 242 45 .000 .000 .000 .000 .000
16:59:27 3.0 .899 .316 222 43 .000 .000 .000 .000 .000
17:00:27 3.0 1.179 .432 242 58 .000 .000 .000 .000 .000
17:01:27 3.0 1.096 .432 227 58 .000 .000 .000 .000 .000
17:02:27 3.0 1.146 .432 242 58 .000 .000 .000 .000 .000
17:03:27 3.0 1.113 .432 234 58 .000 .000 .000 .000 .000
17:04:27 3.0 .965 .349 212 47 .000 .000 .000 .000 .000
17:05:28 3.0 1.065 .366 234 49 .000 .000 .000 .000 .000
17:06:28 3.0 .781 .316 191 43 .000 .000 .000 .000 .000
17:07:28 3.0 .998 .316 228 43 .000 .000 .000 .000 .000
17:08:28 3.0 1.096 .316 233 43 .000 .000 .000 .000 .000

Command ==>
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return
```

Figure 142. Layout of Linux Network Activity Log Screen (FCX248)

The information shown is based on APPLDATA domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'Linux Userid' field (and the 'Nr of users' field which is omitted from this screen), all data columns are identical to the ones of the 'General Linux Networking Activity' (LXNETWRK) display. Please see “FCX245, General Linux Networking Activity Screen – LXNETWRK” on page 716 for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.



# FCX249, SCSI Device – SCSI

**Performance Toolkit Subcommand:** “SCSI” on page 282.

The SCSI Device report (FCX249):

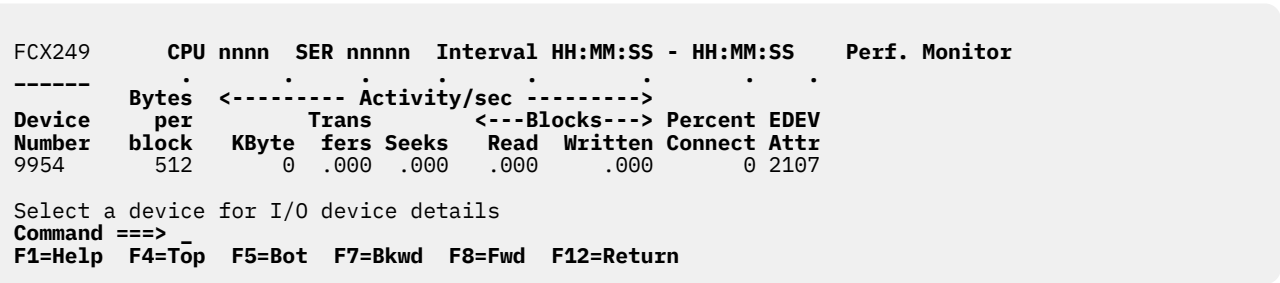


Figure 143. Layout of SCSI Device report (FCX249)

The sample shows a report with SCSI devices sorted in ascending device address sequence (indicated by the underscore characters above the Device Number header), the initial sequence. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the SORT subcommand with the string shown in the lowest line of the column heading as the sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing Enter.

**Field Descriptions:**

**Device Number**

The real emulated FBA device (EDEV) number.

**Bytes per block**

The number of bytes in a block for this DASD.

**KByte**

The calculated number of kilobytes per second transfer rate of the device, calculated from written and read blocks, and bytes per block data.

**Transfers**

The number of data transfers per second to and from the device.

**Seeks**

The number of seek operations per second associated with the device. Note that the transfer rate is only meaningful for devices that support discrete seeks. SCSI DASDs do not support discrete seeks and return a value of 0 for this field.

**Blocks Read**

The number of blocks read per second during the selected interval.

**Blocks Written**

The number of blocks written per second during the selected interval.

**Percent Connect**

The percentage of time that the device was busy.

**EDEV Attr**

The value of this device's EDEV attribute.

**Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages are shown for all load fields.

# FCX250, SCSI Resource Utilization Log – SCSILOG

**Performance Toolkit Subcommand:** “SCSILOG” on page 283.

The SCSI Resource Utilization Log report (FCX250) shows a "by time" log of SCSI performance data for the selected device (which must have been previously selected for benchmarking).

FCX250	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor
Resource Utilization Log for SCSI Device 9954				
Interval	Bytes	<----- Activity/sec ----->		
End Time	per	Trans	<---Blocks---	Percent EDEV
>>Mean>>	block	KByte fers Seeks	Read Written Connect Attr	
11:57:39	512	19693 1769 .000	7902 31484	47 2107
11:58:39	512	4995 80.45 .000	4992 4998	4 2107
11:59:39	512	5004 80.57 .000	5001 5007	3 2107
12:00:39	512	7505 120.9 .000	7501 7510	4 2107
12:01:39	512	7505 120.9 .000	7501 7510	4 2107
12:02:39	512	18782 1547 .000	6739 30826	40 2107
12:03:39	512	34269 3300 .000	11256 57283	83 2107
12:04:39	512	32465 3357 .000	8879 56052	83 2107
12:05:39	512	34163 3283 .000	11248 57079	83 2107
12:06:39	512	28747 3022 .000	8412 49083	85 2107
12:07:39	512	28443 3084 .000	7801 49085	86 2107
12:08:39	512	26953 3115 .000	8000 45906	85 2107
Command ===>				
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return				

Figure 144. Layout of SCSI Resource Utilization Log report (FCX250)

The information shown is based on MONITOR domain and I/O domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the Interval End Time column.

**Field Descriptions:**

**Interval End Time**

The end time of the monitor interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the Interval End Time field, which replaces the Device Number field, all data columns are identical to those in “FCX249, SCSI Device – SCSI” on page 721

**Effect of AVERAGE and INTERIM subcommands:**

None.

## FCX251, QDIO Activity Screen – QDIO

Performance Toolkit Subcommand: “QDIO” on page 264.

The QDIO Activity Screen (FCX251):

```

FCX251      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor

-----
Device  <-Attached-->  QDIO  <Queues>  <----- SIGA Instructions/sec ----->  <---- Data Transfer/sec ---->  <- Intrpts/sec -->
Number  Userid  Vdev  Fmt  In  Out  Gst-r  Gst-w  Gst-s  CP-r  CP-w  CP-s  WBuffer WBytes RBuffer RBytes  Total  Input  Dis-
>> All QDIO << -----
4113 DTCVSW2 4113 QDIO 1 1 .000 .000 .000 .000 .269 .000 .333 255 .238 39 .269 .269 .000
503D DTCENS1 503D INMN 1 1 .000 .000 .000 .000 .283 .000 .350 268 .250 41 .250 .250 .000
F020 DTCVSW2 F020 HBRDG 2 2 .000 .000 .000 .000 .000 .000 .000 0 .000 0 .019 .019 .000
F620 DTCVSW2 F620 HBRDG 2 2 .000 .000 .000 .000 .000 .000 .000 0 .000 0 .024 .024 .000
F623 DTCVSW2 F623 HBRDG 2 2 .000 .000 .000 .000 .000 .000 .000 0 .000 0 .000 .000 .000

Select a device for I/O device details or UQDIO for user interrupts
Command ==>
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

```

Figure 145. Layout of QDIO Activity Screen (FCX251)

This display shows QDIO device activity. It is based on I/O domain SAMPLE data.

**Note:** Because of the limits in the FCP protocol, if the selected QDIO device is an FCP device, the count fields in this screen might not be accurate.

The sample shows a screen with the data lines for different QDIO devices sorted in ascending device number sequence (the initial setting), indicated by the underscore characters above the Device Number header. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the SORT command with the string shown in the lowest line of the column heading as the sort argument or by positioning the cursor on the header of the field that is to be sorted and then pressing Enter.

### Field Descriptions:

The following field descriptions apply to the QDIO and QDIOLOG displays, with the exception of the Device Number field, which is replaced by the Interval End Time field on the QDIOLOG display.

#### Device Number

The QDIO device number.

#### Attached Userid

The user ID of the virtual machine to which the QDIO device is dedicated.

#### Attached Vdev

The virtual device number that is assigned to the QDIO device.

#### QDIO Fmt

The format type of the QDIO device:

##### Hiper

HiperSockets

##### FCP

Fibre Channel Protocol

##### QDIO

Queued direct I/O (OSA, OSD, OSE)

##### OSN

QDIO - OSA for NCP device

##### INMN

QDIO - OSM CHPID

**IEDN**

QDIO - OSX CHPID

**HBRDG**

IQD Hiper-Bridge CHPID

**HIEDN**

IQD Hiper-IEDN CHPID

**????**

Unrecognized format.

**Queues In**

The number of input queues.

**Queues Out**

The number of output queues.

*SIGA Instructions/sec:***By Guest Gst-r**

The number of SIGA-r instructions executed per second by the guest.

**By Guest Gst-w**

The number of SIGA-w instructions executed per second by the guest.

**By Guest Gst-s**

The number of SIGA-s instructions executed per second by the guest.

**By CP CP-r**

The number of SIGA-r instructions executed per second by CP.

**By CP CP-w**

The number of SIGA-w instructions executed per second by CP.

**By CP CP-s**

The number of SIGA-s instructions executed per second by CP.

*Data Transfer:***Written WBuffer**

The number of QDIO buffers written to the adapter per second.

**Written WBytes**

The number of bytes written to the adapter per second.

If the current device is an FCP device, Performance Toolkit displays - - - in this field.

**Read RBuffer**

The number of QDIO buffers read from the adapter per second.

**Read RBytes**

The number of bytes read from the adapter per second.

If the current device is an FCP device, Performance Toolkit displays - - - in this field.

*Int/sec:***Total**

The total number of interrupts (including PCI and AI) per second for the device.

**Input Req'd**

The number of interrupts (including PCI and AI) per second that required a response and caused an input operation.

**Discarded**

The number of interrupts (including PCI and AI) per second that were discarded because they did not require a response.

**Effect of AVERAGE and INTERIM subcommands:** Averages are shown as requested for all activity fields. Current data are inserted for owner, vdev, format, and number of queues.

## FCX252, QDIO Activity Log Screen – QDIOLOG

**Performance Toolkit Subcommand:** “QDIOLOG” on page 265.

The QDIO Activity Log Screen (FCX252) shows a 'by time' log of the selected device's activity. Note that the QDIO device selected must have been previously selected for benchmarking.

FCX252	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor
QDIO Activity Log for Device 2308				
Interval	<-Attached-->	QDIO	<Queues>	<----- SIGA Instructions/sec ----->
End Time	Userid	Vdev	Fmt	<-- By Guest ---> <---- By CP ---->
>>Mean>>	LANTZY	2308	QDIO	Gst-r Gst-w Gst-s CP-r CP-w CP-s
12:24:34	LANTZY	2308	QDIO	.181 .038 2.030 .000 .038 .000
12:25:34	LANTZY	2308	QDIO	.033 .033 1.433 .000 .033 .000
12:26:34	LANTZY	2308	QDIO	.017 .000 1.217 .000 .000 .000
12:27:34	LANTZY	2308	QDIO	.333 .067 2.817 .000 .067 .000
12:28:34	LANTZY	2308	QDIO	.500 .400 4.683 .000 .400 .000
12:29:34	LANTZY	2308	QDIO	.317 .050 2.683 .000 .050 .000
12:30:34	LANTZY	2308	QDIO	.017 .000 1.217 .000 .000 .000
12:31:34	LANTZY	2308	QDIO	.067 .017 1.483 .000 .017 .000
12:32:34	LANTZY	2308	QDIO	.000 .000 1.150 .000 .000 .000
12:33:35	LANTZY	2308	QDIO	.033 .000 1.283 .000 .000 .000
12:34:34	LANTZY	2308	QDIO	.167 .000 1.800 .000 .000 .000
12:35:34	LANTZY	2308	QDIO	.200 .000 1.967 .000 .000 .000
12:36:34	LANTZY	2308	QDIO	.533 .033 3.483 .000 .033 .000
12:36:34	LANTZY	2308	QDIO	.067 .000 1.417 .000 .000 .000
Command ==>				
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return				

Figure 146. Layout of QDIO Activity Log Screen (FCX252)

The information shown is based on I/O domain SAMPLE data.

**Note:** Because of the limits in the FCP protocol, if the selected QDIO device is an FCP device, the count fields in this screen may not be accurate.

The first entry is reserved for averages since the last RESET command, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Description:

#### Interval End Time

The end time of the monitor interval for which the values were calculated.

In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'Device Number' field, all data columns are identical to the ones of the 'QDIO Activity' (QDIO) display. Please see “FCX251, QDIO Activity Screen – QDIO” on page 723 for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX253, Storage Utilization Log Screen – STORLOG

**Performance Toolkit Subcommand:** “STORLOG” on page 303.

The Storage Utilization Log Screen (FCX253):

FCX253	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor									
<----- Storage Utilization (Page Frames) ----->														
<----- DPA ----->				Stor	<----- Locked ----->									
Interval	<--Pageable-->	Nonpgb	Util	Save	Track	<-LOCK	REAL->	<-SXS	Alias->	Diag	Re-	Resid		
End Time	<2GB	>2GB	<2GB	%	Areas	Cache	<2GB	>2GB	Total	LOCK	X'98'	served	Shared	
>>Mean>>	493420	77482k	30868	101	127	12936	0	0	134	0	540	3864576	8094	
04:39:19	493413	77482k	30875	97	102	15134	0	0	122	0	540	3864576	8043	
04:44:19	493413	77482k	30875	98	339	14899	0	0	142	0	540	3864576	8066	
04:49:19	493413	77482k	30875	98	136	14892	0	0	141	0	540	3864576	8062	
04:54:19	493419	77482k	30869	101	102	13742	0	0	136	0	540	3864576	8056	
04:59:19	493420	77482k	30868	101	166	13718	0	0	136	0	540	3864576	8055	
05:04:19	493421	77482k	30867	100	111	13718	0	0	137	0	540	3864576	8057	
05:09:19	493421	77482k	30867	101	103	13057	0	0	135	0	540	3864576	7928	
05:14:19	493420	77482k	30868	101	109	14098	0	0	138	0	540	3864576	8883	
05:19:19	493420	77482k	30868	102	186	12973	0	0	135	0	540	3864576	7928	
05:24:19	493422	77482k	30866	101	102	12927	0	0	140	0	540	3864576	7933	

See AVAILLOG, DEMNDLOG and PAGELLOG for more storage related information

Command ==> \_  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 147. Layout of Storage Utilization Log Screen (FCX253)

The information shown is based on CP monitor SYSTEM and STORAGE domain SAMPLE data.

The display shows general storage utilization data by time. The leftmost column contains the end time (format hh:mm:ss) of the monitor interval for which the values were calculated.

The first entry is reserved for the overall report period averages. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry you can find the relevant interval in the header line.

**DPA:** (Shows page frame counts in the dynamic paging area)

#### Pageable <2GB

The number of pageable DPA frames below the 2GB line.

#### Pageable >2GB

The number of pageable DPA frames above the 2GB line.

#### Nonpgb <2GB

The number of DPA frames below the 2GB line that have been made non-pageable. This includes pages locked for I/O, CP LOCK command, virtual system lock functions, etc.

#### Stor Util %

The percentage of pageable DPA used by the working sets of in-queue (i.e. non dormant) users.

#### Save Areas

The number of page frames in use for save areas.

#### Track Cache

The number of real storage page frames used by the full track cache. (See “MDCSTOR” on page 234 log for a <2GB and >2GB page count split.)

**Locked:**

#### LOCK REAL <2GB

The number of pages locked below the 2GB line by the CP LOCK REAL command.

**LOCK REAL >2GB**

The number of pages locked above the 2GB line by the CP LOCK REAL command.

**SXS Alias Total**

The total number of System Execution Space alias pages locked in logical.

**SXS Alias LOCK**

The number of System Execution Space alias pages locked in logical via the CP LOCK command. This is a subset of the 'SXS Alias Total' pages.

**Diag X'98'**

The number of pages locked by Diagnose X'98'.

**Reserved**

The number of reserved pages.

**Resid Shared**

The number of resident shared pages.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX254, Available List Log Screen – AVAILLOG

**Performance Toolkit Subcommand:** “AVAILLOG” on page 132.

### The Available List Log Screen (FCX254):

FCX254	CPU nnnn		SER nnnnn		Interval		HH:MM:SS - HH:MM:SS		Perf. Monitor		
<----->						Available List Management					
<---- Thresholds ---->						<----- Page Frames ----->					
Interval	<---Low--->		<--High-->		<Available>		<Obtains/s>		<Returns/s>		<-Times
End Time	<2GB	>2GB	<2GB	>2GB	<2GB	>2GB	<2GB	>2GB	<2GB	>2GB	<-Empty
>>Mean>>	276	....	554	....	4694	....	185.0	...	170.2	...	911
08:30:20	372	....	652	....	9022	....	428.4	...	90.8	...	0
08:35:20	372	....	652	....	9685	....	46.4	...	14.0	...	0
08:40:20	372	....	652	....	9490	....	292.9	...	37.4	...	0
08:45:20	372	....	652	....	9129	....	85.7	...	14.6	...	0
08:50:20	354	....	634	....	3855	....	120.5	...	36.3	...	0
08:55:20	286	....	566	....	2928	....	81.1	...	30.9	...	0
09:00:20	286	....	566	....	3057	....	26.3	...	15.5	...	0
09:05:20	286	....	566	....	3327	....	204.2	...	37.8	...	0
09:10:20	286	....	566	....	3111	....	239.8	...	32.4	...	0
09:15:20	340	....	620	....	1406	....	408.1	...	79.1	...	4
09:20:20	308	....	588	....	2836	....	362.8	...	172.7	...	0
09:25:20	387	....	667	....	7478	....	294.6	...	136.1	...	27
09:30:20	358	....	638	....	10068	....	1530	...	429.5	...	41
09:35:20	514	....	794	....	757	....	1627	...	1014	...	248
09:40:20	540	....	820	....	1119	....	1175	...	721.6	...	570
See DEMNDLOG, PAGELOG and STORLOG for more storage related information											
Command ==>											
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F10=Left	F11=Right	F12=Return				

Figure 148. Layout of Available List Log Screen (FCX254)

This screen applies only to monitor data when running Performance Toolkit against z/VM 6.2 and earlier data.

The information shown is based on CP monitor SYSTEM and STORAGE domain SAMPLE data.

The display shows available list status and activity by time. The leftmost column contains the end time ((in hh:mm:ss format) of the monitor interval for which the values were calculated.

The first entry is reserved for the overall report period averages. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Descriptions:

## Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry you can find the relevant interval in the header line.

**Thresholds:** Lists the minimum frame counts below which a replenishment scan should be started (low), or above which no further scan is necessary (high).

### Low <2GB

The low threshold for frames below 2GB.

**Low >2GB**

The low threshold for frames above 2GB.

## High <2GB

The high threshold for frames below 2GB.

**High >2GB**

The high threshold for frames above 2GB.

**Page Frames:** Shows the number of frames available, and the rate at which they are obtained and returned.

**Available <2GB**

The number of frames available on the below 2GB available list.



**Available >2GB**

The number of frames available on the above 2GB available list.

**Obtains/s <2GB**

The rate per second at which frames were obtained from the below 2GB available list.

**Obtains/s >2GB**

The rate per second at which frames were obtained from the above 2GB available list.

**Returns/s <2GB**

The rate per second at which frames were returned to the below 2GB available list.

**Returns/s >2GB**

The rate per second at which frames were returned to the above 2GB available list.

**Times Empty <2GB**

The number of times a <2G frame request could not be satisfied from the <2G available list because it was empty during the interval.

**Times Empty >2GB**

The number of times a >2G frame request could not be satisfied from the >2G available list because it was empty during the interval.

*Replenishment:* Shows the system's activity for replenishing the available list when the available frame count fell below the low threshold. So as not to impact active users more than absolutely necessary, the demand scan will try to obtain page frames from dormant users first.

All values are totals for the periods displayed, i.e. for the overall reporting period in the case of the >>Mean>> line, or for the last monitor sample interval, or the 'by time' period, in the case of the detail lines.

All of the following field show totals for scans ended after, and pages obtained from, the different user lists. See [“FCX259, Demand Scan Log Screen – DEMNDLOG” on page 736](#) for more details on demand scan operation.

**Scan1 Compl**

The number of times the replenishment scan stopped after scan 1 because it could obtain a sufficient number of page frames.

**Scan1 Pages**

The number of page frames obtained during scan 1.

**Scan2 Compl**

The number of times the replenishment scan stopped after scan 2 because it could obtain a sufficient number of page frames.

**Scan2 Pages**

The number of page frames obtained during scan 2.

**Em-Scan Compl**

The number of times the replenishment scan stopped after the emergency scan because it could obtain a sufficient number of page frames.

**Em-Scan Pages**

The number of page frames obtained during the emergency scan.

**Scan Fail**

The number of times the replenishment scan failed, i.e. it could not obtain a sufficient number of frames even after the emergency scan.

**Perct Emerg Scan**

The percentage of scans that did not end after the first pass. This is a measure of how difficult it is for the system to keep a sufficient number of page frames on the available list: low numbers are good, high numbers indicate insufficient storage.

**Effect of AVERAGE and INTERIM subcommands:**

None.

## FCX255, QEBSM Activity Screen – QEBSM

**Performance Toolkit Subcommand:** “QEBSM” on page 266.

The QEBSM Activity Screen (FCX255):

FCX255	CPU nnnn	SER nnnnn	Interval	HH:MM:SS	- HH:MM:SS	Perf. Monitor	
-----	<-----	-----	Executed/s by z/VM	-----	>	<-QIOBK/QSD-->	<P
Device	<-- SQBS -->	<-- EQBS -->	<---- SIGA Instr ---->	<Lock Fails/s>	<-		
Number	SCompl	SPartl	ECompl	EPartl	SIGA-wt	SIGA-rt	SIGA-lt
>>All<<	.817	.000	1.117	.000	.217	.000	.000
800A	.733	.000	1.167	.000	.100	.000	.000
CE02	.900	.000	1.067	.000	.333	.000	.000

Select a device for I/O device details or UQDIO for user interrupts  
 Command ==>  
 F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return

Figure 149. Layout of QEBSM Activity Screen (FCX255)

This display shows QEBSM device specific activity. It is based on I/O domain SAMPLE data. See also “FCX251, QDIO Activity Screen – QDIO” on page 723 for general QDIO device related information, and “FCX257, User QDIO Activity Screen – UQDIO” on page 733 for QDIO guest interrupt handling.

The sample shows a screen with the data lines for different QEBSM devices sorted in ascending device number sequence (the initial setting), indicated by the underscore characters above the 'Device Number' header. Other possible sort sequences are indicated by dots above the corresponding columns – they can be set using the 'SORT' command with the column heading as sort argument (see “SORT” on page 295), or by placing the cursor on the header for the field which is to be sorted and then pressing the ENTER key.

The first line contains the prefix >>All<< instead of a device number. It provides the average load figures for all QEBSM devices.

Use the scrolling PF-keys or commands to view data not visible on the initial screen, or use the locate command if you want to search for a specific device.

### Field Descriptions:

The following field descriptions apply to both the QEBSM and 'QEBSMLOG devno' displays, with the exception of the 'Device Number' field which will be replaced by the field 'Interval End Time' on the QEBSMLOG display.

#### Device Number

The QEBSM device number.

Executed/s by z/VM:

#### SQBS SCompl

The number of SQBS instructions per second that could not be executed by microcode, and that had to be completely processed by CP.

#### SQBS SPartl

The number of SQBS instructions per second that could not be completely executed by microcode, and that had to be partially processed by CP.

#### EQBS ECompl

The number of EQBS instructions per second that could not be executed by microcode, and that had to be completely processed by CP.

**EQBS EPartI**

The number of EQBS instructions per second that could not be completely executed by microcode, and that had to be partially processed by CP.

**SIGA-wt**

The number of SIGA-wt instructions executed per second by CP.

**SIGA-rt**

The number of SIGA-rt instructions executed per second by CP.

**SIGA-lt**

The number of SIGA-lt instructions executed per second by CP.

*QIOBD/QSD Lock Fails/s:*

**General**

The number of unsuccessful attempts per second to obtain a general QIOBD/QSD lock for serializing input or output.

**Output**

The number of unsuccessful attempts per second to obtain a QIOBD/QSD output lock. This lock is used for serializing a format 1 FCP device used for SCSI communications, but not for OSA express or Hipersocket devices.

*Pin Probl Rate/s:*

**Pin**

The number of 'pin' requests per second that could not be handled by the microcode, and that had to be processed by CP.

**UnPin**

The number of 'unpin' requests per second that could not be handled by the microcode, and that had to be processed by CP.

*FOBLK:*

**Resrce Exh/s**

The number of times per second FOBLK resources were exhausted.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all activity fields.

## FCX256, QEBSM Activity Log Screen – QEBSMLOG

**Performance Toolkit Subcommand:** “QEBSMLOG” on page 267.

The QEBSM Activity Log Screen (FCX256) shows a 'by time' log of the selected device's QEBSM specific activity. Note that the QEBSM device selected must have been previously selected for benchmarking.

FCX256	CPU nnnn SER nnnnn Interval HH:MM:SS - HH:MM:SS							Perf. Monitor	
	<----- Executed/s by z/VM ----->							<-QIOBK/QSD-->	
Interval	<-- SQBS -->		<-- EQBS -->		<---- SIGA Instr ---->			<Lock Fails/s>	
End Time	SCompl	SPartl	ECompl	EPartl	SIGA-wt	SIGA-rt	SIGA-lt	General	Output
>>Mean>>	1.189	.000	1.867	.000	.119	.000	.000	.000	.000
13:37:03	2.367	.000	3.500	.000	.100	.000	.000	.000	.000
13:37:33	1.200	.000	1.900	.000	.133	.000	.000	.000	.000
13:38:03	.733	.000	1.167	.000	.100	.000	.000	.000	.000
13:38:33	.500	.000	.900	.000	.133	.000	.000	.000	.000
13:39:03	2.733	.000	3.933	.000	.033	.000	.000	.000	.000
13:39:33	1.267	.000	2.000	.000	.133	.000	.000	.000	.000
13:40:03	.533	.000	.967	.000	.133	.000	.000	.000	.000
13:40:33	.733	.000	1.333	.000	.167	.000	.000	.000	.000
13:41:03	.633	.000	1.100	.000	.133	.000	.000	.000	.000
Command	====>	_							
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F10=Left	F11=Right	F12=Return		

Figure 150. Layout of QEBSM Activity Log Screen (FCX256)

The information shown is based on I/O domain SAMPLE data.

The first entry is reserved for overall report averages since the last RESET command, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field, which replaces the 'Device Number' field, all other data columns are identical to those in the 'QEBSM Activity' display. See “FCX255, QEBSM Activity Screen – QEBSM” on page 730 for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX257, User QDIO Activity Screen – UQDIO

**Performance Toolkit Subcommand:** “UQDIO” on page 342.

The User QDIO Activity Screen (FCX257) contains detailed information on each user's QDIO related interrupt and instruction intercept activity during the last measuring interval (for current data), or for the current INTERIM or REPORT period (when INTERIM or AVERAGE data has been selected). The information shown is based on CP monitor USER domain SAMPLE data, but the screen only shows users that had some QDIO related activity since logging on – users with any QDIO related activity during the selected interval are shown highlighted on the screen.

```

FCX257      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor
-----
          <Guest Instruct>
          <-Intercepts/s->      Alert <- Adapter Interruptions/s ->
          TVS      SVS      /sec <Delivered to Guest> Stacked      Nr of
Userid      <-Intercepts/s->      /sec ViaTPI      InSIE      ByCP      Users
>System<      .0      6.7      .0      .0      .0      44.3      18
User Class Data:
TCP/IP      .0      .0      .5      .0      .5      .0      1
User Data:
LXM00001      .0      .0      .0      .0      .0      232.2
LX00014      .0      .0      .0      .0      .0      3.9
LX00015      .0      .0      .0      .0      .0      3.4
LX00016      .0      20.2      .0      .0      .0      20.2
LX00017      .0      20.2      .0      .0      .0      20.2
LX00018      .0      19.8      .0      .0      .0      19.8
LX00019      .0      19.9      .0      .0      .0      19.9
LX00020      .0      20.2      .0      .0      .0      20.2
LX00021      .0      20.3      .0      .0      .0      20.3
LX00022      .0      .0      .0      .0      .0      .0
LX00023      .0      .0      .0      .0      .0      .0
Select a user for user details
Command ==>
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return

```

Figure 151. Layout of User QDIO Activity Screen (FCX257)

The sample shows a screen with the user IDs sorted in ascending sequence, the default (indicated by the underscore characters above the Userid header). Other possible sort sequences are indicated by dots above the corresponding columns – they can be set using the SORT command with the column heading as sort argument (see “SORT” on page 295), or by placing the cursor on the header for the field which is to be sorted and then pressing the Enter key.

The first entry is reserved for the overall system average for all users, regardless of the current sorting sequence. It is identified byby: >System<-userid.

Some additional lines with user class averages will be shown at the top when user classes have been defined (see “FCONTROL UCLASS” on page 89). The user ID shown will then be the user class name.

Use the scrolling function keys or commands to view data not visible on the initial screen, or use the locate command if you want to search for a specific user.

For more detailed information about a specific user, do one of the following:

1. Place the cursor on the Userid field of the virtual machine you are interested in and press Enter
2. On the command line, enter the subcommand: USER *userid*

### Field Descriptions:

The following field descriptions apply to the UQDIO and UQDIOLOG *userid* displays, with the exception of the Userid field, which is replaced by the Interval End Time field on the UQDIOLOG display.

#### Userid

The user identification.

*Guest Instruct Intercepts/s:*

**TVS**

The number of guest TVS instructions intercepted per second.

**SVS**

The number of guest SVS instructions intercepted per second.

**Alert Interrupts/sec**

The number of alert interruptions handled per second by the guest.

*Adapter Interruptions/s:*

**Delivered to Guest ViaTPI**

The number of adapter interruptions delivered to the guest per second via TPI.

**Delivered to Guest InSIE**

The number of adapter interruptions delivered to the guest per second inside SIE.

**Stacked by CP**

The number of adapter interruptions that were stacked for the guest per second by CP.

**Nr of Users**

The number of users that contributed to the average shown on this line (values are inserted only for lines with averages; this field is blank for user data lines).

**Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages will be shown for all fields. All of a user's sessions during the selected period contribute to the averages shown.

# FCX258, User QDIO Activity Log Screen – UQDIOLOG

**Performance Toolkit Subcommand:** “UQDIOLOG” on page 343.

The User QDIO Activity Log Screen (FCX258) shows a "by time" log of the selected user's QDIO interrupt and instruction intercept activity. Note that the specified user ID must have been selected previously for benchmarking.

FCX258	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor
QDIO Activity Log for User LNX2C				
Interval	<Guest Instruct>	Alert	<- Adapter Interruptions/s ->	
End Time	<-Intercepts/s->	Interrupts	<Delivered to Guest>	Stacked
	TVS	SVS	ViaTPI	InSIE
>>Mean>>	.0	.0	.0	2037
13:13:02	.0	.0	.0	2029
13:14:02	.0	.0	.0	2037
13:15:02	.0	.0	.0	2037
13:16:02	.0	.0	.0	2043
13:17:02	.0	.0	.0	2035
13:18:02	.0	.0	.0	2044

Command ==>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 152. Layout of User QDIO Activity Log Screen (FCX258)

The information shown is based on USER domain SAMPLE data.

The first entry is reserved for overall report period averages. It is designated by the string >>Mean>> in the 'Interval End Time' column.

**Field Descriptions:**

Except for the 'Interval End Time' field which replaces the 'Userid' field, all data columns are identical to those for the 'User QDIO Activity' (UQDIO) display. Please see “FCX257, User QDIO Activity Screen – UQDIO” on page 733 for their description.

**Interval End Time**

The end time of the monitor interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX259, Demand Scan Log Screen – DEMNDLOG

**Performance Toolkit Subcommand:** “DEMNDLOG” on page 167.

The Demand Scan Log Screen (FCX259):

FCX259	CPU nnnn SER nnnnn Interval HH:MM:SS - HH:MM:SS										Perf. Monitor									
	<----- Demand Scan Pass 1 ----->										<----- Demand Scan Pass 2 ----->									
	<-- Ended After -->					<--- Page Frames --->					<--Ended After-->					<--- Page F				
Interval	Lng	Drm	NSS	Eli	Dsp	Long	Dor-	NSS	Eli	Dsp	Lng	Drm	Eli	Dsp	Long	Dor-	F			
End Time	Drm	ant	Shr	Lst	Lst	Dorm	mant	Shr	Lst	Lst	Drm	ant	Lst	Lst	Dorm	mant	S			
>>Mean>>	0	0	0	0	23	0	1747	0	0	51k	0	3	0	0	0	5608				
07:54:48	0	1	0	0	57	0	1811	0	0	59k	0	3	0	0	0	7561				
07:56:18	0	1	0	0	40	0	2851	0	0	68k	0	3	0	0	0	4119				
07:57:48	0	0	0	0	19	0	658	0	0	59k	0	2	0	0	0	3738				
07:59:18	0	0	0	0	12	1	2677	0	0	44k	0	3	0	0	0	6237				
08:00:48	0	1	0	0	20	1	2133	0	0	50k	0	4	0	0	0	7249				
08:02:18	0	0	0	0	9	3	2277	0	0	49k	0	4	0	0	0	6113				
08:03:48	0	0	0	0	12	0	683	0	0	38k	0	3	0	0	0	5427				
08:04:48	0	0	0	0	10	0	461	0	0	36k	0	4	0	0	0	3829				
See AVAILLOG, PAGELOG and STORLOG for more storage related information																				
Command ==>																				
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return																				

Figure 153. Layout of Demand Scan Log Screen (FCX259)

This screen applies only to monitor data when running Performance Toolkit against z/VM 6.2 and earlier data.

The information shown is based on CP monitor STORAGE domain SAMPLE data.

The display shows detailed demand scan information by time. The leftmost column contains the end time (in hh:mm:ss format) of the monitor interval for which the values were calculated.

The first entry is reserved for the overall report period averages. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

*Demand Scan Pass 1:* Provides demand scan details for the initial pass 1.

#### Ended After

The number of times pass 1 of the demand scan ended after scanning a certain list of users or system pages:

##### Lng Drm

Users on the long term dormant list

##### Drm ant

Users on the dormant list

##### NSS Shr

System VMDBK, NSSs, DCSSs, shared address space

##### Eli Lst

Users on the eligible list

##### Dsp Lst

Users on the dispatch list

#### Page Frames

The number of page frames obtained during pass 1 of the demand scan:

##### Long Dorm

From users on the long term dormant list



**Dormant**

From users on the dormant list

**NSS Shr**

From System VMDBK, NSSs, DCSSs, shared spaces

**Eli Lst**

From users on the eligible list

**Dsp Lst**

From users on the dispatch list

*Demand Scan Pass 2:* Provides demand scan details for the second pass. A second pass may become necessary if sufficient page frames could not be obtained during the first pass.

**Ended After**

The number of times pass 2 of the demand scan ended after scanning a certain list of users:

**Lng Drm**

Users on the long term dormant list

**Drm ant**

Users on the dormant list

**Eli Lst**

Users on the eligible list

**Dsp Lst**

Users on the dispatch list

**Page Frames**

The number of page frames obtained during pass 2 of the demand scan:

**Long Dorm**

From users on the long term dormant list

**Dormant**

From users on the dormant list

**NSS Shr**

From System VMDBK, NSSs, DCSSs, shared spaces

**Eli Lst**

From users on the eligible list

**Dsp Lst**

From users on the dispatch list

*Emergency Scan:* Provides details for the emergency scan. An emergency scan may become necessary if sufficient page frames could not be obtained during pass 1 and pass 2.

**Ended After**

The number of times the emergency scan ended after scanning a certain list of users or system pages:

**Drm ant**

Users on the dormant list

**NSS Shr**

System VMDBK, NSSs, DCSSs, shared address space

**Eli Lst**

Users on the eligible list

**Dsp Lst**

Users on the dispatch list

**Page Frames**

The number of page frames obtained during the emergency scan:

**Dormant**

From users on the dormant list

**NSS Shr**

From System VMDBK, NSSs, DCSSs, shared spaces

**Eli Lst**

From users on the eligible list

**Dsp Lst**

From users on the dispatch list

**Scan Failed**

The number of times the replenishment function failed obtaining a sufficient number of page frames even after the emergency scan.

**Effect of AVERAGE and INTERIM subcommands:**

None.

# FCX260, Storage Management Logs Menu Screen – STORMENU

**Performance Toolkit Subcommand:** “STORMENU” on page 304. The Storage Management Logs Menu Screen (FCX260) shows a list of all available 'by time' logs that are related to storage management.

```
FCX260      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor

  General Storage/Paging
S Log Name  Description
- AVLB2GLG  Below 2G Available list management log
- AVLA2GLG  Above 2G Available list management log
- STEALLOG  Frame steal details log
- AGELLOG   Age List details log
- AVAILLOG  Available list management log
- DEMNDLOG  Demand scan details log
- PAGELOG   Paging activity log
- STORLOG   Storage utilization log
- PINLOG    Pinned storage statistics log
- SUBPLOG   Subpool storage usage log

  System Execution Space
S Log Name  Description
- SXSAVAIL  SXS page queues management log
- SXSDEFER  SXS deferred tasks management log
- SXSPAGE   SXS page management log
- SXSUTIL   SXS storage utilization log

Select performance screen with cursor and press ENTER
Command ==>
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 154. Layout of Storage Management Logs Menu Screen (FCX260)

All of the selectable logs are based on CP monitor SYSTEM and STORAGE domain SAMPLE data.

**Field Descriptions:**

- S**  
Cursor selection field. The field will be highlighted and made cursor-selectable (set as 'unprotected') only if the corresponding storage mgmt. log is available.
- Log Name**  
The log file name, equal to the command with which the log data can be selected for display. The field will be highlighted if the corresponding storage management log is available.
- Description**  
Description of the storage management log.
- Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX261, SXS Available Page Queues Management Log Screen – SXS AVAIL

**Performance Toolkit Subcommand:** “SXS AVAIL” on page 306. The SXS Available Page Queues Management Log Screen:

FCX261	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor														
Interval	<-- Backed <2GB Page Queue -->					<-- Backed >2GB Page Queue -->					<----- Unbacked Page Queue ----->					<---- Replenishment ---->			
End Time	Avail	<-Pages/s-->	<Preferred>			Avail	<-Pages/s-->	<Preferred>			Avail	<-Pages/s-->	<Preferred>			Thres	Att/s	Stolen	MinPgs
Pages	Taken	Return	Used	Empty		Pages	Taken	Return	Used	Empty	Pages	Taken	Return	Used	Empty				
>>Mean>>	21	3.272	3.271	3.272	.013	22	6.740	6.747	5.422	.785	492673	95.64	94.79	94.45	.000	128	.000	128	...
20:14:42	14	3.383	3.583	3.383	.017	21	4.750	4.900	2.150	.283	493813	491.0	486.1	490.0	.000	128	.000	128	...
20:15:42	25	.150	.333	.150	.000	21	5.200	5.200	2.817	.083	493760	19.25	18.37	18.75	.000	128	.000	128	...
20:16:42	25	.200	.200	.200	.000	32	7.233	7.417	5.183	.000	493742	30.38	30.08	29.93	.000	128	.000	128	...
20:17:42	9	1.333	1.067	1.333	.633	15	25.15	24.87	21.75	.000	493505	33.72	29.77	32.98	.000	128	.000	128	...
20:18:42	32	31.58	31.97	31.58	.117	20	10.03	10.12	5.083	.017	493378	69.65	67.53	69.38	.000	128	.000	128	...
20:19:42	28	93.23	93.17	93.23	.000	32	11.47	11.67	8.617	8.283	492991	170.2	163.8	161.3	.000	128	.000	128	...
20:20:42	19	36.47	36.32	36.47	.000	32	11.82	11.82	10.20	.567	492446	131.6	122.6	130.5	.000	128	.000	128	...
20:21:42	28	.850	1.000	.850	.000	22	15.33	15.17	14.02	2.333	492482	979.7	980.3	975.8	.000	128	.000	128	...
20:22:42	26	2.417	2.383	2.417	.000	27	25.92	26.00	23.67	7.650	491474	512.6	495.8	503.4	.000	128	.000	128	...
20:23:42	14	.700	.500	.700	.000	26	16.92	16.90	14.38	25.07	489936	199.6	174.0	172.5	.000	128	.000	128	...
20:24:42	13	.517	.500	.517	.000	29	12.08	12.13	9.750	.233	490129	432.9	436.1	431.4	.000	128	.000	128	...
20:25:42	18	.217	.300	.217	.000	31	10.45	10.48	9.283	.000	489496	172.5	162.0	171.2	.000	128	.000	128	...
20:26:42	15	.500	.450	.500	.000	31	20.62	20.62	18.93	.850	490908	441.9	465.4	440.3	.000	128	.000	128	...
20:27:42	17	.167	.200	.167	.000	32	11.12	11.13	9.700	.000	490983	45.32	46.58	44.23	.000	128	.000	128	...
20:28:42	16	.200	.183	.200	.000	29	30.22	30.17	28.67	.100	491181	117.0	120.3	115.9	.000	128	.000	128	...
20:29:42	13	.233	.183	.233	.000	30	61.70	61.72	60.38	.000	491230	92.05	92.87	90.90	.000	128	.000	128	...
20:30:42	15	.200	.233	.200	.000	14	6.117	5.850	4.117	.133	491405	75.12	78.03	74.58	.000	128	.000	128	...
20:31:42	13	.217	.183	.217	.000	29	7.817	8.067	5.917	.000	491390	46.42	46.15	45.82	.000	128	.000	128	...
20:32:42	12	.550	.533	.550	.000	17	8.233	8.033	6.267	.650	491502	38.55	40.42	37.45	.000	128	.000	128	...
20:33:42	13	.183	.200	.183	.000	32	5.733	5.983	3.467	.117	491517	35.42	35.67	34.98	.000	128	.000	128	...
20:34:42	14	.200	.217	.200	.000	27	14.38	14.30	12.52	.117	491485	417.4	416.9	416.7	.000	128	.000	128	...
20:35:42	17	.133	.183	.133	.000	26	6.067	6.050	3.850	.000	491346	116.2	113.9	115.9	.000	128	.000	128	...
20:36:42	15	.183	.150	.183	.000	10	5.867	5.600	3.850	.117	490076	400.0	378.8	399.4	.000	128	.000	128	...
20:37:42	14	.333	.317	.333	.000	32	5.750	6.117	3.717	.000	491367	31.58	53.10	31.08	.000	128	.000	128	...
20:38:42	14	3.250	3.250	3.250	.000	25	6.050	5.933	2.917	.133	491500	34.63	36.85	33.58	.000	128	.000	128	...
20:39:42	14	.183	.183	.183	.000	31	11.08	11.18	6.300	.000	491403	31.63	30.02	31.43	.000	128	.000	128	...
20:40:42	14	.367	.367	.367	.000	24	7.233	7.117	2.933	.250	491396	44.70	44.58	43.75	.000	128	.000	128	...
20:41:42	15	1.067	1.083	1.067	.000	31	6.583	6.700	2.867	.000	491518	27.85	29.88	26.48	.000	128	.000	128	...
20:42:42	15	3.033	3.033	3.033	.000	32	11.43	11.45	7.883	1.100	491603	39.15	40.57	36.60	.000	128	.000	128	...
20:43:42	15	2.283	2.283	2.283	.000	32	9.233	9.233	4.833	.000	490106	60.92	35.97	60.32	.000	128	.000	128	...
Command ==>																			
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return																			

Figure 155. Layout of SXS Available Page Queues Management Log Screen (FCX261)

The information shown is based on CP monitor SYSTEM and STORAGE domain SAMPLE data.

The display shows SXS available page queues management information by time. The leftmost column contains the end time (in hh:mm:ss format) of the monitor interval for which the values were calculated.

The first entry is reserved for the overall report period averages. It is designated by the string >>Mean>> in the 'Interval End Time' column.

## Field Descriptions:

### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

**Backed <2GB Page Queue:** Information on the backed below 2GB System Execution Space available page queue.

### Avail Pages

The number of pages available on the backed below 2GB SXS page queue.

### Pages/s Taken

The number of pages taken per second from the backed below 2GB SXS page queue.

### Pages/s Return

The number of pages returned per second to the backed below 2GB SXS page queue.

### Preferred Used

The number of times per second the backed below 2GB SXS page queue was the preferred queue, and was used to satisfy the request.

**Preferred Empty**

The number of times per second the backed below 2GB SXS page queue was the preferred queue for the request, but was empty.

*Backed >2GB Page Queue:* Information on the backed above 2GB System Execution Space available page queue.

**Avail Pages**

The number of pages available on the backed above 2GB SXS page queue.

**Pages/s Taken**

The number of pages taken per second from the backed above 2GB SXS page queue.

**Pages/s Return**

The number of pages returned per second to the backed above 2GB SXS page queue.

**Preferred Used**

The number of times per second the backed above 2GB SXS page queue was the preferred queue, and was used to satisfy the request.

**Preferred Empty**

The number of times per second the backed above 2GB SXS page queue was the preferred queue for the request, but was empty.

*Unbacked Page Queue:* Information on the unbacked System Execution Space page queue.

**Avail Pages**

The number of pages available on the unbacked SXS page queue.

**Pages/s Taken**

The number of pages taken per second from the unbacked SXS page queue.

**Pages/s Return**

The number of pages returned per second to the unbacked SXS page queue.

**Preferred Used**

The number of times per second the unbacked SXS page queue was the preferred queue, and was used to satisfy the request.

**Preferred Empty**

The number of times per second the unbacked SXS page queue was the preferred queue for the request, but was empty.

*Replenishment:* Provides information on replenishment management of the SXS unbacked page queue.

**Thres**

The threshold that triggers replenishment of the available unbacked SXS page queue.

**Att/s**

The number of attempts made per second to replenish the available unbacked SXS page queue.

**Stolen**

The number of aliases that are stolen during a single replenishment of the available unbacked SXS page queue.

**MinPgs**

The smallest number of pages on the available, unbacked SXS page queue at the completion of replenishment.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX262, SXS Page Management Log Screen – SXSPAGE

**Performance Toolkit Subcommand:** [“SXSPAGE” on page 308.](#)

### The SXS Page Management Log Screen (FCX262):

CPU	nnnn	SER	nnnnn	Interval	HH:MM:SS	-	HH:MM:SS	Perf.	Monitor			
Interval	<---- Page Requests/s ---->			<----- Page Returns/s ----->			<---- Aliases ----->		<PTE Invalidation>			
End Time	Free	Non-FS	<Contiguous>	Free	Non-FS	<Contiguous>	<Queues of Pgs>	<-- Requests/sec -->	<-- Rates/sec -->	Steal	CSP	IPTE
>>Mean>>	Stor	Single	Requests	Stor	Single	Returns	Pages	Create	Remove	Queue		
20:17:42	1.086	9.342	.375 1.840	1.002	9.215	.374 1.837	.164 .428	92.50	92.07	.000	.032	.000 93.18
20:18:42	1.317	23.72	.700 3.500	1.150	22.48	.700 3.500	.083 .133	32.98	29.72	.000	.000	.000 29.77
20:19:42	.650	36.80	1.017 5.083	1.100	37.27	1.017 5.083	.200 .900	69.38	67.27	.000	.000	.000 67.53
20:20:42	9.350	110.1	.700 3.500	7.400	108.1	.700 3.500	.133 7.183	161.3	157.0	.000	.000	.000 163.8
20:21:42	2.817	47.23	.500 2.500	2.817	46.87	.500 2.500	.117 1.700	130.5	121.3	.000	.317	.000 122.9
20:22:42	1.983	17.20	.633 2.900	1.067	16.03	.633 2.900	.183 .450	975.8	977.5	.000	.067	.000 980.3
20:23:42	5.750	33.73	.867 4.067	.650	28.00	.867 4.067	.083 .250	503.4	492.2	.000	.267	.000 496.1
20:24:42	27.53	40.15	.967 4.833	5.983	17.27	.967 4.833	.050 4.733	172.5	169.3	.000	.300	.000 174.3
20:25:42	.983	10.50	.717 3.583	3.750	13.45	.717 3.583	.200 1.267	431.4	431.7	.000	.050	.000 436.1
20:26:42	.150	9.500	.500 2.500	5.133	14.48	.500 2.500	.233 .333	171.2	155.8	.000	.000	.000 162.0
20:27:42	.133	20.28	.500 2.500	3.250	23.30	.500 2.500	.050 .050	440.3	460.7	.000	.067	.000 465.5
20:28:42	.067	9.867	.500 2.500	3.267	13.18	.500 2.500	.133 .167	44.25	42.25	.000	.000	.000 46.58
20:29:42	.083	28.97	.500 2.500	3.067	31.88	.500 2.500	.067 .083	115.9	116.2	.000	.017	.000 120.3
20:30:42	.317	60.62	.517 2.533	2.183	62.35	.517 2.533	.067 .083	90.90	89.88	.000	.067	.000 92.93
20:31:42	.183	4.450	.500 2.500	1.717	5.917	.500 2.500	.083 .100	74.58	75.70	.000	.100	.000 78.13
20:32:42	.183	6.133	.500 2.500	.983	6.900	.500 2.500	.083 .083	45.80	45.00	.000	.000	.000 46.15
20:33:42	.250	7.467	.500 2.500	.633	7.883	.500 2.500	.050 .067	37.45	38.60	.000	.083	.000 40.50
20:34:42	.317	3.767	.517 2.583	.650	4.117	.517 2.583	.083 .167	34.98	35.15	.000	.000	.000 35.67
20:35:42	.500	12.83	.500 2.500	.367	12.73	.500 2.500	.067 .067	416.7	416.2	.000	.050	.000 416.9
20:36:42	.217	3.983	.500 2.500	.533	4.267	.500 2.500	.117 .183	115.9	113.3	.000	.000	.000 113.9
20:37:42	.217	4.150	.500 2.500	.283	4.333	.500 2.500	.033 .033	399.4	377.7	.000	.033	.000 378.8
20:38:42	.100	4.050	.517 2.533	.383	4.350	.517 2.533	.033 .133	31.08	52.65	.000	.000	.000 53.10
20:39:42	.250	6.300	.833 4.167	.417	6.467	.833 4.167	.133 .150	33.58	35.40	.000	.117	.000 36.97
20:40:42	.133	6.483	1.000 5.000	.117	6.433	1.000 5.000	.050 .050	31.43	29.95	.000	.017	.000 30.03
20:41:42	.367	3.550	1.000 5.000	.350	3.517	1.000 5.000	.050 .050	43.75	43.55	.000	.000	.000 44.58
20:42:42	.083	3.933	1.017 5.083	.067	4.050	1.017 5.083	.033 .033	26.48	28.53	.000	.000	.000 29.88
20:43:42	.267	12.02	1.000 5.000	.200	12.13	1.000 5.000	.067 .067	36.60	37.92	.000	.000	.000 40.57
20:44:42	.133	7.117	1.000 5.000	.067	6.883	1.000 5.000	.050 .050	60.32	35.60	.000	.000	.000 35.97
20:45:42	.133	4.700	1.000 5.000	.050	4.500	1.000 5.000	.050 .050	31.40	31.42	.000	.017	.000 32.22
20:46:42	.100	4.700	.933 4.667	.167	4.817	.933 4.667	.083 .083	32.78	51.80	.000	.000	.000 52.43
Command ==>>												
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return							

Figure 156. Layout of SXS Page Management Log Screen (FCX262)

The information shown is based on CP monitor SYSTEM and STORAGE domain SAMPLE data.

The display shows SXS page management information by time. The leftmost column contains the end time (in *hh:mm:ss* format) of the last monitor interval for which the values were calculated.

The first entry is reserved for the overall report period averages. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Descriptions:

## Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

## Page Requests/s:

## Free Stor

The number of free storage pages requested per second.

### Non-FS Single

The number of non-free storage single (non-contiguous) SXS pages requested per second.

## Contiguous Requests

The number of SXS page requests per second for two or more contiguous pages.

## Contiguous Pages

The number of contiguous SXS pages obtained per second for contiguous page requests.

Page Returns/s:

## Free Stor

The number of free storage pages returned per second.

**Non-FS Single**

The number of non-free storage single (non-contiguous) SXS pages returned per second.

**Contiguous Returns**

The number of SXS page returns per second for two or more contiguous pages.

**Contiguous Pages**

The number of contiguous SXS pages returned per second for contiguous page returns.

**Queues of Pgs Returns**

The number of returns per second of queues of SXS pages.

**Queues of Pgs Pages**

The number of SXS pages returned per second via queues.

*Aliases Requests/sec:***Create**

The number of 'create alias' requests per second.

**Remove**

The number of 'remove alias' requests per second.

**Queue**

The number of 'queue SXSTE' requests per second, to requeue aliases dequeued by the SXS page manager when locked during a search for aliases to steal.

*PTE Invalidation Rates/sec:***Steal**

The number of aliases stolen per second by the SXS page manager.

**CSP**

The number of CSP instructions used per second instead of IPTE when invalidating multiple page table entries.

**IPTE**

The number of IPTE instructions executed per second to invalidate an SXS page table entry.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX263, SXS Deferred Tasks Management Log Screen – SXSDEFER

Performance Toolkit Subcommand: “SXSDEFER” on page 307.

The SXS Deferred Tasks Management Log Screen (FCX263):

FCX263	CPU	nnnn	SER	nnnnn	Interval HH:MM:SS - HH:MM:SS				Perf. Monitor									
<-Req. Deferred/s->		<--Non-Deferrable-->				<--Deferred-->		<---- Deferred Page Requests ----->				<-Redrives/s->						
Interval	On a	On a	Frame	<Requests	Failed/s->	<Create	Alias	Total	<With	Backing>	<For	Frame->	Att-	Ind.Task				
End Time	Page	Any	<2GB	Alias	NoPage	NoFrame	Average	Rate/s	Any		Any	<2GB	empts	Redrives				
>>Mean>>	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:18:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:19:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:20:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:21:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:22:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:23:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:24:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:25:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:26:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:27:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:28:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:29:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:30:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:31:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:32:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:33:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:34:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:35:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:36:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:37:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:38:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:39:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:40:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:41:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
20:42:42	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000				
Command ==>																		
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return													

Figure 157. Layout of SXS Deferred Tasks Management Log Screen (FCX263)

The information shown is based on CP monitor SYSTEM and STORAGE domain SAMPLE data.

The display shows SXS deferred tasks management information by time. The leftmost column contains the end time (in *hh:mm:ss* format) of the last monitor interval for which the values were calculated.

The first entry is reserved for the overall report period averages. It is designated by the string *>>Mean>>* in the 'Interval End Time' column.\

## Field Descriptions:

### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the *>>Mean>>* entry, you can find the relevant interval in the header line.

*Req. Deferred/s:*

### On a Page

The number of times per second an SXS page request deferred on a page.

### On a Frame, Any

The number of times per second an SXS page request deferred on any frame.

### On a Frame <2GB

The number of times per second an SXS page request deferred on a frame below 2GB.

*Non-Deferrable Requests Failed/s:*

### Alias

The number of times per second a non-deferrable SXS alias request could not be fulfilled.

### NoPage

The number of times per second a non-deferrable SXS page request failed due to the lack of a page.

### NoFrame

The number of times per second a non-deferrable SXS page request failed due to the lack of a frame.



*Deferred Create Alias:*

**Average**

The average number of SXS 'create alias' requests deferred at the end of each sample interval during the corresponding period.

**Rate/s**

The number of times per second a 'create alias' request was deferred.

*Deferred Page Requests:*

**Total**

The average total number of SXS page requests that were deferred at the end of each sample interval during the corresponding period.

**With Backing Any**

The average number of SXS page requests with backing=any that were deferred at the end of each sample interval during the corresponding period.

**With Backing <2GB**

The average number of SXS page requests with backing below 2GB that were deferred at the end of each sample interval during the corresponding period.

**For Frame, Any**

The average number of SXS page requests that were deferred for any frame at the end of each sample interval during the corresponding period.

**For Frame <2GB**

The average number of SXS page requests that were deferred for a frame below 2GB at the end of each sample interval during the corresponding period.

*Redrives/s:*

**Attempts**

The number of attempts made per second to redrive tasks waiting for an available page.

**Ind.Task Redrives**

The number of individual redrives per second for tasks waiting for an available page.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX264, SXS Storage Utilization Log Screen – SXSUTIL

Performance Toolkit Subcommand: “SXSUTIL” on page 309.

The SXS Storage Utilization Log Screen (FCX264):

FCX264	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS					Perf. Monitor						
<----- System Execution Space Utilization (Pages) ----->														
<----- Used ----->														
<----- Aliases ----->														
Interval	Total	Avail-	ID			CP		Aliases		No-	<Pages	Backed>	Poten-	Conti-
End Time	Pages	able	Total	Mapped	Free	Other	Total	---	Locked---	Owned	<2GB	>2GB	Steal	iguous
		Queues						Total	CPLock				Queue	Pages
>>Mean>>	524287	492561	31726	4057	6297	14489	6884	87	0	0	16842	14927	6893	1878
20:05:42	524287	494272	30015	4057	6003	14409	5546	74	0	0	16669	13382	5555	1871
20:06:42	524287	494263	30024	4057	6001	14410	5556	72	0	0	16669	13391	5565	1871
20:07:42	524287	493856	30431	4057	6020	14417	5937	76	0	0	16712	13752	5946	1871
20:08:42	524287	493874	30413	4057	5999	14414	5943	85	0	0	16710	13756	5952	1871
20:09:42	524287	493999	30288	4057	5995	14417	5819	76	0	0	16724	13614	5828	1871
20:10:42	524287	493690	30597	4057	6019	14441	6080	78	0	0	16768	13853	6089	1883
20:11:42	524287	494146	30141	4057	6000	14436	5648	71	0	0	16672	13502	5657	1883
20:12:42	524287	494162	30125	4057	5985	14436	5647	69	0	0	16671	13486	5656	1883
20:13:42	524287	494124	30163	4057	6004	14438	5664	70	0	0	16678	13499	5673	1883
20:14:42	524287	493848	30439	4057	6011	14433	5938	75	0	0	16729	13745	5947	1883
20:15:42	524287	493806	30481	4057	6007	14434	5983	71	0	0	16743	13784	5992	1883
20:16:42	524287	493799	30488	4057	6013	14434	5984	71	0	0	16741	13804	5993	1883
20:17:42	524287	493529	30758	4057	6083	14438	6180	79	0	0	16791	13991	6189	1883
20:18:42	524287	493430	30857	4057	6056	14437	6307	73	0	0	16806	14103	6316	1883
20:19:42	524287	493051	31236	4057	6173	14443	6563	74	0	0	16994	14302	6572	1883
20:20:42	524287	492497	31790	4057	6173	14465	7095	69	0	0	17069	14772	7104	1883
20:21:42	524287	492532	31755	4057	6228	14480	6990	81	0	0	17070	14735	6999	1883
20:22:42	524287	491527	32760	4057	6534	14518	7651	85	0	0	17121	15692	7660	1883
20:23:42	524287	489976	34311	4057	7827	14598	7829	116	0	0	16994	17357	7838	1883
20:24:42	524287	490171	34116	4057	7661	14587	7811	127	0	0	17010	17148	7820	1883
20:25:42	524287	489545	34742	4057	7362	14587	8736	137	0	0	17124	17667	8745	1883
20:26:42	524287	490954	33333	4057	7175	14592	7509	144	0	0	16956	16423	7518	1883
20:27:42	524287	491032	33255	4057	6983	14586	7629	143	0	0	16951	16353	7638	1883
20:28:42	524287	491226	33061	4057	6804	14590	7610	82	0	0	16963	16143	7619	1883
20:29:42	524287	491273	33014	4057	6692	14598	7667	85	0	0	16969	16088	7676	1883
20:30:42	524287	491434	32853	4057	6600	14602	7594	86	0	0	16947	15935	7603	1883
20:31:42	524287	491432	32855	4057	6552	14604	7642	150	0	0	16950	15947	7651	1883
20:32:42	524287	491531	32756	4057	6529	14602	7568	74	0	0	16951	15834	7577	1883
20:33:42	524287	491562	32725	4057	6509	14601	7558	149	0	0	16947	15823	7567	1883
20:34:42	524287	491526	32761	4057	6517	14599	7588	77	0	0	16958	15844	7597	1883
20:35:42	524287	491389	32898	4057	6498	14601	7742	81	0	0	16959	15982	7751	1883
20:36:42	524287	490101	34186	4057	6494	14594	9041	85	0	0	17087	17124	9050	1883
20:37:42	524287	491413	32874	4057	6477	14593	7747	84	0	0	16959	15961	7756	1883
20:38:42	524287	491539	32748	4057	6467	14593	7631	78	0	0	16960	15827	7640	1883
20:39:42	524287	491448	32839	4057	6468	14595	7719	147	0	1	16960	15924	7727	1883
20:40:42	524287	491434	32853	4057	6469	14596	7731	152	0	2	16970	15921	7738	1883
20:41:42	524287	491564	32723	4057	6470	14588	7608	80	0	0	16968	15801	7617	1883
20:42:42	524287	491650	32637	4057	6474	14577	7529	75	0	0	16945	15739	7538	1883
20:43:42	524287	490153	34134	4057	6478	14587	9012	92	0	0	17107	17074	9021	1883
20:44:42	524287	490143	34144	4057	6483	14594	9010	78	0	0	17126	17048	9019	1883
20:45:42	524287	491291	32996	4057	6479	14591	7869	146	0	0	16991	16049	7878	1883
20:46:42	524287	491279	33008	4057	6485	14580	7886	144	0	0	16993	16060	7895	1883
20:47:42	524287	491300	32987	4057	6483	14581	7866	75	0	0	16992	16037	7875	1883
20:48:42	524287	490962	33325	4057	6493	14582	8193	148	0	0	17053	16313	8202	1883
Command ==>>														
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return														

Figure 158. Layout of SXS Storage Utilization Log Screen (FCX264)

The information shown is based on CP monitor SYSTEM and STORAGE domain SAMPLE data.

The display shows SXS storage utilization information by time. The leftmost column contains the end time (in hh:mm:ss format) of the last monitor interval for which the values were calculated.

The first entry is reserved for the overall report period averages. It is designated by the string >>Mean>> in the 'Interval End Time' column.

## Field Descriptions:

### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

### *System Execution Space Utilization (Pages):*

#### **Total Pages**

The total system execution space size in pages.

#### **Available Queues**

The total number of SXS pages on the different available page queues (backed<2GB, backed>2GB and unbacked). See [“FCX261, SXS Available Page Queues Management Log Screen – SXS AVAIL”](#) on page 740 for details.

#### *Used:*

#### **Total**

The total number of SXS pages that are in use as CP pages or aliases, including no-owned aliases and SXS pages assigned to tasks deferred on frames in the real storage manager.

#### **CP ID Mapped**

The number of ID-mapped CP pages (host logical = host real), primarily CP nucleus and prefix pages.

#### **CP Free**

The total number of CP free storage pages, with backing frames both below and above 2GB.

#### **CP Other**

Count of other CP type pages (excluding ID-mapped and free storage pages).

#### **Aliases Total**

The number of SXS alias pages associated with 'pageable FRMTes'. They may be stealable if not last translated or locked.

#### **Aliases Locked Total**

The number of locked SXS alias pages.

#### **Aliases Locked CPLock**

The number of locked SXS alias pages that were locked by CP LOCK command.

#### **Aliases No-Owned**

The number of no-owned SXS alias pages.

#### **Pages Backed <2GB**

The number of SXS pages backed below 2GB.

#### **Pages Backed >2GB**

The number of SXS pages backed above 2GB.

#### **Potential Steal Queue**

The number of pages on the SXS potentially stealable queue.

#### **Contiguous Pages**

The number of SXS pages in use as part of a contiguous page request for more than one page. This count does not include pages that were pre-allocated as part of initialization.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX265, Spin Lock Log Screen – LOCKLOG

**Performance Toolkit Subcommand:** “LOCKLOG” on page 219.

The Spin Lock Log Screen (FCX265) shows spin lock information by time.

FCX265      CPU 2817   SER 36D85   Interval 00:00:30 - 07:33:30      Perf. Monitor													
<----- Spin Lock Collision Activity ----->													
<----- Combined ----->				<----- Exclusive ----->				<----- Shared ----->					
Interval	Colls	AvgSpin	Spin	CADs	Colls	AvgSpin	Spin	CADs	Colls	AvgSpin	Spin	CADs	
End Time LockName	/sec	usec	%Busy	/sec	/sec	usec	%Busy	/sec	/sec	usec	%Busy	/sec	
>>Mean>> RCCSFQL	.0	.000	.000	.0	.0	.000	.000	.0	.0	.000	.000	.0	
>>Mean>> RSANOQLK	.0	.000	.000	.0	.0	.000	.000	.0	.0	.000	.000	.0	
>>Mean>> NSUNLSLK	.0	.000	.000	.0	.0	.000	.000	.0	.0	.000	.000	.0	
>>Mean>> HCPPGDML	.0	.000	.000	.0	.0	.000	.000	.0	.0	.000	.000	.0	
>>Mean>> NSUIMGLK	.0	.000	.000	.0	.0	.000	.000	.0	.0	.000	.000	.0	
>>Mean>> FSDVMLK	.0	.000	.000	.0	.0	.000	.000	.0	.0	.000	.000	.0	
>>Mean>> DCTLLOK	.0	.000	.000	.0	.0	.000	.000	.0	.0	.000	.000	.0	
>>Mean>> SYSDATLK	.0	.000	.000	.0	.0	.000	.000	.0	.0	.000	.000	.0	
>>Mean>> RSACALLK	.1	.249	.000	.1	.1	.249	.000	.0	.0	.000	.000	.0	
>>Mean>> RSAAVLLK	.0	.000	.000	.0	.0	.000	.000	.0	.0	.000	.000	.0	
>>Mean>> HCPPGDAL	.0	.000	.000	.0	.0	.000	.000	.0	.0	.000	.000	.0	
>>Mean>> HCPPGDTL	.0	.000	.000	.0	.0	.000	.000	.0	.0	.000	.000	.0	
>>Mean>> >>Total>	.1	.249	.000	.1	.1	.249	.000	.0	.0	.000	.000	.0	
Command ==> _													
F1=Help   F4=Top   F5=Bot   F7=Bkwd   F8=Fwd   F12=Return													

Figure 159. Layout of Spin Lock Log Screen (FCX265)

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### Lock Name

The name of the CP spin lock.

#### Combined Colls/sec

The total exclusive and shared spin lock collisions per second.

#### Combined AvgSpin usec

The total average microseconds spent spinning on both exclusive and shared spin lock collisions.

#### Combined Spin %Busy

The percentage of total elapsed time spent in both exclusive and shared spin lock collisions.

#### Combined CADs/sec

Total number of CAD (Compare and Delay facility) instructions executed when attempting to obtain the lock in shared and exclusive mode, per second.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

When running Performance Toolkit against data from z/VM 6.3 without APAR VM65586 installed, this column contains .0 values.

#### Exclusive Colls/sec

The total exclusive spin lock collisions per second.

#### Exclusive AvgSpin usec

The average total microseconds spent spinning on exclusive spin lock collisions.

#### Exclusive Spin %Busy

The percentage of total elapsed time spent in exclusive spin lock collisions.

#### Exclusive CADs/sec

Total number of CAD (Compare and Delay facility) instructions executed when attempting to obtain the lock in exclusive mode, per second.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

When running Performance Toolkit against z/VM 6.3 data without APAR VM65586 installed, this column contains .0 values.

**Shared Colls/sec**

The total shared spin lock collisions per second.

**Shared AvgSpin usec**

The average total microseconds spent spinning on shared spin lock collisions.

**Shared Spin %Busy**

The percentage of total elapsed time spent in shared spin lock collisions.

**Shared CADs/sec**

Total number of CAD (Compare and Delay facility) instructions executed when attempting to obtain the lock in shared mode, per second.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

When running Performance Toolkit against z/VM 6.3 data without APAR VM65586 installed, this column contains .0 values.

**Effect of AVERAGE and INTERIM subcommands:**

None.

## FCX266, General Virtual Switch Description Screen – GVSWITCH

**Performance Toolkit Subcommand:** “GVSWITCH” on page 193.

In the General Virtual Switch Description Screen (FCX266), all of the virtual switches are shown, by name, without any associated performance data. If you select any row in this screen and press Enter, you will go to the Virtual Switch Activity Screen for that device. (See “FCX240, Virtual Switch Activity Screen – VSWITCH” on page 707.)

FCX266	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor												
----	----	-----	-----	-----	-----												
Addr	Name	Group	Vswitch	Macaddr	< SNMP Network Management >			OSA	<--Config-->			<----- Bridge			Port	----->	
					IP address	Userid	Layer	Port	Isol	Acc	Typ	Parm	Role	State	Flag	MFSs	LPARname
60D5	PEGIVL		02-12-34-00-00-02		.....		2	01	OFF	U	I	...	.....	.....	.....	.....	.....
63D0	GLOB1	PGRP1	02-12-34-00-00-0A		.....		2	01	OFF	U	Q	LSG	.....	.....	.....	.....	.....
63D3	GLOB2	PGRP1	02-12-34-00-00-0E		.....		2	01	OFF	U	Q	.SG	.....	.....	.....	.....	.....
64C0	GLOB1	PGRP1	02-12-34-00-00-0A		.....		2	00	OFF	U	Q	.SG	.....	.....	.....	.....	.....
64C3	GLOB2	PGRP1	02-12-34-00-00-0E		.....		2	00	OFF	U	Q	.SG	.....	.....	.....	.....	.....

See EVSWITCH / GVSWITCH / GLONACT for other Vswitch reports  
 Command ==>  
 F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 160. Layout of General Virtual Switch Description Screen (FCX266)

The information shown is based on CP monitor I/O domain data. For correct data display, the I/O domain (EVENT data) must be enabled.

The sample shows a screen with virtual switches sorted in ascending device number sequence (indicated by the underscore characters above the 'Addr' header). Other possible sort sequences are indicated by dots above the corresponding columns; they can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

### Field Descriptions:

#### Addr

The real device number.

#### Name

The name of the virtual switch. A virtual switch is given a 16-byte identification, or LAN ID, when it is created. The VSWITCH LAN ID consists of two 8-byte character fields: Owner and Name. The Owner component is always SYSTEM. This column is the Name component of the LAN ID.

#### Group

The name of the link aggregation port group in use for this VSWITCH.

#### Vswitch Macaddr

The LAN management MAC address.

#### IP address

The LAN management IP address. This address is used by clients involved in switch management. For example, a Network Management System can request Bridge MIB information from an SNMP agent using this address. If an SNMP agent has not been assigned, the IP address field will contain dots.

#### Userid

The LAN management user ID. This is the user ID of the z/VM TCP/IP stack that has assigned an SNMP agent to the virtual switch in its HOME statement. If an SNMP agent has not been assigned, the Userid field will contain blanks.

#### Layer

The session layer:

#### Layer 2

Ethernet frames

### **Layer 3**

IP datagrams

#### **OSA Port**

The OSA port number specified on the RDEV parameter of the DEFINE VSWITCH command.

For a virtual switch bridge port, you will see dots in in this field.

When running Performance Toolkit against z/VM 5.3 or earlier data, you will see dashes in this field.

#### **Isol**

The isolation level, as specified on the SET VSWITCH or MODIFY VSWITCH command.

When running Performance Toolkit against z/VM 5.3 or earlier data, you will see dashes in this field.

When running Performance Toolkit against z/VM 5.4 data, the value can be FORW, DROP, or OFF.

When running Performance Toolkit against z/VM 6.1 or later data, the value can be ON or OFF.

#### **Config Acc**

The access type for a virtual switch:

##### **U**

User-based. Specifies that the virtual switch will be on a user ID basis. Authorization and configuration will be for the user. Use SET VSWITCH GRANT and REVOKE commands for authorization. Port numbers for guests will be assigned by CP.

##### **P**

Port-based. Specifies that the virtual switch will be port based. All connections will be on a port basis. Each port must be configured with the SET VSWITCH PORTNUMBER command. This configuration will allow a guest to have multiple unique access ports connected to the same VSWITCH instance.

When running Performance Toolkit against z/VM 6.1 or earlier data, you will see dashes in this field.

#### **Config Typ**

The LAN type:

##### **Q**

QDIO

##### **N**

INMN

##### **E**

IEDN

##### **I**

IVL

#### **Config Parm**

The parameters for the global networking object:

##### **L**

The device is the active LAG port controller.

##### **S**

The shared port group in use.

##### **G**

The virtual switch is a global virtual switch.

#### **Role**

The configured bridge port role:

##### **Prim**

Primary

##### **Sec**

Secondary

**State**

The current bridge port state:

**UNDEF**

Undefined

**INACT**

Inactive

**STNBY**

Standby

**ACTIV**

Active

**Flag**

The state of the bridge port connection:

**1**

VSWITCH defined

**2**

Controller not available

**3**

Operator intervention required

**4**

Disconnected

**5**

VDEVs attached to controller

**6**

Initialization in process

**7**

Device not ready

**8**

Device ready

**9**

Devices being detached

**10**

VSWITCH delete pending

**11**

VSWITCH failover recovering

**12**

Autorestart in progress

**MFSs**

CHPID configured MFS size

**LPARname**

When BPBSTATE is Active or Standby, the LPAR name acting as the active bridge port connection. This field contains blanks otherwise.

**Effect of AVERAGE and INTERIM subcommands:**

These subcommands have no effect because all data on the display pertains to current information. There is no activity data on the display.



## FCX267, Extended Virtual Switch Activity Screen – EVSWITCH

**Performance Toolkit Subcommand:** “EVSWITCH” on page 177.

The Extended Virtual Switch Activity Screen (FCX267) shows additional data beyond that which is shown on the Virtual Switch Activity Screen. Note that besides issuing the EVSWITCH subcommand, this screen can also be reached by selecting any row in the Virtual Switch Activity Screen and pressing ENTER. (See “FCX240, Virtual Switch Activity Screen – VSWITCH” on page 707.)

```

FCX267
-----
CPU nnnn SER nnnnn Interval HH:MM:SS - HH:MM:SS Perf. Monitor
-----
<-- Lock and waits /sec --> <----- Total Value for Interval -----> <-Input/s-> <Output/s->
Bal Lock Req Lock Defers TX RX <-NIC Marker PDU Packets-> <- LACPDU -> Buffr Queue Buffr Queue ucode
Int Reqs Defers Defers Defers SENT RSPRCV RSPSNT TIMEO SENT RCVD IPProc IOvfl OProc OOvfl level
>> System >> ..... 1 .0 .0 .0 5 5 1 0 7 202 .8 .0 .2 .0 ...
60D5 PEGIVL 0 1 .0 .0 .0 0 0 0 0 0 0 1.2 .0 .7 .0 07.0F
63D0 GLOB1 300 0 .0 .0 .0 16 16 4 0 33 656 .5 .0 .2 .0 07.0F
63D3 GLOB2 300 1 .0 .0 .0 4 4 0 0 0 213 .9 .0 .0 .0 07.0F
64C0 GLOB1 300 1 .0 .0 .0 4 4 0 0 0 131 1.0 .0 .0 .0 07.0F
64C3 GLOB2 300 1 .0 .0 .0 0 0 0 0 0 10 .9 .0 .0 .0 07.0F

See EVSWITCH / GVSWITCH / GLONACT for other Vswitch reports
Command ==>
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

```

Figure 161. Layout of Extended Virtual Switch Activity Screen (FCX267)

The information shown is based on CP monitor I/O domain SAMPLE data.

The sample shows a screen with virtual switches sorted in ascending device number sequence (indicated by the underscore characters above the 'Addr' header). Other possible sort sequences are indicated by dots above the corresponding columns; they can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER. Note that the information extends beyond the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window so as to view the remaining data.

### Field Description:

#### Addr

Real device number. The first line contains the string ">> System <<" instead of a device number and owner identification. It shows the average load figures for all virtual switches for which performance data was collected.

#### Name

The VSWITCH LAN ID consists of two 8 byte character fields: OWNER and NAME. The OWNER is always SYSTEM. This column is the NAME component of the LAN ID.

#### Load Balance

The load balancing interval. This is the number of seconds between load balancing operations.

#### Lock Requests

Number of lock requests made for the network lock, per second.

#### LAN Defers

Number of times per second that CP waited for the network lock.

#### TX Defers

Number of times per second that CP waited for any lock when sending data from this VSWITCH port.

#### RX Defers

Number of times per second that CP waited for any lock when receiving data on this VSWITCH port.

**PDU Sent**

Number of marker PDUs sent to this port. This is the total value for the monitor interval.

**PDU response**

Number of marker responses received. This is the total value for the monitor interval.

**PDU response**

Number of marker response PDUs sent to this port in response to receiving a marker PDU from the partner port. This is the total value for the monitor interval.

**Timeouts**

Number of times timed out waiting for a marker response PDU for a marker request sent by CP to a partner port. This is the total value for the monitor interval.

**LAC sent**

Number of LACP PDUs sent on this port. This is the total value for the monitor interval.

**LAC rcvd**

Number of LACP PDUs received on this port. This is the total value for the monitor interval.

*Input/s:***Bufrr IProc**

Number of input buffers processed per second.

**Queue IOvfl**

Number of input queue overflows per second.

*Output/s:***Bufrr OProc**

Number of output buffers processed per second.

**Queue OOVfl**

Number of output queue overflows per second.

**ucode level**

OSA device microcode level.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all activity fields.

# FCX268, General Virtual Network Device Description Screen – GVNIC

**Performance Toolkit Subcommand:** “GVNIC” on page 192.

In the General Virtual Network Device Description Screen (FCX268), all of the virtual network devices are shown, by name, without any associated performance data. If you select any row in this screen and press ENTER, you will go to the Virtual Network Device Activity Screen for that device. (See “FCX269, Virtual Network Device Activity Screen – VNIC” on page 757.)

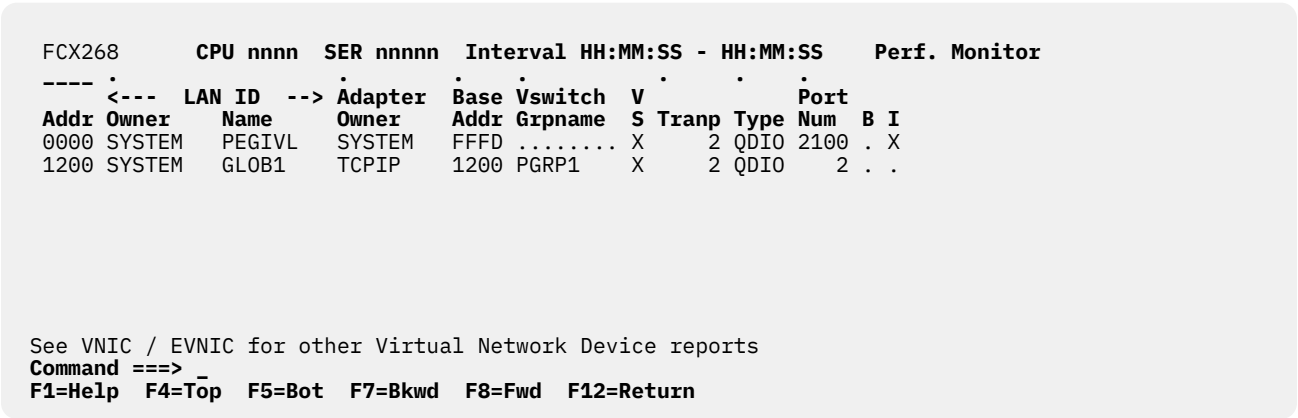


Figure 162. Layout of General Virtual Network Device Description Screen (FCX268)

The sample shows a screen with virtual network devices sorted in ascending device number sequence (indicated by the underscore characters above the 'Addr' header). Other possible sort sequences are indicated by dots above the corresponding columns; they can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

## Field Descriptions:

### Addr

The virtual device address.

### LAN ID

The VNIC LAN ID, consisting of two 8-byte character fields, Owner and Name.

### Adapter Owner

The user ID owning the adapter.

### Base Addr

The NIC base device address.

### Vswitch Grpname

The name of the link aggregation port group in use for this VSWITCH, when the device is a VSWITCH.

### VS

VSWITCH indicator. X indicates that this network device is a VSWITCH.

### Tranp

The transport type used in the communication between the VNIC and CP. This is referred to as the session layer. 2 = Ethernet, 3 = IP.

### Type

Represents the type of the virtual NIC.

### Hiper

HiperSockets

### QDIO

Queued direct I/O adapter (OSD, OSA, OSE)

**INMN**

QDIO - OSM CHPID

**IEDN**

QDIO - OSX CHPID

**VSW**

VSWITCH (QDIO+UPLINK)

**HBRDG**

IQD Hiper-Bridge port

**HIEDN**

Hiper-IEDN device

**????**

Unrecognized type

**Note:** INMN and IEDN are valid only when running Performance Toolkit against data from z/VM 6.1 with the PTF for VM64780, or later. HBRDG is valid only on when running Performance Toolkit against data from z/VM 6.2 with the PTFs for VM65042 and PM46988, or later.

**Port Num**

The port value for the guest connection.

When running Performance Toolkit against z/VM 6.1 or earlier data, you will see dashes in this field.

**B**

Bridge Port indicator. X indicates that this network device is connected to a logical LAN.

**I**

Inter-VSwitch Link (IVL) indicator. X indicates an IVL NIC port.

**Effect of AVERAGE and INTERIM subcommands:**

These subcommands have no effect because all data on the display pertains to current information. There is no activity data on the display.

# FCX269, Virtual Network Device Activity Screen – VNIC

**Performance Toolkit Subcommand:** “VNIC” on page 360.

In the Virtual Network Device Activity Screen (FCX269), the lines for virtual network devices which have been active during the last measuring interval are highlighted on the screen.

Note that besides issuing the VNIC subcommand, this screen can also be reached from the TCP/IP menu by selecting "Virtual Network Devices", or by selecting any row in the General Virtual Network Device Description Screen and pressing ENTER. (See “FCX268, General Virtual Network Device Description Screen – GVNIC” on page 755.)

If you select any row in this screen and press ENTER, you will go to the Extended Virtual Network Device Activity Screen for additional data. (See “FCX270, Extended Virtual Network Device Activity Screen – EVNIC” on page 759.)

FCX269		CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS										Perf. Monitor										
----		.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.		
<--- LAN ID --->		Adapter	Base Vswitch	V	<--- Outbound/s --->					<--- Inbound/s --->					<----- Locks/sec ----->									
Addr	Owner	Name	Owner	Addr	Gxpname	S	L	T	B	Bytes	<	Packets	>	Bytes	<	Packets	>	<	Intrpts	>	Lock	LAN	TX	RX
<<	-----	System	(	2)	-----	>>				T_Byte	T_Pack	T_Disc		R_Byte	R_Pack	R_Disc		Out/s	In/s		Reqs	Defer	Defer	Defer
0000	SYSTEM	PEGIVL	SYSTEM	FFFD	.....	X	2	Q		72	.4	.0		53	.9	.5		.0	.2		.0	.0	.0	.0
1200	SYSTEM	GLOB1	TCPIP	1200	PGRP1	X	2	Q		128	.8	.0		65	1.2	1.0		.0	.0		.0	.0	.0	.0
										15	.1	.0		40	.5	.0		.0	.5		.1	.0	.0	.0

See EVNIC / GVNIC for other Virtual Network Device reports

Command ==>

F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 163. Layout of Virtual Network Device Activity Screen (FCX269)

The information shown is based on CP monitor I/O domain SAMPLE data.

The sample shows a screen with virtual network devices sorted in ascending device number sequence (indicated by the underscore characters above the 'Addr' header). Other possible sort sequences are indicated by dots above the corresponding columns; they can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER. Note that the information extends beyond the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window so as to view the remaining data.

## Addr

The virtual device number. The first line contains the string << System >> instead of a device number and owner identification. It shows the average load figures for all VNICs for which performance data were collected. The string << System >> also includes the total number of VNICs (in parentheses).

## LAN ID

The VNIC LAN ID, consisting of two 8-byte character fields, OWNER and NAME.

## Owner

The userid owning the adapter.

## Base address

NIC base device address.

## Group Name

Name of the link aggregation port group in use for this VSWITCH, when the device is a VSWITCH.

## Vswitch indicator

X indicates that this network device is a VSWITCH.

## **L**

The transport type used in the communication between the VNIC and CP. This is referred to as the session layer. 2 = Ethernet, 3 = IP.

## **T**

Represents the type of the virtual NIC.

### **H**

HiperSockets

### **Q**

Queued direct I/O adapter (OSD)

### **N**

QDIO with OSM CHPID (INMN)

### **E**

QDIO with OSX CHPID (IEDN)

### **B**

IQD bridge port (BRDG)

### **?**

Unrecognized type

Note: INMN and IEDN are valid only on z/VM 6.1 with the PTF for VM64780 or higher. BRDG is valid only on z/VM 6.2 with the PTFs for VM65042 and PM46988 or higher.

## **B**

Bridge port indicator. X indicates that this network device is connected to a logical LAN.

## **T\_Byte**

Bytes sent to OSA per second.

## **T\_Pack**

Packets sent to OSA per second.

## **T\_Disc**

Outbound packets discarded per second,

## **R\_Byte**

Bytes received from OSA per second.

## **R\_Pack**

Packets received from OSA per second.

## **R\_Disc**

Inbound packets discarded per second

## **Int out/s**

Number of PCI (or AIF) events reflected for output queues per second.

## **Int in/s**

Number of PCI (or AIF) events reflected for input queues per second.

## **Lock Requests**

Number of lock requests made for the network lock, per second.

## **LAN Defer**

Number of times per second that CP waited for the network lock.

## **TX Defer**

Number of times per second that CP waited for any lock when sending data from this port.

## **RX Defer**

Number of times per second that CP waited for any lock when receiving data on this port.

## **Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages will be shown for all activity fields.

# FCX270, Extended Virtual Network Device Activity Screen – EVNIC

**Performance Toolkit Subcommand:** “EVNIC” on page 176.

The Extended Virtual Network Device Activity Screen (FCX270) shows additional data beyond that which is shown on the Virtual Network Device Activity Screen. Note that besides issuing the EVNIC subcommand, this screen can also be reached by selecting any row in the Virtual Network Device Activity Screen and pressing ENTER. (See “FCX269, Virtual Network Device Activity Screen – VNIC” on page 757.)

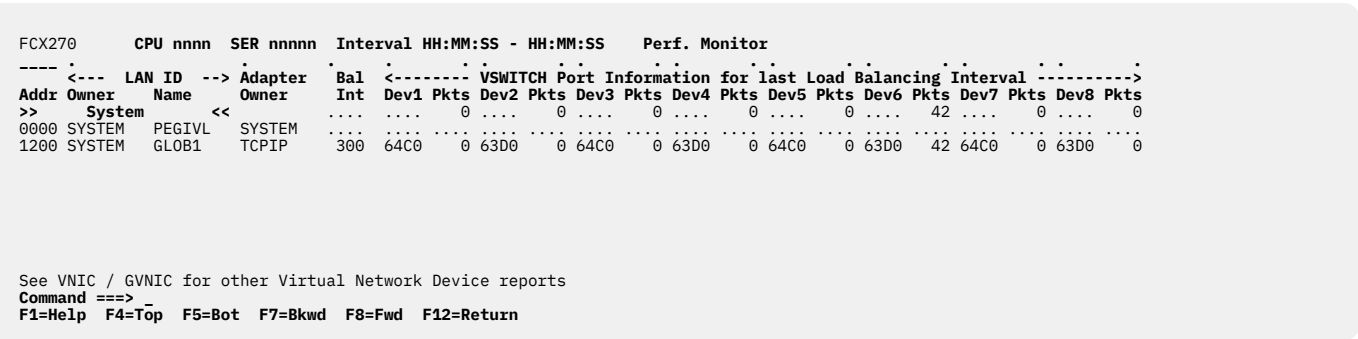


Figure 164. Layout of Extended Virtual Network Device Activity Screen (FCX270)

The information shown is based on CP monitor I/O domain SAMPLE data.

The sample shows a screen with virtual network devices sorted in ascending device number sequence (indicated by the underscore characters above the 'Addr' heading). Other possible sort sequences are indicated by dots above the corresponding columns; they can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER. Note that the information extends beyond the 80 columns of an ordinary screen, use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window so as to view the remaining data.

### Field Description:

#### Addr

The virtual device number.

#### LAN ID

The VNIC LAN ID consists of two 8-byte character fields, OWNER and NAME.

#### Adapter Owner

The userid owning the adapter.

#### Load Balance

Number of seconds between load balancing operations.

#### Address/packets

The port information for the last load balancing interval is displayed for each OSA device. The OSA device number is displayed under the Devn column where n is a value 1-8 to reflect up to 8 ports. The Pkts column displays the number of packets transmitted on this port in the last load balancing interval.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all activity fields.

# FCX271, SSI Data Menu Screen – SSIMENU

**Performance Toolkit Subcommand:** “SSIMENU” on page 298.

The SSI Data Menu Screen (FCX271) displays a selection menu of available SSI / ISFC reports.

```
FCX271      CPU nnnn  SER nnnnn  SSI Data Menu      Perf. Monitor

  SSI performance reports
S Command Description
- SSICONF  SSI configuration
- SSISCHLG SSI State Change Synchronization Activity log
- SSISMILG SSI State/Mode Information log

  ISFC performance reports
S Command Description
- ISFECONF ISFC End Point configuration
- ISFEACT  ISFC End Point activity
- ISFLCONF ISFC Logical Link configuration
- ISFLACT  ISFC Logical Link activity state
- ISFLALOG ISFC Logical Link activity log

  LGR performance reports
S Command Description
- LGRDATA  Live Guest Relocation data
- LGRELOG  Live Guest Relocation event log

Command ==> _
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 165. Layout of SSI Data Menu Screen (FCX271)

**Field Description:**

**S**  
Cursor selection field. The field will be highlighted and made cursor-selectable (set as 'unprotected') only if the corresponding SSI, ISFC, or LGR performance screen is available.

**Command**  
The command name equal to the screen name that can be selected for display. The field will be highlighted if the corresponding SSI, ISFC, or LGR performance screen is available.

**Description**  
Description of the performance display.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.



## FCX272, ISFC End Point Configuration Screen – ISFECONF

**Performance Toolkit Subcommand:** “ISFECONF” on page 205.

The ISFC End Point Configuration Screen (FCX272) displays the ISFC endpoints present on this system. There is one row for each endpoint.

```
FCX272      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor

Initial Status on 2020/08/13 at 16:03:14

EndPoint Typ Act PortName AssocNam OthrNode OthrPort <----- End point name ----->
00000001 DGR BND *SYNC --- --- --- 'Synchronization services'
00000002 DGR BND *PLMVOTE --- --- --- ''
00000003 DGR BND *PLMDCSN --- --- --- ''
00000004 DGR BND *PLMSUSP --- --- --- ''
00000005 DGR BND *PLMBITM --- --- --- ''
00000006 DGR BND *PLMBEAT --- --- --- ''
00000007 LSN BND *XSCIF --- --- --- 'Cross system SCIF listen'
00000008 LSN BND *QITEMS --- --- --- 'Shared SPOOL et al listen'
00000009 DGR BND *XDISK --- --- --- ''
0000000A DGR BND *NETWORK --- --- --- 'SSI Network Services'
0000000B LSN BND *LGR --- --- --- 'Live Guest Relocation Listen'
0000000C DGR BND *RELODOM --- --- --- 'Relocation Domain Synch Port'
0000000D LSN BND *AT --- --- --- 'AT command listen'
0000000E DGR BND *USRBKS --- --- --- 'USRBK synchronization'
0000000F DGR BND *LXCHECK --- --- --- 'LOGON cross check'
00000010 DGR BND *RPI --- --- --- ''

Changed Config Status
Date Time Changed
..... No status changes received

Command ==> _
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return
```

Figure 166. Layout of ISFC End Point Configuration Screen (FCX272)

The information shown is based on MONITOR domain (SAMPLE) and ISFC domain (EVENT) data. It is initially provided by a MONITOR domain record when Performance Toolkit connects to the \*MONITOR service. If the program has been set up to collect the monitor data from another machine's disk file, the program will try to obtain the monitor domain data from the first monitor sample in the disk file; the data cannot be shown if the required information is not found in that first sample.

ISFECONF primarily uses the following system records:

- Domain 1, Record 22 - MRMTRISC, ISFC End Point Configuration, Sample Config
- Domain 9, Record 1 - MRISFISC, ISFC End Point Status Change, Event

### Field Descriptions:

#### {Initial|Refresh} Status on yyyy/mm/dd at hh:mm:ss

Indicates when the status was obtained, that is, when monitoring was last activated (or reactivated). Initial Status is displayed the first time status information is shown. Refresh Status is displayed when a new configuration record arrives.

#### Endpoint

The unique sequence number that is associated with this endpoint.

#### Typ

The endpoint type: Datagram (DGR), Listen (LSN), or Stream (STR).

#### Act

The endpoint activity: Accept (ACC), Bind (BND), Close (CLS), or Connect (CON).

**PortName**

The bound port name, port name, or name generated by CP if no Bind.

**AssocNam**

The name of the service (PortName) that resulted in the creation of this endpoint.

**OthrNode**

The node at the other end of the connection that is represented by this endpoint.

**OthrPort**

The service at the other end of the connection that is represented by this endpoint.

**End point name**

The free-form name (up to 120 characters) that is assigned to this endpoint by its creator.

**Changed Config Status**

Indicates the date, time, and nature of any changes in the state of an ISFC endpoint.

**Date**

The date of the status change, in *mm/dd* format.

**Time**

The time of the status change, in *hh:mm* format.

**Changed**

Describes how the status changed, as follows:

```
EndPoint= endpoint typ/act
OthrNode=node
OthrPort=port AssocNam=name
PortName=port Name= tttt
```

where:

**endpoint**

Is the 4-byte binary numeric unique identifier that is associated with this endpoint.

**typ**

Is the endpoint type: Datagram (DGR), Listen (LSN), or Stream (STR).

**act**

Is the endpoint activity: Accept (ACC), Bind (BND), Close (CLS), or Connect (CON).

**node**

Is the name of the node on the other end of the connection.

**port**

Is a CP-generated EBCDIC character string name for the endpoint.

**name**

For an ACC event only, this is the port name of the associated Listen endpoint. For all other events, this field is not meaningful.

**tttt**

Is a free-form name (up to 120 characters) that is assigned to this endpoint.

**Messages:**

These messages could appear in the Changed Config Status area:

**Not applicable in Trend Scan mode**

is displayed when viewing data using the TRNDSCAN command. Because the corresponding change status events are not written to the history files, they are not available for redisplay.

**... (oldest lines not displayed) ...**

is displayed when the count of the change status lines exceeds the REDISP number (set using the FC MONCOLL REDISP command).

**\*\*\* Status refreshed \*\*\***

is displayed when the new configuration record for this particular report has been received.

**Effect of AVERAGE and INTERIM subcommands:**

None.

## FCX273, ISFC End Point Activity Screen – ISFEACT

**Performance Toolkit Subcommand:** “ISFEACT” on page 204.

The ISFC End Point Activity Screen displays the traffic on ISFC Transport, by endpoint. There is one row for each endpoint. The information shown is based on ISFC domain (SAMPLE, EVENT) data.

FCX273	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor							
-----	.	<-----	Sent	----->	<-----	Received	----->	<--	Others	-->	...	
EndPoint	LifeT	SByte	SBRat	SMsgs	SMRat	RByte	RBRat	RMsgs	RMRat	Buffd	Discd	...
00000001	...	8928	37.20	62	.26	9072	37.80	63	.26	0	0	
00000002	...	586	2.44	7	.03	716	2.98	7	.03	0	0	
00000003	...	160	.67	2	.01	848	3.53	8	.03	0	0	
00000004	...	0	.00	0	.00	0	.00	0	.00	0	0	
00000005	...	952	3.97	1	.00	952	3.97	1	.00	0	0	
00000006	...	848	3.53	8	.03	1590	6.62	15	.06	0	0	
00000007	...	0	.00	0	.00	80	.33	1	.00	0	0	
00000008	...	0	.00	0	.00	80	.33	1	.00	0	0	
00000009	...	22952	95.63	167	.70	22016	91.73	168	.70	0	0	
0000000A	...	4536	18.90	3	.01	4536	18.90	3	.01	0	0	
0000000B	...	0	.00	0	.00	0	.00	0	.00	0	0	
0000000C	...	80	.33	1	.00	112	.47	1	.00	0	0	
0000000D	...	0	.00	0	.00	0	.00	0	.00	0	0	
0000000E	...	912	3.80	10	.04	1328	5.53	11	.05	0	0	
0000000F	...	0	.00	0	.00	0	.00	0	.00	0	0	
00000010	...	87	.36	1	.00	94	.39	1	.00	0	0	
00000013	...	80	.74	1	.01	80	.74	1	.01	0	0	
00000014	...	80	.74	1	.01	80	.74	1	.01	0	0	
00000015	...	0	.00	0	.00	0	.00	0	.00	0	0	
00000016	...	96	.40	1	.00	96	.40	1	.00	0	0	
00000017	20.00	624	31.20	3	.15	256	12.80	3	.15	0	0	
00000018	137.2	80	.58	1	.01	0	.00	0	.00	0	0	
00000019	137.2	912	6.65	8	.06	968	7.06	8	.06	0	0	

Command ==> \_  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 167. Layout of ISFC End Point Activity Screen (FCX273)

The sample shows a screen with endpoints sorted in ascending unique sequence number associated with this endpoint sequence (indicated by the underscore characters above the EndPoint heading), the initial sequence. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

ISFEACT primarily uses the following System records:

- Domain 9, Record 1 - MRISFISC - ISFC End Point Status Change (Event)
- Domain 9, Record 2 - MRISFISA - ISFC End Point Activity (Sample)

### Field Descriptions:

#### Endpoint

The unique sequence number associated with this endpoint.

#### LifeT

The total time, in seconds, of endpoint existence (from the time of creation up to closure, or up to present time if the endpoint is not closed yet). For endpoints for which Performance Toolkit can't determine moment of creation, the field contains dots.

#### Sent

Includes the following information:

**SByte**

The number of bytes sent.

**SBRat**

The number of bytes sent per second.

**SMsgs**

The number of messages sent.

**SMRat**

The number of messages sent per second.

**Received**

Includes the following information:

**RByte**

The number of bytes received.

**RBRat**

The number of bytes received per second.

**RMsgs**

The number of messages received.

**RMRat**

The number of messages received per second.

**Others**

Includes the following information:

**Buffd**

Sent messages awaiting recovery.

**Discd**

Messages that could not be sent, typically due to the loss of a link.

**Description**

The short description of the end point (40 characters or less). The description contains the endpoint name, if present. Otherwise, Performance Toolkit generates a temporary name.

For closed endpoints, the description contains the flag (CLS) before all text.

**Effect of AVERAGE and INTERIM subcommands:**

Averages will be shown for all activity fields as requested.

# FCX274, ISFC Logical Link Activity Screen – ISFLACT

**Performance Toolkit Subcommand:** “ISFLACT” on page 207.

The ISFC Logical Link Activity Screen displays ISFC logical link transport activity. There is one row for each ISFC logical link. The information shown is based on ISFC domain (SAMPLE, EVENT) data.

FCX274	CPU	nnnn	SER	nnnnn	Interval	HH:MM:SS	-	HH:MM:SS	Perf.	Monitor
-----	.	.	.	.	.	.	.	.	.	.
	<----- Sent ----->				<----- Received ----->					
Partner	Devs	LSByte	LSBRt	LSMsg	LSMR	LRByte	LRBRt	LRMsg	LRMR	
AAACH00	1.0	71	1.18	1	.02	71	1.18	1	.02	
BCT12NDA	1.0	71	1.18	1	.02	71	1.18	1	.02	
CFT2NDA	1.0	71	1.18	1	.02	71	1.18	1	.02	
GDLBCT1	1.0	654	10.89	4	.07	766	12.75	4	.07	
GDLLCPX1	9.0	487	8.11	3	.05	571	9.50	3	.05	
GDLRCTS1	1.0	487	8.11	3	.05	571	9.50	3	.05	
GDLRCT1	1.0	487	8.11	3	.05	571	9.50	3	.05	

Command ==> \_  
**F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return**

Figure 168. Layout of ISFC Logical Link Activity Screen (FCX274)

The sample shows a screen with logical links sorted in ascending partner node name sequence (indicated by the underscore characters above the Partner heading), the initial sequence. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

ISFLACT primarily uses the System records:

- Domain 9, Record 3 - MRISFILC - ISFC Logical Link Definition Change (Event)
- Domain 9, Record 4 - MRISFNOD - ISFC Logical Link Activity (Sample)

## Field Description:

### Partner

The node on the other end of this logical link.

### Devs

The number of real devices comprising the logical link to the specified node.

### Sent

Includes the following information:

#### LSByte

The number of bytes sent on this link originating from this system.

#### LSBRt

The number of bytes sent on this link originating from this system per second.

#### LSMsg

The number of messages sent on this link originating from this system.

#### LSMR

The number of messages sent on this link originating from this system per second.

### Received

Includes the following information:

**LRByte**

The number of bytes received on this link destined for this system.

**LRBRt**

The number of bytes received on this link destined for this system per second.

**LRMsg**

The number of messages received on this link destined for this system.

**LRMR**

The number of messages received on this link destined for this system per second.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Averages will be shown for all activity fields as requested.

## FCX275, ISFC Logical Link Configuration Screen – ISFLCONF

**Performance Toolkit Subcommand:** “ISFLCONF” on page 209.

The ISFC Logical Link Configuration Screen displays the configuration of ISFC logical links. There is one row for each ISFC logical link. The information shown is based on MONITOR domain (SAMPLE) and ISFC domain (EVENT) data.

```
FCX275      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor

Initial Status on 2020/12/05 at 11:55

Partner  Devs  Rdevs
CFT2NDA   1    96A1
PLXA01    0    ...
PLXA02    1    912A
Changed Config Status
Date Time      Changed
12/06 16:06:28  Rdev 912A has been deactivated for Node PLXA02 (last)

Not applicable in Trend Scan mode
... (oldest lines not displayed) ...
*** Status refreshed ***

Command ==>
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 169. Layout of ISFC Logical Link Configuration Screen (FCX275)

The information shown is initially provided by a MONITOR domain record when Performance Toolkit connects to the \*MONITOR service. If the program has been set up to collect the monitor data from another machine's disk file, the program will attempt to obtain the monitor domain data from the first monitor sample in the disk file; the data cannot be shown if the required information is not found in that first sample.

ISFLCONF primarily uses the following System records:

- Domain 1, Record 23 - MRMTRILC - ISFC Logical Link Configuration Record (Sample configuration)
- Domain 9, Record 3 - MRISFILC - ISFC Logical Link Definition Change (Event)

The primary configuration of ISFC Logical Links is obtained from Domain 1 Record 23 monitor records. Performance Toolkit updates this screen according to logical link configuration changes expressed in Domain 9 Record 3 monitor records.

### Field Descriptions:

#### {Initial|Refresh} Status on yyyy/mm/dd at hh:mm:ss

Indicates when the status was obtained, that is, when monitoring was last activated (or reactivated). Initial Status is displayed the first time status information is shown. Refresh Status is displayed when a new configuration record arrives.

#### Partner

The partner node name.

#### Devs

The number of real devices comprising the logical link to the specified node.

#### Rdevs

The set of real devices comprising the logical link to the specified node.

#### Changed Config Status

Indicates the date, time, and nature of any changes in the state of an ISFC logical link.

#### Date

The date of the status change, in mm/dd format.



**Time**

The time of the status change, in *hh:mm* format.

**Changed**

Describes how the status changed, as follows:

```
Rdev xxxx has been activated Node=node  
Rdev xxxx has been activated Node=node (first)  
Rdev xxxx has been deactivated Node=node  
Rdev xxxx has been deactivated Node=node (last)
```

where:

***xxxx***

Is the real device number of the ISFC link device comprising the logical link to the specified node.

***node***

Is the partner node name.

**Messages:**

These messages could appear in the Changed Config Status area:

**Not applicable in Trend Scan mode**

is displayed when viewing data using the TRNDSCAN command. Because the corresponding change status events are not written to the history files, they are not available for redisplay.

**... (oldest lines not displayed) ...**

is displayed when the count of the change status lines exceeds the REDISP number (set using the FC MONCOLL REDISP command).

**\*\*\* Status refreshed \*\*\***

is displayed when the new configuration record for this particular report has been received.

**Effect of AVERAGE and INTERIM subcommands:**

None.

# FCX276, SSI Configuration Screen – SSICONF

**Performance Toolkit Subcommand:** “SSICONF” on page 297.

The SSI Configuration Screen displays the SSI configuration of this system. The information shown is based on MONITOR domain (SAMPLE) and SSI domain (EVENT) data.

```
FCX276      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor

Initial Status on 2020/12/09 at 11:55

SSI Name                DWSSI
Number of slots configured 4
Number of slots in use    2

Members Information
Slot SystemID
1 PLXA01
2 PLXA02
3 ...
4 ...

Changed Config Status
Date Time Changed
12/10 09:58:16 System PLXA02 state: Joined -> Unknown
                SSI mode: Stable -> Safe
12/10 09:58:16 System PLXA01 state: Joined -> Suspended
                SSI mode: Safe

Not applicable in Trend Scan mode
... (oldest lines not displayed) ...
*** Status refreshed ***

Command ==>_
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 170. Layout of SSI Configuration Screen (FCX276)

The information shown is initially provided by a MONITOR domain record when Performance Toolkit connects to the \*MONITOR service. If the program has been set up to collect the monitor data from another machine's disk file, the program will attempt to obtain the monitor domain data from the first monitor sample in the disk file; the data cannot be shown if the required information is not found in that first sample.

SSI configuration uses the following system records:

- Domain 1, Record 25 - MRMTRSSI, SSI Configuration, Sample Config
- Domain 11, Record 3 - MRSSISCH, State Change Event, Event
- Domain 11, Record 4 - MRSSISLT, Slot Definition, Event

**Field Descriptions:**

**{Initial|Refresh} Status on yyyy/mm/dd at hh:mm:ss**

Indicates when the status was obtained, that is, when monitoring was last activated (or reactivated). Initial Status is displayed the first time status information is shown. Refresh Status is displayed when a new configuration record arrives.

**SSI Name**

The SSI name as defined in the system configuration.

**Number of slots configured**

The number of slots configured for the system members.

**Number of slots in use**

The number of slots currently used by the system members.

**Slot**

The slot number.

**SystemID**

The member system name.

**Changed Config Status**

Indicates the date, time, and nature of any changes in the SSI state.

**Date**

The date of the status change, in *mm/dd* format.

**Time**

The time of the status change, in *hh:mm* format.

**Changed**

Describes how the status changed, as follows:

```
System name state: state1 -> state2  
SSI mode: mode1 -> mode2
```

where:

***name***

Is the member system name.

***state1***

Is the old (or current, if *state2* is missing) member system state.

***state2***

Is the new member system state.

***mode1***

Is the old (or current, if *mode2* is missing) SSI mode.

***mode2***

Is the new SSI mode.

Possible system states are: Down, Isolated, Joined, Joining, Leaving, Not valid, Suspended, and Unknown (xx), where xx is the hexadecimal state value.

Possible SSI modes are: Influx, Not\_In\_SSI, Safe, Stable, and Unknown (xx), where xx is the hexadecimal mode value.

**Messages:**

These messages could appear in the Changed Config Status area:

**Not applicable in Trend Scan mode**

is displayed when viewing data using the TRNDSCAN command. Because the corresponding change status events are not written to the history files, they are not available for redisplay.

**... (oldest lines not displayed) ...**

is displayed when the count of the change status lines exceeds the REDISP number (set using the FC MONCOLL REDISP command).

**\*\*\* Status refreshed \*\*\***

is displayed when the new configuration record for this particular report has been received.

**Effect of AVERAGE and INTERIM subcommands:**

None.

# FCX277, SSI State Change Synchronization Activity Log Screen – SSISCHLG

**Performance Toolkit Subcommand:** “SSISCHLG” on page 299.

The SSI State Change Synchronization Activity Log Screen displays the current SSI state change synchronization activity. The information shown is based on SSI domain SAMPLE data.

FCX277	CPU	nnnn	SER	nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor
Interval	Init	<- Count	of Xmit	->	<- Sync	Duration	->
End Time	Sync	PWRep	Proc	Abrt	PWRep	P&Ab	Vote
>>Mean>>	.0	.0	.0	.0	.0	.0	.0
16:59:16	0	0	0	0	0	0	0
16:59:31	0	0	0	0	0	0	0
16:59:46	0	0	0	0	0	0	0
Command ==> _							
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return		

Figure 171. Layout of SSI State Change Synchronization Activity Log Screen (FCX277)

SSI State Change Synchronization Activity log uses the following System record:

- Domain 11, Record 1 - MRSSISCS, State Change Synchronization Activity, Sample

**Field Description:**

**Interval End Time**

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

**Init Synch**

Cumulative number of times synchronization was initiated.

**Count of Xmit**

Include the following information:

**PWRep**

Cumulative number of times Proceed With Reply was transmitted.

**Proc**

Cumulative number of times Proceed was transmitted.

**Abort**

Cumulative number of times Abort was transmitted.

**Sync Duration**

Includes the following information:

**PWRep**

Cumulative duration in microseconds of Proceed With Reply phases.

**P&Ab**

Cumulative duration in microseconds of Proceed and Abort phases.

**Vote**

Cumulative duration in microseconds of Vote phases.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX278, SSI State/Mode Information Log Screen – SSISMILG

**Performance Toolkit Subcommand:** “SSISMILG” on page 300.

The SSI State/Mode Information Log Screen displays SSI configuration of this system. The information shown is based on SSI domain SAMPLE data.

FCX278	CPU nnnn		SER nnnnn		Interval HH:MM:SS - HH:MM:SS						Perf. Monitor							
<----- SSI State ----->															<-- SSI Mode -->			
Interval	<-Joining-->		<--Joined-->		<-Leaving-->		<-Isolated-->		<-Suspended-->		<--Stable-->		...					
End Time	Count	%Time	Count	%Time	Count	%Time	Count	%Time	Count	%Time	Count	%Time	Count	%Time	...			
>>Mean>>	.0	.000	.3	73.955	.0	.000	.0	.000	.0	26.045	0	0	0	0	...			
16:59:16	0	.000	1	21.867	0	.000	0	.000	0	78.133	0	0	0	0	...			
16:59:31	0	.000	0	100.00	0	.000	0	.000	0	.000	0	0	0	0	...			
16:59:46	0	.000	0	100.00	0	.000	0	.000	0	.000	0	0	0	0	...			
Command ==> _																		
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return													

Figure 172. Layout of SSI State/Mode Information Log Screen (FCX278)

SSI State/Mode Information Log uses the following System record:

- Domain 11, Record 2 - MRSSISMI, State/Mode Information, Sample

Note that the information extends beyond the 80 columns of an ordinary screen. Use the LEFT/RIGHT subcommands, or the corresponding PF-keys, for shifting the display window so as to view the remaining data.

## Field Descriptions:

### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

### SSI State Joining

Includes the following information:

#### Count

Cumulative number of times joining state entered.

#### %Time

Total percentage of time spent in joining state.

### SSI State Joined

Includes the following information:

#### Count

Cumulative number of times joined state entered.

#### %Time

Total percentage of time spent in joined state.

### SSI State Leaving

Includes the following information:

#### Count

Cumulative number of times leaving state entered.

#### %Time

Total percentage of time spent in leaving state.

### SSI State Isolated

Includes the following information:

#### Count

Cumulative number of times isolated state entered.

**%Time**

Total percentage of time spent in isolated state.

**SSI State Suspended**

Includes the following information:

**Count**

Cumulative number of times suspended state entered.

**%Time**

Total percentage of time spent in suspended state.

**SSI Mode Stable**

Includes the following information:

**Count**

Cumulative number of times stable mode was entered.

**%Time**

Total percentage of time spent in stable mode.

**SSI Mode Influx**

Includes the following information:

**Count**

Cumulative number of times influx mode was entered.

**%Time**

Total percentage of time spent in influx mode.

**SSI Mode Safe**

Includes the following information:

**Count**

Cumulative number of times safe mode was entered.

**%Time**

Total percentage of time spent in safe mode.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX279, ISFC End Point Activity Log Screen – ISFELOG

**Performance Toolkit Subcommand:** “ISFELOG ” on page 206.

The ISFC End Point Activity Log Screen shows a "per sample" log of selected ISFC endpoints. When requesting an ISFC endpoint activity log for an ISFC endpoint, the *epoint* must have been previously selected for benchmarking and for which monitor records should be collected.

Figure 173 on page 775 shows a sample ISFC endpoint activity log for endpoint 00000001:

FCX279      CPU nnnn SER nnnnn Interval HH:MM:SS - HH:MM:SS Perf. Monitor

Activity log for End Point  
00000001

Interval		<----- Sent ----->				<----- Received ----->				<-- Others -->		...
End Time	LifeT	SByte	SBRat	SMsgs	SMRat	RByte	RBRat	RMsgs	RMRat	Buffd	Discd	...
>>Mean>>	...	9216	19.20	64	.13	9216	19.20	64	.13	0	0	
02:19:26	...	1152	19.24	8	.13	1152	19.24	8	.13	0	0	
02:20:26	...	1152	19.16	8	.13	1152	19.16	8	.13	0	0	
02:21:26	...	1152	19.23	8	.13	1152	19.23	8	.13	0	0	
02:22:26	...	1152	19.20	8	.13	1152	19.20	8	.13	0	0	
02:23:26	...	1152	19.20	8	.13	1152	19.20	8	.13	0	0	
02:24:26	...	1152	19.20	8	.13	1152	19.20	8	.13	0	0	
02:25:26	...	1152	19.19	8	.13	1152	19.19	8	.13	0	0	
02:26:26	...	1152	19.21	8	.13	1152	19.21	8	.13	0	0	

Command ==>>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return

Figure 173. Layout of ISFC End Point Activity Log Screen (FCX279)

The information shown is based on MONITOR and ISFC domain SAMPLE and EVENT data.

The leftmost column contains the end time (in *hh:mm:ss* format) of the monitor interval for which the values were calculated.

The first entry is reserved for averages since the last RESET subcommand or since data collection was activated. It is designated by the string >>Mean>> in the Interval End Time column.

**Field Descriptions:**

**Interval End Time**

The end time of the monitor interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the Interval End Time field, which replaces the EPoint field, all data columns are identical to the fields in FCX273. See “FCX273, ISFC End Point Activity Screen – ISFEACT” on page 764 for the descriptions of these fields.

**Effect of AVERAGE and INTERIM subcommands:**

None.

# FCX280, ISFC Logical Link Activity Log Screen – ISFLLOG

**Performance Toolkit Subcommand:** “ISFLLOG” on page 210.

The ISFC Logical Link Activity Log Screen displays a 'per sample' log of the selected ISFC logical link. When requesting an ISFC logical link activity log for an ISFC logical link, the *link* must have been previously selected for benchmarking and for which monitor records should be collected.

Figure 174 on page 776 shows a 'per sample' log for the logical link CFT2NDA1:

FCX280      **FCX280 CPU nnnn SER nnnnn Interval HH:MM:SS - HH:MM:SS Perf. Monitor**

Activity log for logical link CFT2NDA

Interval		<----- Sent ----->				<----- Received ----->			
End Time	Devs	LSByte	LSBRt	LSMsg	LSMR	LRByte	LRBRt	LRMsg	LRMR
>>Mean>>	1.0	48191	114.8	167	.40	20456	48.72	164	.39
02:19:26	1.0	0	.00	0	.00	0	.00	0	.00
02:20:26	1.0	4809	79.99	47	.78	4800	79.84	48	.80
02:21:26	1.0	24688	412.2	57	.95	8091	135.1	54	.90
02:22:26	1.0	7397	123.3	14	.23	1897	31.62	14	.23
02:23:26	1.0	11226	187.1	48	.80	5597	93.29	47	.78
02:24:26	1.0	71	1.18	1	.02	71	1.18	1	.02
02:25:26	1.0	0	.00	0	.00	0	.00	0	.00

Command ==> \_  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 174. Layout of ISFC Logical Link Activity Log Screen (FCX280)

The information shown is based on MONITOR and ISFC domain SAMPLE and EVENT data.

The leftmost column contains the end time (in the format *hh:mm:ss*) of the monitor interval for which the values were calculated.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the Interval End Time column.

**Field Descriptions:**

**Interval End Time**

The end time of the monitor interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the Interval End Time field, which replaces the Partner field, all data columns are identical to the fields in FCX274. See “FCX274, ISFC Logical Link Activity Screen – ISFLACT” on page 766 for the descriptions of these fields.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.



## FCX281, ISFC Logical Link Activity By-Time Log Screen – ISFLALOG

**Performance Toolkit Subcommand:** “ISFLALOG” on page 208.

The ISFC Logical Link Activity By-Time Log Screen displays overall performance data for all ISFC Logical Links existing in the system, by time. Each entry consists of a group of lines for every ISFC Logical Link per interval, as shown in [Figure 175 on page 777](#).

FCX281 FCX281 CPU nnnn SER nnnnn Interval HH:MM:SS - HH:MM:SS Perf. Monitor											
Interval			<----- Sent ----->				<----- Received ----->				
End Time	Partner	Devs	LSByte	LSBrt	LSMsg	LSMR	LRByte	LRBrt	LRMsg	LRMR	
>>Mean>>	AAACH00	1.0	2152	7.18	21	.07	1815	6.05	19	.06	
>>Mean>>	BCT12NDA	1.0	2212	7.38	22	.07	1815	6.05	19	.06	
>>Mean>>	CFT2NDA	1.0	48120	160.5	166	.55	20385	67.98	163	.54	
>>Mean>>	GDLBCT1	1.0	4832	16.11	27	.09	4823	16.08	24	.08	
>>Mean>>	GDLCPX1	9.0	14251	47.52	72	.24	15042	50.16	69	.23	
>>Mean>>	GDLRCTS1	1.0	7361	24.55	39	.13	7652	25.52	36	.12	
>>Mean>>	GDLRCT1	1.0	436751	1456	70	.23	28386	94.66	69	.23	
02:20:26	AAACH00	1.0	689	5.74	7	.06	622	5.18	7	.06	
02:20:26	BCT12NDA	1.0	678	5.65	7	.06	491	4.09	5	.04	
02:20:26	CFT2NDA	1.0	4809	40.08	47	.39	4800	40.00	48	.40	
02:20:26	GDLBCT1	1.0	1739	14.49	10	.08	1592	13.27	8	.07	
02:20:26	GDLCPX1	9.0	4522	37.69	23	.19	4963	41.36	23	.19	
02:20:26	GDLRCTS1	1.0	2575	21.46	14	.12	2708	22.57	13	.11	
02:20:26	GDLRCT1	1.0	231687	1931	33	.28	14998	125.0	33	.28	
02:23:26	AAACH00	1.0	1463	8.13	14	.08	1193	6.63	12	.07	
02:23:26	BCT12NDA	1.0	1534	8.53	15	.08	1324	7.36	14	.08	
02:23:26	CFT2NDA	1.0	43311	240.8	119	.66	15585	86.63	115	.64	
02:23:26	GDLBCT1	1.0	3093	17.19	17	.09	3231	17.96	16	.09	
02:23:26	GDLCPX1	9.0	9729	54.08	49	.27	10079	56.03	46	.26	
02:23:26	GDLRCTS1	1.0	4786	26.60	25	.14	4944	27.48	23	.13	
02:23:26	GDLRCT1	1.0	205064	1140	37	.21	13388	74.42	36	.20	

Command ==>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 175. Layout of ISFC Logical Link Activity By-Time Log Screen (FCX281)

The information shown is based on MONITOR and ISFC domain SAMPLE and EVENT data.

The leftmost column contains the end time (in the format *hh:mm:ss*) of the monitor interval for which the values were calculated.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the Interval End Time column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Except for the Interval End Time field, all data columns are identical to the fields in FCX274. See “FCX274, ISFC Logical Link Activity Screen – ISFLACT” on page 766 for the descriptions of these fields.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX282, HPF I/O Device Screen – DEVICE HPF

**Performance Toolkit Subcommand:** “DEVICE” on page 169.

The HPF I/O Device Screen (FCX282) is shown when you enter the DEVICE subcommand with the HPF parameter. Only zHPF-capable DASDs are included on this screen. Devices that were active during the last measuring interval are highlighted on the screen.

FCX282		CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS		Perf. Monitor												
<-- Device Descr. -->			Mdisk	Pa-	<----- zHPF ----->		<----- Time (msec) ----->						Req.	-----		Pct	LCU	
Addr	Type	Label/ID	Links	ths	Mask	Featr	Read/s	Write/s	Pend	Disc	Conn	Serv	Resp	CUWt	Qued	I/O	Busy	SSID
>>A11 DASD<< ( 3729)			.....	...	..	..	.1	.0	.252	11.6	.744	12.6	22.0	.047	.01	1.3	2	...
C50D	3390-9	ATP111 CP	0	4	F0	EM	18.9	3.4	.252	11.6	.685	12.5	21.4	.000	.21	22.4	28	0607
CD0A	3390-9	ATP220 CP	0	4	F0	EM	18.9	3.3	.257	12.2	.700	13.1	25.3	.000	.27	22.3	29	060F
C707	3390-9	ATP133 CP	0	4	F0	EM	18.9	3.3	.250	11.4	.708	12.4	23.1	.000	.25	22.3	28	0609
C701	3390-9	ATP127 CP	0	4	F0	EM	18.9	3.3	.251	11.6	.656	12.5	24.1	.000	.27	22.3	28	0609
BE00	3390-9	ATP000 CP	0	4	F0	EM	14.2	.8	.264	13.0	.564	13.9	13.9	.087	.00	15.1	21	0600
BE02	3390-9	ATP002 CP	0	4	F0	EM	14.2	.9	.259	11.8	.566	12.6	12.7	.053	.00	15.2	19	0600
BE08	3390-9	ATP008 CP	0	4	F0	EM	14.2	.8	.258	11.7	.560	12.5	12.6	.047	.00	15.1	19	0600
BE07	3390-9	ATP007 CP	0	4	F0	EM	14.1	.8	.272	13.2	.587	14.0	14.1	.081	.00	15.0	21	0600
BE06	3390-9	ATP006 CP	0	4	F0	EM	14.1	.8	.268	13.2	.602	14.0	14.1	.112	.00	14.9	21	0600
BE04	3390-9	ATP004 CP	0	4	F0	EM	14.0	.8	.261	12.0	.614	12.8	12.9	.067	.00	14.8	19	0600
BE01	3390-9	ATP001 CP	0	4	F0	EM	13.9	.8	.272	13.4	.626	14.3	14.5	.110	.00	14.8	21	0600
BE0B	3390-9	ATP011 CP	0	4	F0	EM	13.7	.8	.273	13.2	.623	14.1	14.2	.094	.01	14.6	21	0600
BE03	3390-9	ATP003 CP	0	4	F0	EM	13.6	.8	.281	14.5	.653	15.4	15.5	.138	.00	14.5	22	0600
BE11	HPAV	.....	...	4	F0	EM	6.4	4.0	.257	9.96	3.36	13.6	13.6	2.17	.00	10.4	14	0600
BE13	HPAV	.....	...	4	F0	EM	6.4	4.0	.259	9.97	3.53	13.8	13.8	2.30	.00	10.4	14	0600
BE17	HPAV	.....	...	4	F0	EM	6.3	4.1	.259	9.86	3.50	13.6	13.6	2.27	.00	10.4	14	0600
BE18	HPAV	.....	...	4	F0	EM	6.3	4.0	.261	10.1	3.64	14.0	14.0	2.40	.00	10.3	14	0600
BE10	HPAV	.....	...	4	F0	EM	6.3	4.1	.258	10.1	3.42	13.7	13.7	2.18	.00	10.4	14	0600
BE14	HPAV	.....	...	4	F0	EM	6.3	4.2	.256	9.55	3.49	13.3	13.3	2.26	.00	10.5	14	0600
BE15	HPAV	.....	...	4	F0	EM	6.3	4.1	.257	9.66	3.78	13.7	13.7	2.50	.00	10.4	14	0600
BE12	HPAV	.....	...	4	F0	EM	6.3	4.0	.260	10.1	3.70	14.1	14.1	2.49	.00	10.3	14	0600
BE16	HPAV	.....	...	4	F0	EM	6.3	4.0	.257	9.81	3.69	13.8	13.8	2.44	.00	10.3	14	0600
B000	3390-3	PERF1 CP	2	4	F0	EM	.0	.0	.100	.000	.077	.176	.176	.000	.00	.1	0	0323

Command ==>>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 176. Layout of HPF I/O Device Screen (FCX282)

The information shown is based on CP monitor MONITOR domain and I/O domain SAMPLE data.

Figure 176 on page 778 shows a screen with DASDs sorted in ascending order of the real device number, as indicated by the underscore characters above the Addr field. Other possible sort sequences are indicated by dots above the corresponding columns. These sequences can be set using the SORT subcommand with the string shown in the lowest line of the column heading as the sort argument (see “SORT” on page 295) or by positioning the cursor on the heading of the field that is to be sorted and then pressing ENTER.

Some additional lines with device group averages will be shown at the top when the device groups have been defined. (See “FCONTROL DEVGROUP” on page 26.) The device group name will be displayed in place of the device number and device type. The number of devices field (the Label/ID field in the >>A11 DASD<< and device-group lines) contains the number of active devices belonging to the device group on the current screen.

For more detailed information about a specific DASD, you can perform one of the following actions:

- Place the cursor on the Addr field of the device you are interested in and press Enter.
- Enter this subcommand on the command line:

```
DEVICE devno
```

The I/O Device Details screen will then be shown with detailed data for the selected device. See “FCX110, I/O Device Details Screen – DEVICE” on page 420 for more information.

## Field Descriptions:

*Device Descr.:*

### Addr

The real device number.

The first line, which contains the prefix >> A11 DASD << instead of a device number, provides the average load figures for all zHPF-capable DASDs for which performance data were collected.

### Type

The device type and model. Special cases are as follows:

- For PAV alias devices of an IBM DASD subsystem, the device type is replaced by the following string:

```
-> xxxx
```

where xxxx is the device number of the base RDEV.

- For HyperPAV aliases, the device type is replaced by the following string:

```
HPAV
```

### Label/ID

The label or ID, as follows (note that items 2, 3, and 4 apply to DASDs that are not attached to guests):

1. For a device that is attached to a guest, this field contains:

```
>userid
```

where *userid* is the user ID of the guest to which the device is attached.

2. For a HyperPAV alias device, this field contains a label that consists of dots.
3. For a PAV alias device, this field contains the disk label of the base RDEV.
4. For all other DASDs, this field contains the disk label, and, if the device is in the CP-owned list, the label will have a suffix of CP.

### Mdisk Links

For DASDs that are:

- Attached to SYSTEM, this field contains the number of minidisk links.
- Not attached to SYSTEM, this field always contains a 0.

When the device is other than a non-PAV DASD RDEV, a PAV base RDEV, or a HyperPAV base RDEV, dots are displayed in this field.

### Paths

The number of available channel paths to the device. If the number of channel paths is low, and their usage is high, this could cause excessive reconnect misses, which would result in abnormally high disconnect times.

*zHPF:*

### Mask

The mask of paths that support zHPF.

### Featr

The state of the zHPF features on the control unit:

#### D

The zHPF feature has been disabled.

#### E

The zHPF feature has been enabled.

## **M**

Multi-track operations have been enabled.

### **Read/s**

The rate of read operations in zHPF mode.

### **Write/s**

The rate of write operations in zHPF mode.

*Time (msec):*

### **Pend**

The "function pending" time. This is the average amount of time that an I/O request remained pending in the I/O subsystem due to path-busy conditions.

This value and the Disc, Conn, Serv, Busy, and Resp values are calculated based on information obtained from the subchannel measurement block. This information is important especially for disks. Dots are inserted for devices where the information is not available or is not calculated.

### **Disc**

The "device disconnected" time. This is the average amount of time that the device remained disconnected from the channel while executing I/O requests. For DASDs, it includes the time spent on executing the SEEK and SET SECTOR orders and also any re-connect misses. High values could indicate overloaded paths or many long SEEKS to data at opposite ends of the device.

The value also includes the "device-active-only" time where that information is available.

### **Conn**

The "device connected" time. This is the average time that the device was connected to the channel while executing I/O requests, that is, receiving commands or actually transferring data. For DASDs, it also includes the SEARCH time.

### **Serv**

The service time for the device. This is the sum of the function pending, connected, and disconnected times.

### **Resp**

The device response time. This is the device service time plus the time during which an I/O request was waiting to be started.

### **CUWt**

The control unit queuing time, that is, the additional wait time caused when the target device was already busy with an I/O operation that was started from another system.

### **Req. Qued**

The average I/O queue length in the selected interval.

The value is taken from CP monitor high-frequency sample data; it is an indicator for I/O contention on the device.

The I/O queue length (Req. Qued) is not meaningful for devices that are dedicated (attached) to guest systems. The guest system will not try to start another I/O operation before an active one has ended, but rather will do its own queuing. Guest system queuing cannot be shown by host performance monitors, so they will have to be analyzed using a performance monitor on the guest system.

The queue length reported here is the number of real I/O operations (channel programs) that CP's real I/O layer is waiting to start to the device or volume. This includes I/O operations performed for all reasons: for minidisk I/O, for paging I/O, and for spooling I/O, for example. For devices in the CP-owned list, additional queueing information can be found on one of these screens:

- [“FCX109, CP-Owned Device Screen – DEVICE CPOWNED” on page 416](#)
- [“FCX170, CPOWNED Device Log Screen – CPOWNLOG” on page 553.](#)

**I/O**

The number of SSCH and RSCH operations executed per second from this system during the selected interval.

**Pct Busy**

The percentage of time during which the device was busy.

**LCU SSID**

The subsystem ID (SSID) of the logical control unit (LCU).

**Usage notes:**

1. Statistics depicted on the rows of this screen are calculated using only the measurement data for the listed device itself. The effects of PAV or HyperPAV are not included. For a depiction of volume-level statistics that account for the effect of PAV or HyperPAV, use the FCX329 screen (VOLUME) or the FCX330 screen (VOLLOG).
2. On the >>All DASD<< (xxxxx) line, xxxxx indicates the number of devices in the report.

**Effect of AVERAGE and INTERIM subcommands:** Overall or interim period averages will be shown for all load fields. Real averages will be shown for all fields that display device load figures. The Label/ID and Mdisk Links fields will always contain current information.

## FCX283, HPF I/O Device Performance Log Screen – HPFLOG

**Performance Toolkit Subcommand:** “HPFLOG” on page 197.

The HPF I/O Device Performance Log Screen (FCX283) shows a "by time" log of HPF I/O performance data for the selected device. The device data log is available only for devices that you have benchmarked. For each benchmarked device, the interval and line for that device from the FCX282 (INTERIM DEVICE HPF) screen appears on this screen. This allows you to look at a device over time and see how it is performing.

The FCX283 screen looks like this for a HyperPAV base device or a PAV base device:

FCX283		CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS				Perf. Monitor										
HPF I/O Data Log for Device BE00																		
Interval	<- Dev.	Descr.->	Mdisk	Pa-	<----- zHPF ----->			<----- Time (msec)				<----->		Req.	Pct		LCU	
End Time	Type	Label/ID	Links	ths	Mask	Featr	Read/s	Write/s	Pend	Disc	Conn	Serv	Resp	CUWt	Qued	I/O	Busy	SSID
>>Mean>>	3390-9	ATP000 CP	0	4	F0	EM	14.2	.8	.264	13.0	.564	13.9	13.9	.087	.00	15.1	21	0600
14:34:16	3390-9	ATP000 CP	0	4	F0	EM	16.7	.8	.270	13.5	.535	14.3	14.3	.048	.00	17.6	25	0600
14:34:46	3390-9	ATP000 CP	0	4	F0	EM	9.3	.6	.327	17.0	.720	18.0	18.0	.185	.00	10.0	18	0600
14:35:16	3390-9	ATP000 CP	0	4	F0	EM	13.0	.7	.283	17.1	.457	17.9	17.9	.010	.00	13.8	25	0600
14:35:46	3390-9	ATP000 CP	0	4	F0	EM	9.4	.7	.392	29.3	.811	30.5	30.5	.339	.00	10.2	31	0600
14:36:16	3390-9	ATP000 CP	0	4	F0	EM	14.8	.8	.258	12.2	.488	12.9	12.9	.055	.00	15.6	20	0600
14:36:46	3390-9	ATP000 CP	0	4	F0	EM	9.8	.8	.312	16.2	.846	17.4	17.4	.222	.00	10.7	19	0600
14:37:16	3390-9	ATP000 CP	0	4	F0	EM	16.4	.5	.280	14.7	.540	15.5	15.6	.138	.00	17.0	26	0600
14:37:46	3390-9	ATP000 CP	0	4	F0	EM	12.4	.6	.305	14.6	.620	15.5	15.5	.043	.00	13.1	20	0600
14:38:16	3390-9	ATP000 CP	0	4	F0	EM	18.6	1.0	.247	11.9	.588	12.7	12.7	.130	.00	19.7	25	0600
14:38:46	3390-9	ATP000 CP	0	4	F0	EM	13.8	1.0	.258	11.9	.763	12.9	12.9	.203	.00	14.8	19	0600
14:39:16	3390-9	ATP000 CP	0	4	F0	EM	16.7	1.0	.232	9.60	.499	10.3	10.3	.059	.00	17.7	18	0600
14:39:46	3390-9	ATP000 CP	0	4	F0	EM	10.8	1.2	.272	11.7	.755	12.7	12.7	.090	.00	12.0	15	0600
14:40:16	3390-9	ATP000 CP	0	4	F0	EM	19.1	.5	.206	8.54	.345	9.09	9.09	.017	.00	19.7	18	0600
14:40:46	3390-9	ATP000 CP	0	4	F0	EM	12.4	.8	.327	16.9	.695	18.0	18.0	.087	.00	13.3	24	0600
14:41:16	3390-9	ATP000 CP	0	4	F0	EM	15.3	.7	.270	13.1	.580	14.0	14.0	.092	.00	16.1	22	0600
14:41:46	3390-9	ATP000 CP	0	4	F0	EM	9.4	.9	.364	21.1	.687	22.2	22.2	.067	.00	10.4	23	0600
14:42:16	3390-9	ATP000 CP	0	4	F0	EM	15.0	.7	.279	14.1	.607	15.0	15.0	.094	.00	15.7	24	0600
14:42:46	3390-9	ATP000 CP	0	4	F0	EM	13.2	1.0	.307	15.1	.735	16.1	16.1	.050	.00	14.3	23	0600
14:43:16	3390-9	ATP000 CP	0	4	F0	EM	15.5	.9	.260	13.1	.567	13.9	13.9	.069	.00	16.5	23	0600
14:43:46	3390-9	ATP000 CP	0	4	F0	EM	9.6	.8	.309	14.6	.664	15.6	15.6	.038	.00	10.5	16	0600
14:44:16	3390-9	ATP000 CP	0	4	F0	EM	27.9	.4	.098	1.24	.145	1.48	1.48	.005	.00	28.4	4	0600
Command ==>																		
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return													

Figure 177. Layout of HPF I/O Device Performance Log Screen (FCX283)

The information shown is based on CP monitor MONITOR domain and I/O domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the Interval End Time column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

Dev. Descr.:

#### Type

The device type and model. Special cases are as follows:

- For PAV alias devices of an IBM DASD subsystem, the device type is replaced by the following string:

```
-> xxxx
```

where xxxx is the device number of the base RDEV.

- For HyperPAV aliases, the device type is replaced by the following string:

**Label/ID**

The label or ID, as follows (note that items 2, 3, and 4 apply to DASDs that are not attached to guests):

1. For a device that is attached to a guest, this field contains:

```
>userid
```

where *userid* is the user ID of the guest to which the device is attached.

2. For a HyperPAV alias device, this field contains a label that consists of dots.
3. For a PAV alias device, this field contains the disk label of the base RDEV.
4. For all other DASDs, this field contains the disk label, and, if the device is in the CP-owned list, the label will have a suffix of CP.

**Mdisk Links**

For DASDs that are:

- Attached to SYSTEM, this field contains the number of minidisk links.
- Not attached to SYSTEM, this field always contains a 0.

When the device is other than a non-PAV DASD RDEV, a PAV base RDEV, or a HyperPAV base RDEV, dots are displayed in this field.

**Paths**

The number of available channel paths to the device. If the number of channel paths is low, and their usage is high, this could cause excessive reconnect misses, which would result in abnormally high disconnect times.

*zHPF:*

**Mask**

The mask of paths that support zHPF.

**Featr**

The state of the zHPF features on the control unit:

**D**

The zHPF feature has been disabled.

**E**

The zHPF feature has been enabled.

**M**

Multi-track operations have been enabled.

**Read/s**

The rate of read operations in zHPF mode.

**Write/s**

The rate of write operations in zHPF mode.

*Time (msec):*

**Pend**

The "function pending" time. This is the average amount of time that an I/O request remained pending in the I/O subsystem due to path-busy conditions.

This value and the Disc, Conn, Serv, Busy, and Resp values are calculated based on information obtained from the subchannel measurement block. This information is important especially for disks. Dots are inserted for devices where the information is not available or is not calculated.

**Disc**

The "device disconnected" time. This is the average amount of time that the device remained disconnected from the channel while executing I/O requests. For DASDs, it includes the time spent on executing the SEEK and SET SECTOR orders and also any re-connect misses. High values could indicate overloaded paths or many long SEEKS to data at opposite ends of the device.

The value also includes the "device-active-only" time where that information is available.

**Conn**

The "device connected" time. This is the average time that the device was connected to the channel while executing I/O requests, that is, receiving commands or actually transferring data. For DASDs, it also includes the SEARCH time.

**Serv**

The service time for the device. This is the sum of the function pending, connected, and disconnected times.

**Resp**

The device response time. This is the device service time plus the time during which an I/O request was waiting to be started.

**CUWt**

The control unit queuing time, that is, the additional wait time caused when the target device was already busy with an I/O operation that was started from another system.

**Req. Qued**

The average I/O queue length in the selected interval.

The value is taken from CP monitor high-frequency sample data; it is an indicator for I/O contention on the device.

The I/O queue length (Req. Qued) is not meaningful for devices that are dedicated (attached) to guest systems. The guest system will not try to start another I/O operation before an active one has ended, but rather will do its own queuing. Guest system queuing cannot be shown by host performance monitors, so they will have to be analyzed using a performance monitor on the guest system.

The queue length reported here is the number of real I/O operations (channel programs) that CP's real I/O layer is waiting to start to the device or volume. This includes I/O operations performed for all reasons: for minidisk I/O, for paging I/O, and for spooling I/O, for example. For devices in the CP-owned list, additional queueing information can be found on one of these screens:

- [“FCX109, CP-Owned Device Screen – DEVICE CPOWNED” on page 416](#)
- [“FCX170, CPOWNED Device Log Screen – CPOWNLOG” on page 553.](#)

**I/O**

The number of SSCH and RSCH operations executed per second from this system during the selected interval.

**Pct Busy**

The percentage of time during which the device was busy.

**LCU SSID**

The subsystem ID (SSID) of the logical control unit (LCU).

**Usage note:**

Statistics depicted on the rows of this screen are calculated using only the measurement data for the listed device itself. The effects of PAV or HyperPAV are not included. For a depiction of volume-level statistics that account for the effect of PAV or HyperPAV, use the FCX329 screen (VOLUME) or the FCX330 screen (VOLLOG).

**Effect of AVERAGE and INTERIM subcommands:**

None.



# FCX284, I/O Device Data Selection Menu Screen – DEVMENU

**Performance Toolkit Subcommand:** “DEVMENU” on page 171.

The I/O Device Data Selection Menu Screen (FCX284) displays a selection menu of available device reports.

```
FCX284      CPU nnnn  SER nnnnn  I/O data menu  Perf. Monitor

General I/O device performance reports
S Command      Description
- VOLUME              DASD Volume Performance
- DEVICE              I/O device load
- DEVICE  DASD        General DASD load and performance

Specific DASD performance reports
S Command      Description
- DEVICE  CPOWNERD  CPOWNERD DASD load and performance
- DEVICE  HPF       HPF DASD load and performance
- SCSI    EDEV      SCSI EDEV load and performance
- HPALIAS  Alias     HyperPAV Alias Activity
- HPSHARE  Share     HyperPAV Share Activity

Select performance screen with cursor and hit ENTER
Command ==> _
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 178. Layout of I/O Device Data Selection Menu Screen (FCX284)

**Field Descriptions:**

**S**  
The cursor selection field. The field will be highlighted and made cursor-selectable (set as "unprotected") only if the corresponding device performance screen is available.

**Command**  
The command name equal to the screen name that can be selected for display. The field will be highlighted if the corresponding device performance screen is available.

**Description**  
Description of the performance display.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX285, LGR Event Log Screen – LGRELOG

**Performance Toolkit Subcommand:** “LGRELOG” on page 214.

The LGR Event Log screen (FCX285) shows information about the Live Guest Relocation events that have occurred on the system:

FCX285	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor
LGR Event Log Log started at 08:39:19 on 2010/02/23 User event data initially enabled for selected users. See MONSET display for later changes.				
Date	Time	Changed		
02/23	08:40:23	Guest relocation for WASAPP1 from GDLEST3 started by MAINT at 08:40:23.367 with options: FORCE ARCH, ASYNCH, MAXTOTAL=NOLIMIT, MAXQUIESCE= 0 VMRelocation for this user is ALLOWED.		
02/23	08:40:28	Guest relocation for WASAPP1 from GDLEST3 ended successfully at 08:40:28.167, completed in 00:00:04.799.		
02/23	08:40:32	Guest relocation for WASAPP2 from GDLEST3 started by MAINT at 08:40:32.199 with options: FORCE ARCH, ASYNCH, MAXTOTAL=NOLIMIT, MAXQUIESCE= 0 VMRelocation for this user is ALLOWED.		
02/23	08:40:38	Guest relocation for WASAPP2 from GDLEST3 ended successfully at 08:40:38.789, completed in 00:00:06.590. The user has been selected for benchmarking, detailed timings:		
		Time to establish connection	08:40:32.200	1298 usec
		Initial eligibility checks	08:40:32.212	11308 usec
		Create skeleton	08:40:32.214	2659 usec
		Storage management set up	08:40:32.221	6753 usec
		Memory transfer (pre-q. passes)	...	... usec
		FCP I/O quiesce	...	... usec
		TIME to quiesce	08:40:37.689	... usec
		I/O relocation	08:40:37.771	82971 usec
		Final stor.management elig.checks	08:40:37.705	16730 usec
		Penultimate memory pass	...	... usec
		Final VSIM eligibility checks	08:40:37.695	6666 usec
		Machine state relocation	08:40:37.705	10124 usec
		Crypto relocation	...	... usec
		Last memory pass	08:40:38.773	1001k usec
Command ==> F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return				

Figure 179. Layout of LGR Event Log Screen (FCX285)

The information shown is based on CP monitor USER domain EVENT data.

An event message displayed on the screen can contain either general information about the guest relocation or detailed relocation phase timings. To display the detailed Live Guest Relocation phase timings the user identifier of the guest being relocated should be selected for benchmarking using the FC BENCHMRK USER subcommand.

### Field Description:

#### Log started

Time and date when data collection was started.

#### User event data

The initial data collect setting for the USER domain. LGR events can be logged only for periods where USER event data were collected for the affected user(s).

#### Live Guest Relocation data

Shows date, time and nature of any Live Guest Relocation event.

**Date**

Month and day (format *mm/dd*) of LGR event.

**Time**

Time (format *hh:mm:ss*) of the LGR event.

**Changed**

Live Guest Relocation status changes as follows:

- Guest Relocation for *userid* from *srcnode* started by *issuer* at *start\_time* with options:

```
relocation_options
VMRelocation for this user is ALLOWED.
```

- Guest Relocation for *userid* to *destnode* started by *issuer* at *start\_time* with options:

```
relocation_options
VMRelocation for this user is ALLOWED.
```

- Guest Relocation for *userid* to *destnode* ended successfully at *end\_time*, completed in *interval* seconds.
- Guest Relocation for *userid* from *destnode* ended successfully at *end\_time*, completed in *interval* seconds.
- Guest Relocation for *userid* from *srcnode* ended with *error* (*nn*, *error\_message*) at *end\_time*, completed in *interval* seconds.
- Guest Relocation for *userid* to *destnode* ended successfully at *start\_time*, completed in *interval* seconds. The user has been selected for benchmarking, detailed timings:

Time to establish connection	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
Initial eligibility checks	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
Create skeleton	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
Storage management set up	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
Memory transfer (pre-q. Passes)	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
FCP I/O quiesce	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
Time to quiesce	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
I/O relocation	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
Final stor.management elig.checks	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
Penultimate memory pass	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
Final VSIM eligibility checks	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
Machine state relocation	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
Crypto relocation	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
Last memory pass	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
Final I/O eligibility checks	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
Resume guest	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
Clean-up	<i>hh:mm:ss.sss</i>	<i>vvv usec</i>
Total quiesce time	<i>HH:MM:SS.SSS</i>	
Total relocation time	<i>HH:MM:SS.SSS</i>	

Where:

***userid***

User identifier of the virtual machine being relocated.

***srcnode***

The system identifier from which the virtual machine relocation was started.

***destnode***

The system identifier to which the virtual machine is being relocated or was relocated to.

***start\_time***

Time when the relocation was started.

***relocation\_options***

Options specified for the virtual machine relocation.

***VMRelocation***

Relocation mode value specified for the virtual machine.

**issuer**

The user identifier from which the relocation was started.

**end\_time**

Time when the virtual machine relocation was ended.

**interval**

The time interval length in which the relocation was completed.

**nn**

CP error return code.

**error\_message**

Error message text:

```
01 Canceled by VMRELOCATE CANCEL command
02 Canceled by CPHX command
03 Canceled due to lost ISFC connection
04 Canceled due to MAXTOTAL time limit exceeded
05 Canceled due to MAXQUIESCE time limit exceeded
06 Canceled due to eligibility violation
07 Canceled due to virtual machine reset or logoff
08 Canceled due to internal processing error
09 Canceled because CP exit rejected this relocation
11 Canceled because CP exit gave a return code that is not valid
```

**hh:mm:ss.sss**

Time when a relocation phase was ended for a user identifier which was selected for benchmarking.

**vvv**

The number of microseconds in which a relocation phase was completed.

**HH:MM:SS.SSS**

Time interval spent for relocation or quiesce.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX286, LGR Data Screen – LGRDATA

**Performance Toolkit Subcommand:** LGRDATA on page 213.

The LGR Data screen (FCX286) shows information about the Live Guest Relocations that have occurred on the system:

FCX286		CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS		Perf.		Monitor																	
-----		.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	

Figure 180. Layout of LGR Data Screen (FCX286)

The information shown is based on CP monitor USER domain EVENT data.

The sample shows a screen with information about Live Guest Relocation sorted in ascending Time of LGR start/end sequence (indicated by the underscore characters above the EndTime heading), the initial sequence. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see [“SORT” on page 295](#)) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

A Live Guest Relocation process is described on the LGR Data screen by a row containing the data from the LGR start event and LGR end event MONITOR records. The Performance Toolkit displays a row for each started relocation and then updates the row fields with the data from the relocation end event message. The maximum number of the finished LGR events kept on this report is limited by the REDISP value set by the FC MONCOLL REDISP subcommand. The oldest events are permanently removed.

### Field Description:

**Date**

Date of LGR start or LGR end.

**EndTime**

Time of LGR start or LGR end.

**Issuer**

VMRELOCATE command issuer.

**Userid**

User identifier of relocation target.

### Source

Name of the source system.

**Dest**

Name of the destination system.

**FC**

Code indicating reason why the relocation ended:

00

Completed successfully.

01

Canceled by VMRELOCATE CANCEL command.

02

Canceled by CPHX command.

- 03**  
Canceled due to lost ISFC connection.
- 04**  
Canceled due to MAXTOTAL time limit exceeded.
- 05**  
Canceled due to MAXQUIESCE time limit exceeded.
- 06**  
Canceled due to eligibility violation.
- 07**  
Canceled due to virtual machine reset or logoff.
- 08**  
Canceled due to internal processing error.
- 09**  
Canceled because CP exit rejected this relocation.
- 11**  
Canceled because CP exit gave a return code that is not valid.
- \*\***  
Relocation has been started and still is in progress.

*Time, sec:*

**Reloc**  
Total relocation time.

**Quisc**  
Total quiesce time.

**Pass**  
Number of memory passes made during the relocation.

*Pages moved by Pass...:*

**First**  
Number of guest pages transferred in the first pass.

**Inter**  
Average number of guest pages transferred during pass 2 through pass n-2.

**N-1**  
Number of guest pages transferred in the next to the last (penultimate) pass.

**N**  
Number of guest pages transferred in the final pass.

**VDEV**  
Count of transferred virtual devices.

**Actv I/O**  
Count of active I/O operations encountered.

*Cleared I/Os:*

**nonQ**  
Count of non-QDIO type I/O operations cleared.

**QDIO**  
Count of QDIO type I/O operations cleared.

*Reloc Options:*

**SYN**  
SYNCHRONOUS option was specified in the VMRELOCATE command (Y=yes, N=no).

**IMM**  
IMMEDIATE option was specified in the VMRELOCATE command (Y=yes, N=no).

**MQ**

MAXQUIESCE NOLIMIT option was specified in the VMRELOCATE command (Y=yes, N=no).

**MR**

MAXTOTAL NOLIMIT option was specified in the VMRELOCATE command (Y=yes, N=no).

**FA**

FORCE ARCHITECTURE option was specified in the VMRELOCATE command (Y=yes, N=no).

**FD**

FORCE DOMAIN option was specified in the VMRELOCATE command (Y=yes, N=no).

**FS**

FORCE STORAGE option was specified in the VMRELOCATE command (Y=yes, N=no).

*Set Opts:*

**OFF**

OFF was specified in the SET VMRELOCATE command (Y=yes, N=no).

**CAF**

CPU affinity suppression (Y=yes, N=no).

**AFC**

CPU affinity was changed (Y=yes, N=no).

**SRC**

The MONITOR record is the source system record (Y=yes, N=no).

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX287, System Topology Machine Organization Screen – TOPOLOG

**Performance Toolkit Subcommand:** “TOPOLOG” on page 330.

The System Topology Machine Organization screen (FCX287) displays LPAR topology information.

```
FCX287      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor

Topology configuration for GDLFCFT

Topology-list geometry (0,0,0,4,6,6) MNest= 3
Multithreading         Facility is not installed

Date Time      Location  Type  Count  Core List  PP/Entitl. Ded.
04/11 14:53:55 Node      3
04/11 14:53:55 Chip      1  ZIIP      2  0C-0D      V-medium    ...
04/11 14:53:55 Node      4
04/11 14:53:55 Chip      1  CP        2  05-06      V-medium    ...
04/11 14:53:55 Chip      2  CP        5  00-04      V-high      No
04/11 14:53:55 Chip      5  CP        2  07-08      V-low       ...
04/11 14:53:55          IFL        1  0E         V-medium    ...
04/11 14:53:55          IFL        1  0F         V-low       ...
04/11 14:53:55 Chip      6  CP        3  09-0B      V-low       ...
04/11 14:53:55          IFL        2  10-11      V-low       ...

Command ===>
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 181. Layout of System Topology Machine Organization Screen (FCX287)

The information shown is based on CP monitor USER domain EVENT data.

### Field Descriptions:

This screen is available when running Performance Toolkit against z/VM 6.2 or later data. These fields are also described in *z/Architecture Principles of Operations*, SA22-7832, under "SYSIB 15.1.2-15.1.6 (Configuration Topology)".

### Topology-list geometry (*mag6,mag5,mag4,mag3,mag2,mag1*) MNest= *m*

The maximum number of container-type topology entries (*mag2–mag6*, core-type entries *mag1*, and maximum nesting level of the topology (*m*).

**Note:** When running Performance Toolkit against data from z/VM 6.3 with CP APAR VM65586 installed, this screen shows the real hardware topology for machines, with up to three nesting levels. For machines with four or more nesting levels, this screen shows the topology layout recalculated into a three-level layout.

### Multithreading

The multithreading mode. Possible values are:

- Enabled
- Facility is not installed
- No CP support
- Disabled, as requested
- Disabled, no MULTITHREADING statement
- Disabled, ESA/390 LPAR
- Disabled, SIGP failure
- Disabled, IPL CPU address is too high
- Disabled, horizontal polarization
- Disabled, corequisite facilities *nn* were not available



- Disabled, LINUX only LPAR with CP processors
- Disabled, SET SRM REBALANCE is active

For more information, see the **Multithreading** field description in [“FCX180, System Configuration – SYSCONF” on page 577.](#)

#### *Geometry Log Details:*

##### **Date**

The date of the status change (in *mm/dd* format).

##### **Time**

The time of the status change (in *hh:mm:ss* format).

##### **Location**

The actual physical location of listed logical cores.

##### **Type**

The core type. Possible values are:

##### **CP**

Central Processor

##### **ICF**

Internal Coupling Facility

##### **IFL**

Integrated Facility for Linux

##### **ZIIP**

IBM Z Integrated Information Processor

##### **Count**

The number of listed logical cores.

##### **Core List**

The list of logical core IDs (in hexadecimal format) of the topology-list entry.

##### **PP/Entitl.**

The polarization (PP) and entitlement for the present topology-list entry. Possible values are:

For vertical polarization:

V-high

V-medium

V-low

For horizontal polarization:

Horizontal

##### **Ded.**

The dedication for the present topology-list entry. The field has a sense only for cores with horizontal polarization or vertical polarization with high entitlement. Otherwise, this field contains dots.

Yes – listed cores are dedicated physical cores.

No – all cores are shared physical cores.

##### **Usage note:**

The term *core* describes a unit that contains one or more processors. All processors of a core have the same processor type. For machine types that do not have the multithreading (MT) facility, the machine has physical processors and the partitions have logical processors. For machine types that have the MT facility, the machine has physical cores, each of which contains two physical processors. A partition that has not activated the MT facility continues to have logical processors. A partition that has activated the MT facility has logical cores, each of which contains one or more logical processors.

**Effect of AVERAGE and INTERIM subcommands:** None.

## FCX288, Multiprocessor User Activity Log Screen – USRMPLOG

**Performance Toolkit Subcommand:** “USRMPLOG” on page 348.

The Multiprocessor User Activity Log Screen (FCX288) shows a 'by time' log of the selected user's virtual CPU consumption. Note that the specified user must have been previously selected for benchmarking.

FCX288	CPU nnnn		SER nnnnn	Interval HH:MM:SS - HH:MM:SS						Perf. Monitor										
Virtual Processors Usage Log for User VSB00001																				
Interval	<vCPU->			<-- All Processors -->					<-Primary Processors ->				<-Secondary Processors>							
End Time	ID	Type	Aff	TTime	TCPU%	VTTime	VCPU%	STime	SCPU%	VTTime	VCPU%	STime	SCPU%	VTTime	VCPU%	IO/s	PgF/s	Share	L	MxShar
>>Mean>>	00	CP	---	1.860	31.00	1.327	22.11	.533	8.889	1.327	22.11	.000	.000	.000	.000	.000	.000	---	-	---
>>Mean>>	01	CP	---	1.860	31.00	1.332	22.21	.528	8.794	1.332	22.21	.000	.000	.000	.000	.000	.000	---	-	---
>>Mean>>	02	CP	---	1.860	31.00	1.326	22.11	.534	8.893	1.326	22.11	.000	.000	.000	.000	.000	.000	---	-	---
>>Mean>>	0F	CP	---	1.860	31.00	1.315	21.92	.545	9.083	1.315	21.92	.000	.000	.000	.000	.000	.000	---	-	---
10:57:32	00	CP	ON	1.860	31.00	1.327	22.11	.533	8.889	1.327	22.11	.000	.000	.000	.000	.000	.000	25	.	...
10:57:32	01	CP	ON	1.860	31.00	1.332	22.21	.528	8.794	1.332	22.21	.000	.000	.000	.000	.000	.000	25	.	...
10:57:32	02	CP	ON	1.860	31.00	1.326	22.11	.534	8.893	1.326	22.11	.000	.000	.000	.000	.000	.000	25	.	...
10:57:32	0F	CP	ON	1.860	31.00	1.315	21.92	.545	9.083	1.315	21.92	.000	.000	.000	.000	.000	.000	25	.	...
Command ==> _																				
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return																				

Figure 182. Layout of Multiprocessor User Activity Log Screen (FCX288)

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval for which the values were calculated. In the case of the ">>Mean>>" entry, you can find the relevant interval in the header line.

#### vCPU ID

Identifies the virtual processor address, in hexadecimal.

#### vCPU Type

Indicates the type of virtual processor:

##### CP

Central Processor

##### ICF

Internal Coupling Facility Processor

##### IFL

Integrated Facility for Linux Processor

##### ZIIP

IBM Z Integrated Information Processor

##### MIX

The MP user has different PU types.

This field contains a value at the end of the interval.

#### Aff

Indicates the affinity and/or status of the virtual processor. The field can be as follows:

##### ON

Indicates CPU affinity has been set on for this user with CPUAFFINITY command and this virtual CPU is dispatched on of the same type.

##### OFF

Indicates CPU affinity has not been set on for this user. When OFF is in effect the virtual CPU is dispatched on only real CPUs of primary type.

##### SUP

Indicates CPU affinity is suppressed. Although CPU affinity set on for this user with the CP SET CPUAFFINITY command, the currently no matching real processor on which to run.

##### DED

Shows that this virtual processor is dedicated

**DEDS**

Shows that this virtual processor is dedicated but in stopped state

**STOP**

The virtual processor is in stopped state.

The field contains a value at the end of the interval in the main log line(s). It always contains dashes in the Mean line(s), because no average data can be calculated for this field.

**All Processors:****TTime**

Total CPU utilization for all logical processors by current virtual processor as sum of virtual and simulation times, in seconds.

**TCPU%**

Total CPU utilization for all logical processors by current virtual processor as sum of virtual and simulation times, in percentages.

**VTime**

Total virtual time for user on all logical processors, in seconds.

**VCPU%**

Total virtual time for user on all logical processors, in percentages.

All fields might contain dots if the timer for Total CPU usage has not been changed during measured interval.

**Primary Processors:****STime**

Total simulation time for user on primary processors, in seconds.

**SCPU%**

Total simulation time for user on primary processors, in percentages.

**VTime**

Total virtual time for user on primary processors, in seconds.

**VCPU%**

Total virtual time for user on primary processors, in percentages.

All fields might contain dots if the timer for Total CPU usage has not been changed during measured interval.

**Secondary Processors:****STime**

Total simulation time for user on secondary processors, in seconds.

**SCPU%**

Total simulation time for user on secondary processors, in percentages.

**VTime**

Total virtual time for user on secondary processors, in seconds.

**VCPU%**

Total virtual time for user on secondary processors, in percentages.

All fields might contain dots if the timer for Total CPU usage has not been changed during measured interval.

**IO/s**

Virtual IO operations per second for the processor.

The field might contain dots if the timer for Total CPU usage has not been changed during measured interval.

**PgF/s**

Tells how many page faults occurred for the processor, per second.

## Share

Appears for users with a relative or absolute share. User's share, either in the format *nnnnn* for relative share, or as *nnnnn%* for an absolute share percentage, which is used by CP for this virtual processor, not for the entire user.

Dots will be displayed if the current processor is in stopped state or dedicated.

The field contains a value at the end of the interval in the main log line(s). It always contains dashes in the Mean line(s), because no average data can be calculated for this field.

## L

Specifies Maximum Share option value if set. Possible values are:

### S

The option LIMITSOFT used for Maximum share value

### H

The option LIMITHARD used for Maximum share value

.

Maximum share was not set, the option NOLIMIT used for that virtual processor.

The field contains a value at the end of the interval in the main log line(s). It always contains a dash in the Mean line(s), because no average data can be calculated for this field.

## MxShar

Specifies the value of Maximum share for that virtual processor. User's max. share, if available, either in the format *nnnnn* for a relative share, or as *nnn%* for an absolute maximum share percentage.

Dots will be displayed if the current processor is in stopped state or the maximum share is not set (set as NOLIMIT).

The field contains a value at the end of the interval in the main log line(s). It always contains dashes in the Mean line(s), because no average data can be calculated for this field.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX289, User Paging Menu Screen – UPGMENU

**Performance Toolkit Subcommand:** “UPGMENU” on page 339.

In the FCX289, User Paging Menu screen (FCX289), shows a list of all the available user paging data displays.

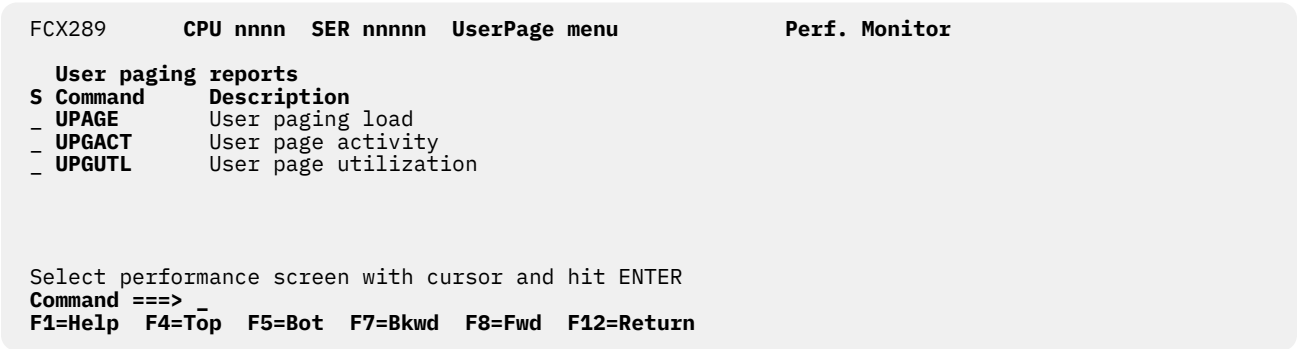


Figure 183. Layout of User Paging Menu Screen (FCX289)

The information on this screen is based on CP monitor USER domain SAMPLE data.

**Field Descriptions:**

**S**

Cursor selection field. The field will be highlighted and made cursor-selectable (set as unprotected) only if the corresponding user page performance screen is available.

**Command**

The screen name, equal to the command that can be selected for display. The field will be highlighted if the corresponding user page performance screen is available.

**Description**

Description of the performance display.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX290, User Page Activity Screen – UPGACT

**Performance Toolkit Subcommand:** “UPGACT” on page 337.

The User Page Activity screen (FCX290) displays detailed information on z/VM's management of each virtual machine's memory management activities during the last measuring interval.

FCX290	CPU	nnnn	SER	nnnnn	Interval	HH:MM:SS	-	HH:MM:SS	Perf.	Monitor					
-----	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
	<----- Storage ----->														
	Stl	<----	Transition/s	----	<-Steal/s->	<----- Movement/s ----->					<Migrate/s>				
-----	Wt	Inst	Relse	Inval	Reval	Ready	NoRdy	PGIN	PGOUT	Reads	Write	MWrit	Xrel	Nr of	Users
>>Mean>>	1.0	901.6	979.1	528K	461K	192K	.0	.0	.0	204K	139K	.0	.0	56	
User Class Data:															
PERF	1.0	983.0	969.4	97662	95874	1516	.0	.0	.0	52511	50927	.0	.0	1	
Service	1.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	1	
User Data:															
APCSCIF	1	0	0	0	0	0	0	0	0	0	0	0	0	0	
BK200001	1	0	0	249K	249K	0	0	0	0	0	0	0	0	0	
BK200007	1	0	0	239K	239K	0	0	0	0	0	0	0	0	0	
BK200008	1	54.6	54.6	244K	244K	0	0	0	0	0	0	0	0	0	
BK300012	1	0	0	309K	1452K	2805K	0	0	0	5285K	1413K	0	0	0	
BK400015	1	0	0	3678K	0	2579K	0	0	0	923K	2799K	0	0	0	
CFT2NDA	1	0	0	35294	34488	532.5	0	0	0	355.0	846.5	0	0	0	
CHG00001	1	9571	9161	54955	53084	1065	0	0	0	0	805.5	0	0	0	
CHG00006	1	6171	6062	58627	56498	13.7	0	0	0	27.3	4410	0	0	0	
CHG00008	1	41.0	41.0	3973	3932	218.5	0	0	0	0	122.9	0	0	0	
Command ==>															
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return										

Figure 184. Layout of User Page Activity Screen (FCX290)

The information shown is based on CP monitor USER domain SAMPLE data.

The sample shows a screen with memory management information for each virtual machine sorted in ascending userid sequence (indicated by the underscore characters above the Userid heading), the initial sequence. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

**Field Descriptions:** The following field descriptions apply to both the UPGACT and UPGACTLG *userid* displays, with the exception of the "Userid" field which will be replaced by the "Interval End Time" field on the FCX291 display.

### Userid

User identification.

### Interval End Time

The end time of the monitor interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

### Stl Wt

Steal weight factor. Used to determine the level of aggressiveness in reclaiming frames from this user.

### Storage Transition/s:

#### Inst

Rate of private storage instantiation, in bytes per second.

#### Relse

Rate of private released storage, in bytes per second.

#### Inval

Rate of private storage invalidation, in bytes per second.

**Reval**

Rate of private IBR storage revalidation, in bytes per second.

*Storage Steal/s:*

**Ready**

Rate of private IBR storage reclaimed, in bytes per second. At reclaim time, the pages' latest content was already preserved in auxiliary storage.

**NoRdy**

Rate of private IBR storage reclaimed but only after the contents are preserved in XSTORE or auxiliary storage at reclaim time, in bytes per second.

*Storage Movement/s:*

**PGIN**

Rate of private storage reads from XSTORE, in bytes per second.

**PGOUT**

Rate of private storage writes to XSTORE, in bytes per second.

**Reads**

Rate of private storage reads from auxiliary storage, in bytes per second.

**Write**

Rate of private storage writes to auxiliary storage, in bytes per second.

*Storage Migrate/s:*

**MWrit**

Rate of private storage writes by the migration task from XSTORE to auxiliary storage, in bytes per second.

**Xrel**

Rate of private storage releases from XSTORE blocks by the migration task, without having to back up the contents in auxiliary storage, because the latest content is already in auxiliary storage, in bytes per second.

**Nr of Users**

The number of users which contributed to the average shown on this line (values are inserted only for lines with averages; the field remains empty for normal user data lines).

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all data fields. All of a user's sessions during the selected period contribute to the averages shown. Note also that all of the activity fields (rates) are based on each user's connect time within the overall period, not the total elapsed time for the whole period.

# FCX291, User Page Activity Log Screen – UPGACTLG

**Performance Toolkit Subcommand:** “UPGACTLG” on page 338.

The User Page Activity Log screen (FCX291) displays a log of user page activity.

FCX291	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor
Page Activity Log for User TH200006				
<----- Storage ----->				
<----- Movement/s ----->				
Interval	Stl	<--- Transition/s --->		<-Steal/s->
End Time	Wt	Inst	Relse	Inval
>>Mean>>	1.0	0	0	4536K
06:09:19	1	0	0	109K
06:14:19	1	0	0	3678K
06:19:19	1	0	0	88446
06:24:19	1	0	0	44128
Command ==> _				
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return				

Figure 185. Layout of User Page Activity Log Screen (FCX291)

The information shown is based on CP monitor USER domain SAMPLE data.

**Field Descriptions:** For a description of the fields on this screen, see “FCX290, User Page Activity Screen – UPGACT” on page 798.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.



## FCX292, User Page Utilization Data Screen – UPGUTL

### Performance Toolkit Subcommand: “UPGUTL” on page 340

The User Page Utilization Data screen (FCX292) displays detailed information on each virtual machine's utilization of z/VM paging resources during the last measuring interval.

FCX292	CPU	nnnn	SER	nnnnn	Interval	HH:MM:SS	-	HH:MM:SS	Perf.	Monitor										
											Storage									
											Resident									
											Invalid But Resident									
											Base									
Userid	Data Spaces Owned	WSS	Inst	Resvd	T_All	T<2G	T>2G	L<2G	L>2G	U<2G	U>2G	P<2G	P>2G	A<2G	A>2G	XSTOR	AUX	Space Size	Nr of Users	
>>Mean>>	10.6	5645M	6113M	256M	5391M	34M	5356M	585.1	41472	.0	2743K	292.6	46M	73.1	6008K	.0	3305M	6341M	56	
User Class Data:																				
PERF	.0	16M	35M	.0	16M	.0	16M	.0	.0	.0	512K	.0	15M	.0	24576	.0	34M	3072M	1	
Service	.0	1972K	5196K	2048K	2048K	4096	2044K	.0	77824	.0	1204K	.0	.0	.0	376K	.0	3640K	32768K	1	
User Data:																				
APCSCIF	0	952K	972K	0	4K	0	4K	0	0	0	0	0	0	0	0	0	968K	64M		
BK200001	0	3318M	3074M	0	3072M	0	3072M	0	0	0	0	0	0	0	0	0	2160K	3328M		
BK200007	0	3318M	3074M	0	3072M	12M	3060M	0	0	0	0	0	0	0	0	0	3074M	3328M		
BK200008	0	3318M	3074M	0	3072M	76K	3072M	0	0	0	0	0	0	0	0	0	3074M	3328M		
BK300012	0	4194M	16G	0	6814M	164K	6814M	0	0	0	8728K	4K	265M	4K	273M	0	16G	16640M		
BK400015	0	100G	100G	0	86G	760M	85G	0	0	0	47M	0	1085M	0	50M	0	60G	100G		
Command ==> _																				
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return																				

Figure 186. Layout of User Page Utilization Data Screen (FCX292)

The information shown is based on CP monitor USER domain SAMPLE data.

The sample shows a screen with utilization of z/VM paging resources for each virtual machine sorted in ascending userid sequence (indicated by the underscore characters above the Userid heading), the initial sequence. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

**Field Descriptions:** The following field descriptions apply to both the UPGUTL and UPGUTLLG *userid* displays, with the exception of the "Userid" field which will be replaced by the "Interval End Time" field on the FCX293 display.

#### Userid

User identification.

#### Interval End Time

The end time of the monitor interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### Data Spaces Owned

Number of non-base spaces owned by this user.

Storage:

#### WSS

User's projected working-set size, in bytes.

#### Inst

Instantiated private storage, in bytes.

#### Resvd

Reserved storage setting, in bytes.

Storage Resident Total:

#### T\_All

Total resident private storage, in bytes. Individual pages may be valid or in the IBR state.

#### T<2G

Resident private storage below 2G, in bytes.

**T>2G**

Resident private storage above 2G, in bytes.

*Storage Resident Locked:*

**L<2G**

Locked private storage below 2G, in bytes.

**L>2G**

Locked private storage above 2G, in bytes.

*Storage Resident Invalid But Resident UFO:*

**U<2G**

Invalid but resident private storage on the User Frame Owned List that is backed by frames below 2G, in bytes.

**U>2G**

Invalid but resident private storage on the User Frame Owned List that is backed by frames above 2G, in bytes.

*Storage Resident Invalid But Resident PNR:*

**P<2G**

Invalid but resident private storage that has not been referenced since it was read-in by a block read and is backed by frames below 2G, in bytes.

**P>2G**

Invalid but resident private storage that has not been referenced since it was read-in by a block read and is backed by frames above 2G, in bytes.

*Storage Resident Invalid But Resident Age List:*

**A<2G**

Invalid but resident private storage on the Age List that is backed by frames below 2G, in bytes.

**A>2G**

Invalid but resident private storage on the Age List that is backed by frames above 2G, in bytes.

*Storage:*

**XSTOR**

Amount of XSTORE occupied by pages belonging to this user, in bytes.

**AUX**

Amount of auxiliary paging storage occupied by pages belonging to this user, in bytes. Minidisk pages mapped in an XC virtual machine's address space using the MAPMDISK interface are *not* included in this count.

**Base Space Size**

Size of the user's base address space, in bytes.

**Nr of Users**

The number of users which contributed to the average shown on this line (values are inserted only for lines with averages; the field remains empty for normal user data lines).

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all data fields. All of a user's sessions during the selected period contribute to the averages shown. Note also that all of the activity fields (rates) are based on each user's connect time within the overall period, not the total elapsed time for the whole period.

# FCX293, User Page Utilization Data Log Screen – UPGUTLLG

**Performance Toolkit Subcommand:** “UPGUTLLG” on page 341.

The User Page Utilization Data Log screen (FCX293) displays a log of user page utilization data.

FCX293	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS				Perf. Monitor													
Page Utilization Log for User TH200006																				
<----- Storage ----->																				
<----- Resident ----->																				
<----- Invalid But Resident ----->																				
Interval	Data																	Base		
End Time	Spaces	Owned	WSS	Inst	Resvd	T_All	Total	T<2G	T>2G	L<2G	L>2G	U<2G	U>2G	P<2G	P>2G	A<2G	A>2G	XSTOR	AUX	Space
>>Mean>>	0	3072M	3072M	2048M	3072M	20K	3072M	0	0	0	0	0	3168K	0	8K	0	0	0	1762M	3072M
06:09:19	0	3072M	3072M	2048M	3072M	20K	3072M	0	0	0	0	0	0	0	0	0	0	0	1762M	3072M
06:14:19	0	3072M	3072M	2048M	3072M	20K	3072M	0	0	0	0	0	512K	0	128K	0	0	0	1762M	3072M
06:19:19	0	3072M	3072M	2048M	3072M	20K	3072M	0	0	0	0	0	0	0	0	0	0	0	1762M	3072M
06:24:19	0	3072M	3072M	2048M	3072M	20K	3072M	0	0	0	0	0	0	0	0	0	0	0	1762M	3072M
Command ==> _																				
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return															

Figure 187. Layout of User Page Utilization Data Log Screen (FCX293)

The information shown is based on CP monitor USER domain SAMPLE data.

**Field Descriptions:** For a description of the fields on this screen, see “FCX292, User Page Utilization Data Screen – UPGUTL” on page 801.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX294, Available List Data Below 2G Screen – AVLB2GLG

**Performance Toolkit Subcommand:** “AVLB2GLG” on page 135.

The Available List Data Below 2G screen (FCX294) displays 'by time' information on the real memory below 2G Available List.

FCX294	CPU nnnn	SER nnnnn	Interval	HH:MM:SS	- HH:MM:SS	Perf. Monitor
<div> <div> <div>Interval</div> <div>End Time</div> </div> <div> <div>&lt;Available&gt;</div> <div>Sing</div> <div>Cont</div> </div> <div> <div>&lt;Requests/s&gt;</div> <div>Sing</div> <div>Cont</div> </div> <div> <div>&lt;Returns/s&gt;</div> <div>Sing</div> <div>Cont</div> </div> <div> <div>&lt;--Times--&gt;</div> <div>&lt;Empty/s--&gt;</div> <div>Sing</div> <div>Cont</div> </div> <div> <div>&lt;Frame Thresh&gt;</div> <div>&lt;- Singles --&gt;</div> <div>Low</div> <div>Prot</div> </div> </div>						
>>Mean>>	8K	149K	350.0	343.5	271.6	322.6
04:39:19	0	88K	204.8	.0	.0	.0
04:44:19	4K	148K	2772	4096	2895	3919
04:49:19	4K	196K	.0	.0	81.9	81.9
04:54:19	88K	156K	136.5	122.9	341.3	68.3
04:59:19	24K	224K	368.6	68.3	245.8	191.1
05:04:19	0	276K	382.3	273.1	300.4	368.6
05:09:19	0	120K	245.8	314.0	.0	.0
05:14:19	0	108K	614.4	163.8	81.9	150.2
05:19:19	0	120K	273.1	191.1	204.8	218.5
05:24:19	8K	228K	245.8	109.2	300.4	273.1
05:29:19	4K	244K	.0	.0	13.7	27.3
05:34:19	4K	252K	.0	.0	.0	27.3
05:39:19	0	120K	245.8	286.7	.0	41.0
<div> <div>Command ==&gt;</div> <div>_</div> </div> <div> <div>F1=Help</div> <div>F4=Top</div> <div>F5=Bot</div> <div>F7=Bkwd</div> <div>F8=Fwd</div> <div>F12=Return</div> </div>						

Figure 188. Layout of Available List Data Below 2G Screen (FCX294)

The information shown is based on CP monitor SYSTEM and STORAGE domains SAMPLE data.

**Field Descriptions:** The following field descriptions apply to both the AVLB2GLG and AVLA2GLG displays, except where the references to '<2G' in the fields for AVLB2GLG below should be read as for '>2G' in the fields for AVLA2GLG.

### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

*Storage:*

### Available Sing

Amount of storage on the singles available list, in bytes.

### Available Cont

Amount of storage on the contig available list, in bytes. The frames comprising the list are in groups of two or more contiguous frames.

### Requests/s Sing

Amount of storage obtained from the singles list, in bytes per second.

### Requests/s Cont

Amount of storage obtained from the contig list, in bytes per second.

### Returns/s Sing

Amount of storage returned to the singles list, in bytes per second.

### Returns/s Cont

Amount of storage returned to the contig list, in bytes per second.

*Times:*

### Empty/s Sing

Count of times a <2G frame request could not be satisfied from the <2G singles available list because it was empty, per second.

**Empty/s Cont**

Count of times a <2G frame request could not be satisfied from the <2G contig available list because it was empty, per second.

*Frame Thresh:*

**Singles Low**

Low threshold for singles available list, in frames.

**Singles Prot**

Protect threshold for < 2G singles available list, in frames.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX295, Available List Data Above 2G Screen – AVLA2GLG

**Performance Toolkit Subcommand:** “AVLA2GLG” on page 134.

The Available List Data Above 2G screen (FCX295) displays 'by time' information on the real memory above 2G Available List.

FCX295	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS				Perf. Monitor				
Interval End Time	<----- Storage ----->				<---Times-->		<-Frame Thresh-->				
	<Available> Sing	Cont	<Requests/s> Sing	Cont	<Returns/s> Sing	Cont	<---Empty--> Sing	Cont	Sing	<-Contigs-> Low	Prot
>>Mean>>	52M	203M	6333K	5532K	6466K	5355K	.0	.0	16	15	15
04:39:19	53M	240M	14M	10M	15M	9852K	.0	.0	16	15	15
04:44:19	231M	241M	23M	20M	24M	20M	.0	.0	17	15	15
04:49:19	0	3308K	7058K	5570K	6416K	4613K	.0	.0	21	15	15
04:54:19	114M	515M	3296K	2119K	3878K	3673K	.0	.0	19	15	15
04:59:19	495M	836M	3242K	4261K	4844K	5056K	.0	.0	15	15	15
05:04:19	7416K	244M	5155K	4255K	3345K	2380K	.0	.0	15	15	15
05:09:19	0	24M	6297K	5585K	6338K	4766K	.0	.0	15	15	15
05:14:19	0	239M	9749K	7655K	9993K	8148K	.0	.0	15	15	15
05:19:19	0	44M	7429K	7280K	7551K	6490K	.0	.0	17	15	15
05:24:19	0	205M	5807K	4930K	5963K	5325K	.0	.0	15	15	15
05:29:19	0	211M	2993K	3067K	3070K	3007K	.0	.0	15	15	15
05:34:19	664K	247M	2289K	2220K	2396K	2237K	.0	.0	15	15	15
05:39:19	0	242M	2573K	2186K	2615K	2126K	.0	.0	15	15	15

Command ==>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 189. Layout of Available List Data Above 2G Screen (FCX295)

The information shown is based on CP monitor SYSTEM domain and STORAGE domain SAMPLE data.

**Field Descriptions:** For a description of the fields on this screen, see “FCX294, Available List Data Below 2G Screen – AVLB2GLG ” on page 804.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX296, Steal Statistics Screen – STEALLOG

**Performance Toolkit Subcommand:** “STEALLOG” on page 301.

The Steal Statistics screen (FCX296) displays 'by time' information on z/VM's memory management stealing statistics.

FCX296	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS										Perf. Monitor															
Interval	Pct	Storage/s																		<-Completions/s>					<-- Age List --->			
Time	Total	Write	<-User-->			<-Shared-->			<Pvt Vdisk>			AgeL	Need	Time	Sys	Travs	<-Users/s>		<--Stor		Skip/s-->							
End Time	Actv	Stoln	OnDmd	Inval	Reval	Inval	Reval	Inval	Reval	Inval	Reval	Reval	Met	Limit	Req	/s	Visit	Skip	Pin	Ser	Resv							
>>Mean>>	4.2	11M	.0	67M	42M	.0	.0	15859	6713	22M	.6	1.9	.0	1.9	.0	.4	9.5	1.9	.0	.0	.0							
04:39:19	16.2	23M	.0	416M	269M	.0	.0	23866	8219	144M	.7	10.5	.0	1.6	34.3	7.4	.0	.0	.0	.0								
04:44:19	13.4	43M	.0	264M	129M	.0	.0	19770	6526	134M	1.0	8.7	.0	2.0	40.6	9.3	.0	.0	.0	.0								
04:49:19	8.6	9756K	.0	227M	149M	.0	.0	24358	6185	76M	.0	5.8	.0	.9	19.7	3.9	.0	.0	.0	.0								
04:54:19	1.6	5746K	.0	35M	22M	.0	.0	11401	2840	9553K	.0	.9	.0	.2	4.4	.8	.0	.0	.0	.0								
04:59:19	2.3	9197K	.0	31M	22M	.0	.0	8411	4683	6212K	.1	.8	.0	.2	3.6	.6	.0	.0	.0	.0								
05:04:19	1.6	4755K	.0	18M	14M	.0	.0	9243	4028	4629K	.0	.6	.0	.2	4.1	.7	.0	.0	.0	.0								
05:09:19	2.5	9578K	.0	15M	7938K	.0	.0	29368	7987	1732K	.6	.7	.0	.3	7.2	1.3	.0	.0	.0	.0								
05:14:19	4.4	17M	.0	42M	30M	.0	.0	21723	3400	7576K	2.3	1.1	.0	.5	11.5	2.1	.0	.0	.0	.0								
05:19:19	3.9	12M	.0	34M	24M	.0	.0	10172	4137	8524K	.8	.9	.0	.3	6.6	1.3	.0	.0	.0	.0								
05:24:19	5.0	11M	.0	28M	20M	.0	.0	65700	54272	4361K	.9	.8	.0	.3	6.5	1.3	.0	.0	.0	.0								
05:29:19	1.2	5923K	.0	6831K	5284K	.0	.0	2976	1475	992K	.0	.3	.0	.1	1.9	.3	.0	.0	.0	.0								
05:34:19	2.2	4462K	.0	5996K	3546K	.0	.0	5734	2908	567K	.0	.4	.0	.1	2.6	.4	.0	.0	.0	.0								
05:39:19	1.3	4592K	.0	3935K	3142K	.0	.0	4560	2089	98K	.1	.2	.0	.1	1.8	.3	.0	.0	.0	.0								
05:44:19	.7	2549K	.0	3358K	2303K	.0	.0	7892	6485	1245K	.0	.2	.0	.0	1.3	.2	.0	.0	.0	.0								
05:49:19	2.1	6525K	.0	8039K	5330K	.0	.0	19442	6881	358K	.6	.3	.0	.2	4.2	.8	.0	.0	.0	.0								
Command ==> -																												
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return																												

Figure 190. Layout of Steal Statistics Screen (FCX296)

The information shown is based on CP monitor STORAGE domain SAMPLE data.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### Pct Time Actv

The percentage of the interval in which demand scan was active.

#### Storage/s:

#### Total Stoln

The total amount of storage reclaimed, in bytes per second. Includes storage taken via write-on demand.

#### Write OnDmd

The total amount of storage reclaimed immediately after being written to auxiliary, in bytes per second.

#### User Inval

The amount of storage invalidated (page-table entries invalidated, with pages left resident) on all User Frame Owned lists, in bytes per second. Pages in this invalid-but-resident state are referred to as IBR. Pages are made IBR prior to frame reclaim. Non-faulted pages read in by a block read are added to the User Frame Owned list in the IBR state. Those pages are not included in this count.

#### User Reval

The amount of storage revalidated (because of guest access) on all User Frame Owned lists, in bytes per second. Revalidations of non-faulted pages read in by a block read are included in this count.

#### Shared Inval

The amount of storage invalidated (page-table entries invalidated, with pages left resident) on the Shared Frame Owned list, in bytes per second. Non-faulted pages read in by a block read are added to the Shared Frame Owned list in the IBR state. Those pages are not included in this count.

#### Shared Reval

The amount of storage revalidated (because of guest access) on the Shared Frame Owned list, in bytes per second. Revalidations of non-faulted pages read in by a block read are included in this count.

**Pvt VDisk Inval**

The amount of storage invalidated (page-table entries invalidated, with pages left resident) on the Private VDisk Frame list, in bytes per second. Non-faulted pages read in by a block read are added to the Private VDisk Frame list in the IBR state. Those pages are not included in this count.

**Pvt VDisk Reval**

The amount of storage revalidated (because of guest access) on the Private VDisk Frame List, in bytes per second. Revalidations of non-faulted pages read in by a block read are included in this count.

**AgeL Reval**

The amount of storage revalidated (because of guest access) on the age list, in bytes per second. Revalidations of non-faulted pages read in by a block read are included in this count.

*Completions/s:*

**Need Met**

Number of times demand scan stopped because the need was met, per second.

**Time Limit**

Number of times demand scan stopped because it was running for too long, per second.

**Sys Req**

Number of times demand scan was told to stop because of "external" factors such as a paging overrun or because it was running on an inappropriate CPU, per second.

**Travs/s**

The number of full traversals through all UFOs, the VUFO, and the SUFO, per second.

*User/s:*

**Visit**

Number of UFO lists from which pages were made invalid but resident (IBR), per second.

**Skip**

Number of UFO lists that could not be processed because (PTIL, UFO, or Shadow Table) serialization could not be obtained, or because of the guests' reserved setting, per second.

*Age List:* A staging area for reclaiming frames containing pageable content.

**Stor Skip/s Pin**

The amount of storage skipped from reclaim at the bottom of the age list because the frames were pinned, in bytes per second.

**Stor Skip/s Ser**

The amount of storage skipped because serialization could not be obtained, in bytes per second.

**Stor Skip/s Resv**

The amount of storage requeued from bottom of age list to UFO list, in bytes per second.

**Note:** The count of revalidations can exceed the count of invalidations because revalidations include non-faulted pages read in by a block read that were added to the frame-owned list as IBR pages; the invalidation count includes only those active pages that were invalidated (made IBR) by demand scan.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.



## FCX297, Age List Log Screen – AGELLOG

**Performance Toolkit Subcommand:** “AGELLOG” on page 130.

The Age List Log screen (FCX297) displays "by time" information about z/VM's management of the age list in real memory.

```
FCX297      CPU nnnn   SER nnnnn Interval HH:MM:ss - HH:MM:ss       Perf. Monitor
----->          <------ Storage ----->                <----- Revalidation
                   <-- Steal Ready ---> <--- Not Ready ---->    %Of <----- Storage/s
----->
Interval Size S E K <-List Size--> <--RefOnly--> <--Changed--> <Evaluating-> Pages <--RefOnly--> <--
Changed-->
End Time %DPA Z W S Target Current NoWrt Write Write PndWrt Refd Change Eval NoWrt Write
NoWrt Write
>>Mean>> 2.0 V Y . 5197M 5196M 51191K 0 468M 4256M 1177K 71K 10 16056K .0
17218K 1055K
14:08:57 2.0 V Y N 5197M 5196M 53376K 0 468M 4177M 0 0 10 16903K .0
21555K 1286K
14:09:27 2.0 V Y N 5197M 5197M 48920K 0 470M 4261M 2148K 88K 10 18190K .0
18869K 1367K
14:09:57 2.0 V Y Y 5197M 5197M 45108K 0 475M 4281M 556K 0 10 16552K .0 17909K
412331
14:10:27 2.0 V Y Y 5197M 5197M 35956K 0 483M 4344M 2112K 0 10 16928K .0
18566K 1230K
14:10:57 2.0 V Y Y 5197M 5197M 39816K 0 481M 4350M 0 0 10 14082K .0
12441K 1048K
14:11:27 2.0 V Y Y 5197M 5196M 38084K 0 483M 4320M 0 0 10 15374K .0 10907K
490701

Command ==> _
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return
```

Figure 191. Layout of Age List Log Screen (FCX297)

The information shown is based on CP monitor STORAGE domain SAMPLE data.

### Field Descriptions:

## Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the **>>Mean>>** entry, you can find the relevant interval in the header line.

### Size %DPA

The target age list size as a percentage of the current dynamic paging area (DPA).

## SZ

Indicates whether the global age list is a fixed (F) or variable (V) size.

A dot indicates that the flag setting has changed during the observed interval. Dots appear on the **>>Mean>>** line of the screen only.

## EW

Indicates whether the early writes are enabled (Y) or disabled (N).

A dot indicates that the flag setting has changed during the observed interval. Dots appear on the **>>Mean>>** line of the screen only.

**KS**

Indicates whether the KEEPSLOT option is enabled (Y) or disabled (N).

A dot indicates that the flag setting has changed during the observed interval. Dots appear on the **>>Mean>>** line of the screen only.

A dash appears in this field when running Performance Toolkit against z/VM 6.3 or earlier data.

A dot indicates that the flag setting has changed during the observed interval. Dots appear on the >>Mean>> line of the screen only.

*Storage List Size:*

**Target**

The target storage size of the age list, in bytes.

**Current**

The current storage size of the age list, in bytes.

*Storage Steal Ready:*

**RefOnly NoWrt**

Referenced-only storage that can be reclaimed immediately, in bytes. The page content was not rewritten to auxiliary storage.

**RefOnly Write**

Referenced-only storage that can be reclaimed immediately, in bytes. The page content has been rewritten to auxiliary storage for load balancing.

**Changed Write**

Changed storage that can be reclaimed immediately, in bytes. The page content has been backed up in auxiliary storage.

**Changed PndWrt**

Changed storage that cannot be reclaimed until its page content is backed up in auxiliary storage, in bytes.

*Storage Not Ready:*

**Evaluating Refd**

Referenced-only storage that is not yet reclaimable because changed pages in the same paging block are queued for a write to auxiliary storage, in bytes.

**Evaluating Change**

Changed storage that is not yet reclaimable because the changed page content is queued for a write to auxiliary storage, in bytes.

**%Of Pages Eval**

The percentage of storage in the age list that has been evaluated for write processing.

*Revalidation Storage/s:*

**RefOnly NoWrt**

Referenced-only storage on the age list that was revalidated, in bytes per second. The page content was not written to auxiliary storage before it was revalidated.

**RefOnly Write**

Referenced-only storage on the age list that was revalidated, in bytes per second. The page content was written to auxiliary storage for load balancing before it was revalidated.

**Changed NoWrt**

Changed storage on the age list that was revalidated, in bytes per second. The changed content was not written to auxiliary storage before it was revalidated.

**Changed Write**

Changed storage on the age list that was revalidated, in bytes per second. The changed content was written to auxiliary before it was revalidated.

**Usage note:**

The Steal Ready and Not Ready counts might not add up to the current size of the age list. This is because there might be referenced-only pages in the Not Ready section of the age list.

**Effect of AVERAGE and INTERIM subcommands:**

None.

# FCX298, Logical Core Organization Log Screen – PUORGLOG

**Performance Toolkit Subcommand:** “PUORGLOG” on page 263.

The Logical Core Organization Log Screen (FCX298) shows CPU organization information for a partition.

FCX298	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor
Logical Core organization for Partition FCFT (GDLFCFT )					
Date	Time	Core	Type	PPD	Ent. Location
04/11	14:53:55	00	CP	Vh	... 4:2
04/11	14:53:55	01	CP	Vh	... 4:2
04/11	14:53:55	02	CP	Vh	... 4:2
04/11	14:53:55	03	CP	Vh	... 4:2
04/11	14:53:55	04	CP	Vh	... 4:2
04/11	14:53:55	05	CP	Vm	... 4:1
04/11	14:53:55	06	CP	Vm	... 4:1
04/11	14:53:55	07	CP	Vl	... 4:5
04/11	14:53:55	08	CP	Vl	... 4:5
04/11	14:53:55	09	CP	Vl	... 4:6
04/11	14:53:55	0A	CP	Vl	... 4:6
04/11	14:53:55	0B	CP	Vl	... 4:6
04/11	14:53:55	0C	ZIIP	Vm	... 3:1
04/11	14:53:55	0D	ZIIP	Vm	... 3:1
04/11	14:53:55	0E	IFL	Vm	... 4:5
04/11	14:53:55	0F	IFL	Vl	... 4:5
04/11	14:53:55	10	IFL	Vl	... 4:6
04/11	14:53:55	11	IFL	Vl	... 4:6
Command ==>					
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return					

Figure 192. Layout of Logical Core Organization Log Screen (FCX298)

The information shown is based on CP monitor MONITOR domain and PROCESSOR domain SAMPLE data.

## Field Descriptions:

*Logical Core organization for Partition xxxx (xxxxxxxx):*

The name of the partition that collected the processor unit (core) organization data.

## Logical Core Organization for Partition

### Date

The date (in *mm/dd* format) of the status change.

### Time

The time (in *hh:mm:ss* format) of the status change.

### Core

The logical core ID in hexadecimal format.

### Type

The core type. Possible values are:

#### CP

Central Processor

#### ICF

Internal Coupling Facility

#### IFL

Integrated Facility for Linux

#### ZIIP

IBM Z Integrated Information Processor

### PPD

The polarization and entitlement for the present topology-list entry. Possible values are:

**Vh**

Polarization is vertical, with high entitlement.

**VhD**

Polarization is vertical, with high entitlement, and the partition is a dedicated partition.

**Vm**

Polarization is vertical, with medium entitlement.

**VI**

Polarization is vertical, with low entitlement.

**Hz**

Polarization is horizontal.

**HzD**

Polarization is horizontal, and the partition is a dedicated partition.

**Ent.**

The average processor entitlement for the listed cores as a percentage, where 100% represents one whole physical core's worth of processing capacity.

**Location**

The actual physical location of listed core as it is defined by the architecture. The format depends on the current CEC hardware model. Possible values are:

***m***

The CEC's node ID.

***m:n***

The CEC's node ID (*m*) and chip number (*n*) inside this node.

**Usage note:**

The term *core* describes a unit that contains one or more processors. All processors of a core have the same processor type. For machine types that do not have the multithreading (MT) facility, the machine has physical processors and the partitions have logical processors. For machine types that have the MT facility, the machine has physical cores, each of which contains two physical processors. A partition that has not activated the MT facility continues to have logical processors. A partition that has activated the MT facility has logical cores, each of which contains one or more logical processors.

**Effect of AVERAGE and INTERIM subcommands:**

None.

# FCX299, Processor Unit (Core and Threads) Configuration Log Screen – PUCFGLOG

**Performance Toolkit Subcommand:** “PUCFGLOG” on page 262.

The Processor Unit (Core and Threads) Configuration Log Screen (FCX299) logs the calculations and decisions z/VM makes regarding how many logical processor units of each type are (and are not) actively doing work.

FCX299		CPU nnnn		SER nnnnn		Interval HH:MM:SS - HH:MM:SS		Perf. Monitor										
Date	Time	Type	OnL	Entitl	Type	Cap	CPUPAD	EX	Load	XP	XPF	T/V	LCei	XPF	T/V	N	NotVh	UpCap
02/21	16:09:42	ICF	1	80.0	...	100.0	70		.0	17.9	17.9	256.0	.0	17.3	256.0	1	97.3	97.3
02/21	16:09:42	ZIIP	1	120.0	...	100.0	70		.3	179.8	179.8	256.0	.6	178.9	256.0	1	.0	100.0
02/21 16:09:43 * Processor 00 varied offline																		
02/21 16:09:43 * System Topology changed. Check/Change count is 10472/6																		
02/21 16:09:43 * CPU Park/Unpark State changed																		
02/21	16:09:43	CP	15	571.4	...	100.0	70		150.4	2149.3	2149.3	2.365	267.5	2107.9	9.966	5	.0	500.0
02/21	16:09:43	IFL	2	251.0	...	100.0	70		.0	974.4	974.4	256.0	2.0	965.1	256.0	2	.0	200.0
02/21	16:09:43	ICF	1	80.0	...	100.0	70		.0	17.8	17.8	256.0	.1	17.4	256.0	1	97.4	97.4
02/21	16:09:43	ZIIP	1	120.0	...	100.0	70		.1	179.9	179.9	256.0	.2	179.4	256.0	1	.0	100.0
02/21 16:09:45 * CPU Park/Unpark State changed																		
02/21	16:09:45	CP	15	571.4	...	100.0	70		184.1	2078.9	2078.9	1.711	261.2	2008.5	9.981	5	.0	500.0
02/21	16:09:45	IFL	2	251.0	...	100.0	70		.0	968.9	968.9	256.0	1.9	963.3	256.0	2	.0	200.0
02/21	16:09:45	ICF	1	80.0	...	100.0	70		.0	17.7	17.7	256.0	.0	17.4	256.0	1	97.4	97.4
02/21	16:09:45	ZIIP	1	120.0	...	100.0	70		.0	180.0	180.0	256.0	.2	179.9	256.0	1	.0	100.0
02/21	16:10:03	CP	15	571.4	...	100.0	70		234.7	2126.7	2126.7	2.019	257.0	2108.5	9.808	5	.0	500.0
02/21	16:10:03	IFL	2	251.0	...	100.0	70		1.6	969.9	969.9	256.0	3.1	964.7	256.0	2	.0	200.0
02/21	16:10:03	ICF	1	80.0	...	100.0	70		.0	17.8	17.8	256.0	.0	17.2	256.0	1	97.2	97.2
02/21	16:10:03	ZIIP	1	120.0	...	100.0	70		.0	180.0	180.0	256.0	.0	180.0	256.0	1	.0	100.0
02/21 16:10:05 * System Topology changed. Check/Change count is 10483/7																		
02/21 16:10:05 * Processor 00 varied online																		
02/21 16:10:05 * CPU Park/Unpark State changed																		
02/21	16:10:05	CP	16	571.4	...	100.0	70		227.9	2168.0	2168.0	2.146	261.7	2128.5	9.798	5	.0	500.0
02/21	16:10:05	IFL	2	251.0	...	100.0	70		2.6	979.8	979.8	256.0	3.6	965.0	256.0	2	.0	200.0
02/21	16:10:05	ICF	1	80.0	...	100.0	70		.0	17.7	17.7	256.0	.0	17.2	256.0	1	97.2	97.2
02/21	16:10:05	ZIIP	1	120.0	...	100.0	70		.0	180.0	180.0	256.0	.0	180.0	256.0	1	.0	100.0
Command ==>																		
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return																		

Figure 193. Processor Unit (Core and Threads) Configuration Log Screen (FCX299)

The information shown is based on CP monitor PROCESSOR domain EVENT data.

## Field Descriptions:

### Date

The date (in *mm/dd* format) of the status change.

### Time

The time (in *hh:mm:ss* format) of the status change.

**Note:** Every record (group of lines, by type, if applicable) always appears in a two-second interval. If FC SET BYTIME sets to any nonzero value, Performance Toolkit keeps only specified timestamps, rounded to the nearest next sample interval, plus one before and one after the significant external event. (See the note in the description of the Type field).

### Type

The core type. Possible values are:

#### CP

Central Processor

#### ICF

Internal Coupling Facility

#### IFL

Integrated Facility for Linux

**ZIIP**

IBM Z Integrated Information Processor

**Note:** When the Type field contains an asterisk, the line denotes a significant external event. Possible values are:

- \* CPU Park/Unpark State changed
- \* Processor *nn* varied offline (where *nn* is a hexadecimal value)
- \* Processor *nn* varied online (where *nn* is a hexadecimal value)
- \* System Topology changed. Check/Change count is *nnnnn/m*

**OnL**

The number of logical processors that are online.

**Entitl**

The current entitlement for the listed core type as a percentage, where 100% represents one whole physical core's worth of processing capacity.

This field is available in an environment that provides the information needed to make excess capacity predictions. Otherwise, dots are displayed.

**Type Cap**

The maximum amount of core capacity that this partition is permitted to use, as a percentage, where 100% represents one whole physical core's worth of processing capacity.

**CPUPAD**

The CPUPAD value set by the administrator using the SET SRM CPUPAD command.

**EX**

The confidence percentage used to make floor projection of unentitled capacity. The value is driven by the SET SRM EXCESSUSE command for this processor type.

**Last Load**

The last actual core utilization measurement in this partition, as a percentage, where 100% represents one whole physical core's worth of processing capacity.

**Last XP**

The calculated amount of excess core capacity available on shared physical cores in the CEC over the interval: the amount of physical core capacity of this type that is available to all shared LPARs that are competing for core capacity beyond their entitlement (including core capacity above entitlement used and unused capacity).

Shown as a percentage, where 100% represents one whole physical core's worth of processing capacity.

This field is available in an environment that provides the information needed to make excess capacity predictions. Otherwise, dots are displayed.

**Last XPF**

The calculated amount of excess core capacity in the interval: the amount of XP that this partition could have used, had it wanted to use it.

Shown as a percentage, where 100% represents one whole physical core's worth of processing capacity.

This field is available in an environment that provides the information needed to make excess capacity predictions. Otherwise, dots are displayed.

**Last T/V**

The calculated T/V ratio in the interval. This field contains dots when T/V was not calculated for the most recent interval.

**Next LCei**

The core utilization ceiling prediction for the next interval, as a percentage, where 100% represents one whole physical core's worth of processing capacity.

**Next XPF**

The excess core capacity floor projection for the next interval, as a percentage, where 100% represents one whole physical core's worth of processing capacity.

**Next T/V**

The ceiling projection for T/V ratio for the next interval. This field contains dots when a prediction was not made for the next interval.

**Next N**

The number of logical processors that will be in unparked state for the next interval.

**Next NotVh**

Displays the amount of power a non-Vh core is bringing to this partition in the next interval.

**Next UpCap**

Displays what this partition's unparked core capacity will be in the next interval.

**Usage note:**

The term *core* describes a unit that contains one or more processors. All processors of a core have the same processor type. For machine types that do not have the multithreading (MT) facility, the machine has physical processors and the partitions have logical processors. For machine types that have the MT facility, the machine has physical cores, each of which contains two physical processors. A partition that has not activated the MT facility continues to have logical processors. A partition that has activated the MT facility has logical cores, each of which contains one or more logical processors.

**Effect of AVERAGE and INTERIM subcommands:**

None.

# FCX300, Dispatch Vector Configuration Change Log Screen – DSVCLLOG

**Performance Toolkit Subcommand:** “DSVCLLOG” on page 174.

The Dispatch Vector Configuration Change Log Screen (FCX300) shows information about each dispatch vector.

FCX300      CPU nnnn   SER nnnnn   Interval HH:MM:SS - HH:MM:SS   Perf. Monitor

Dispatch Vector Configuration Change log

Date	Time	DVID	Type	Location	Count	Pz
04/22	09:32:42	0000	CP	2	1	Vh
04/22	09:32:42	0001	CP	2	1	Vh
04/22	09:32:42	0002	CP	2	1	Vh
04/22	09:32:42	0003	CP	2	1	Vh
04/22	09:32:42	0004	CP	2	1	Vh
04/22	09:32:42	0005	CP	2	1	Vm
04/22	09:32:42	0006	CP	2	1	Vl
04/22	09:32:42	0007	CP	2	1	Vl
04/22	09:32:42	000A	ICF	4	1	Vm
04/22	09:32:42	000B	IFL	2	1	Vh
04/22	09:32:42	000C	IFL	2	1	Vh
04/22	09:32:42	000D	CP	2	1	Vl
04/22	09:32:42	000E	CP	2	1	Vl
04/22	09:32:42	000F	CP	2	1	Vl
04/22	09:32:42	0010	CP	2	1	Vl
04/22	09:32:42	0011	CP	2	1	Vl
04/22	09:32:42	0012	CP	2	1	Vl
04/22	09:32:42	0013	CP	2	1	Vl
04/22	09:32:42	0014	CP	4	1	Vl

Command ==>  
F1=Help   F4=Top   F5=Bot   F7=Bkwd   F8=Fwd   F12=Return

Figure 194. Dispatch Vector Configuration Change Log Screen (FCX300)

The information shown is based on CP monitor PROCESSOR domain EVENT data.

**Field Descriptions:**

**Date**

The date (in *mm/dd* format) of the status change.

**Time**

The time (in *hh:mm:ss* format) of the status change.

**DVID**

The dispatch vector ID.

**Type**

The core type. Possible values are:

**CP**

Central Processor

**ICF**

Internal Coupling Facility

**IFL**

Integrated Facility for Linux

**ZIIP**

IBM Z Integrated Information Processor

**Location**

The actual physical location of the listed core as it is defined by the architecture. The format depends on the current CEC hardware model. Possible values are:



***m***

The CEC's node ID.

***m:n***

The CEC's node ID (*m*) and chip number (*n*) inside this node.

**Count**

The total number of CPUs associated with the dispatch vector at the moment the sample was snapped.

**Pz**

The polarization and entitlement for the present topology-list entry. Possible values are:

For Vertical polarization:

**Vh**

High entitlement

**Vm**

Medium entitlement

**Vl**

Low entitlement

For horizontal polarization:

**Hz**

Horizontal

**Effect of AVERAGE and INTERIM subcommands:**

None.

# FCX301, Dispatch Vector Activity Screen – DSVBKACT

**Performance Toolkit Subcommand:** “DSVBKACT” on page 173.

The Dispatch Vector Activity Screen (FCX301) shows the activity of the DSVBKs.

FCX301	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor
Total reassignments:	20		Life Time:	Max= 3000	
Maximum users per DSV:	14			Min= 3000	
				Avg= 3000	
-----	.	.	<-- LPU -->	Pct	Mean
Epoch	DVID	Type	Assoc Unprk Empty	Non0	
20	0000	CP	1 1	76.7	1.000
20	0001	CP	1 1	100.0	...
20	0002	CP	1 1	100.0	...
20	0003	CP	1 1	96.7	1.000
20	0004	CP	1 1	100.0	...
20	0005	CP	1 1	100.0	...
20	0006	CP	1 1	100.0	...
20	0007	CP	1 1	100.0	...
20	0008	CP	1 1	100.0	...
20	0009	CP	1 1	100.0	...
20	000A	CP	1 1	100.0	...
20	000B	CP	1 1	100.0	...
20	000C	ZIIP	1 1	100.0	...
20	000D	ZIIP	1 1	100.0	...
20	000E	IPL	1 1	100.0	...
20	000F	IPL	1 1	100.0	...
20	0010	IPL	1 1	100.0	...
20	0011	IPL	1 0	100.0	...
Command ==>					
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return					

Figure 195. Dispatch Vector Activity Screen (FCX301)

The information shown is based on CP monitor PROCESSOR domain EVENT and SAMPLE data.

The sample shows a screen with DSVBKs sorted in ascending the counter of reassignments sequence (indicated by the underscore characters above the Epoch heading), the initial sequence. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

### Field Descriptions:

#### Total reassignments

The number of times CPU-to-DSVBK assignments have changed since the last IPL.

#### Life Time

The lifetime for displayed dispatch vectors, in seconds. When the value exceeds 86400 seconds, an approximate value is displayed as *nnnd*, where *nnnn* is the number of days. This field shows maximum, minimum, and average times.

#### Maximum users per DSV

The maximum number of users allowed in one dispatch vector.

#### Epoch

The counter of reassignments for this dispatch vector.

#### DVID

The dispatch vector ID. The special ID x'FFFF' shows the master-only dispatch vector.

#### Type

The core type. Possible values are:

**CP**

Central Processor

**ICF**

Internal Coupling Facility

**IFL**

Integrated Facility for Linux

**ZIIP**

IBM Z Integrated Information Processor

**LPU Assoc**

The total number of CPUs associated with the dispatch vector at the moment the sample was snapped.

**LPU Unprk**

The number of unparked CPUs associated with the dispatch vector at the moment the sample was snapped.

**Pct Empty**

The percentage of high-frequency samples that showed that the dispatch vector had no VMDBKs in it.

**Mean Non0**

The mean number of elements (VMDBKs) in the dispatch vector for those samples that showed a nonzero number of VMDBKs.

**Effect of AVERAGE, CURRENT, and INTERIM subcommands:**

Overall or interim period averages will be shown for all data fields.

## FCX302, Real Core Utilization Log Screen – PHYSLOG

**Performance Toolkit Subcommand:** “PHYSLOG” on page 249.

The Real Core Utilization Log Screen (FCX302) shows real core utilization by time.

```
FCX302      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS      Perf. Monitor

Interval    <PhCore> Shrd  Total
End Time   Type Conf Ded Log. Weight %LgcLC %Ovrhd LCoT/L %LPmgt %Total TypeT/L
>>Mean>> CP      18   8   79    320 1370.2  5.171  1.004 17.462 1392.8  1.017
>>Mean>> IFL     16   0   35     80  5.692   .315  1.055  2.156  8.162  1.434
>>Mean>> ICF       1   0   0      0  .000   .000   ...  3.371  3.371   ...
>>Mean>> ZIIP      3   0   0      0  .000   .000   ...  .530  .530   ...
>>Mean>> >Sum     38   8  114    400 1375.9  5.486  1.004 23.520 1404.9  1.021

18:14:14 CP      18   8   79    320 1372.9  5.693  1.004 19.661 1398.3  1.018
18:14:14 IFL     16   0   35     80  8.057   .317  1.039  1.307  9.681  1.202
18:14:14 ICF       1   0   0      0  .000   .000   ...  3.176  3.176   ...
18:14:14 ZIIP      3   0   0      0  .000   .000   ...  .102  .102   ...
18:14:14 >Sum     38   8  114    400 1381.0  6.010  1.004 24.247 1411.2  1.022

18:15:14 CP      18   8   79    320 1370.8  5.999  1.004 19.770 1396.6  1.019
18:15:14 IFL     16   0   35     80  5.687   .321  1.056  1.994  8.002  1.407
18:15:14 ICF       1   0   0      0  .000   .000   ...  3.280  3.280   ...
18:15:14 ZIIP      3   0   0      0  .000   .000   ...  .105  .105   ...
18:15:14 >Sum     38   8  114    400 1376.5  6.319  1.005 25.149 1407.9  1.023

18:16:14 CP      18   8   79    320 1379.4  4.313  1.003 15.476 1399.2  1.014
18:16:14 IFL     16   0   35     80  5.949   .302  1.051  2.778  9.029  1.518
18:16:14 ICF       1   0   0      0  .000   .000   ...  3.342  3.342   ...
18:16:14 ZIIP      3   0   0      0  .000   .000   ...  .099  .099   ...
18:16:14 >Sum     38   8  114    400 1385.4  4.615  1.003 21.695 1411.7  1.019

Command ==>
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 196. Real Core Utilization Log Screen (FCX302)

The information shown is based on CP monitor PROCESSOR domain EVENT and SAMPLE data.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. The >>Mean>> line shows the overall average values per core type.

#### Type

The core type of this type-pool of physical cores. Possible values are:

##### CP

Central Processor

##### ICF

Internal Coupling Facility

##### IFL

Integrated Facility for Linux

##### ZIIP

IBM Z Integrated Information Processor

A line with >Sum in the Type field shows the sum over all core types.

#### PhCore Conf

The number of configured physical cores of this type.

#### PhCore Ded

The number of dedicated physical cores of this type.

#### Shrd Log.

The number of shared logical cores of this type.

**Total Weight**

For this type, the sum of the partition weights.

**%Lgc1C**

For this physical core type, the physical capacity that is used running logical cores of the corresponding type.

The value ranges from 0% to  $n \times 100\%$ , where  $n$  is the number of physical cores of this type.

**%Ovrhd**

For this physical core type, the physical capacity that is used performing LPAR management activities that are directly chargeable or attributable to the actions of any specific logical processors of the corresponding type.

The value ranges from 0% to  $n \times 100\%$ , where  $n$  is the number of physical cores of this type.

**LCot/L**

For this physical core type, the calculated value  $(\%Lgc1C + \%Ovrhd) / \%Lgc1p$ .

**%LPmgt**

For this physical core type, the physical capacity that is used performing LPAR management activities that are not directly chargeable or attributable to the actions of any specific logical processors of the corresponding type.

The value ranges from 0% to  $n \times 100\%$ , where  $n$  is the number of physical cores of this type.

**%Total**

For this physical core type, the calculated value  $\%Lgc1C + \%Ovrhd + \%LPmgt$ .

**TypeT/L**

For this physical core type, the calculated value  $\%Total / \%Lgc1C$ .

**Usage note:**

The term *core* describes a unit that contains one or more processors. All processors of a core have the same processor type. For machine types that do not have the multithreading (MT) facility, the machine has physical processors and the partitions have logical processors. For machine types that have the MT facility, the machine has physical cores, each of which contains two physical processors. A partition that has not activated the MT facility continues to have logical processors. A partition that has activated the MT facility has logical cores, each of which contains one or more logical processors.

**Effect of AVERAGE and INTERIM subcommands:**

None.

## FCX303, DSVBK Steals per Processor Log Screen – DSVSLOG

**Performance Toolkit Subcommand:** “DSVSLOG” on page 175.

The DSVBK Steals per Processor Log Screen (FCX303) displays information about DSVBK steal rates. It should be used in conjunction with “FCX304, Processor Log Screen – PRCLOG” on page 825.

FCX303	CPU nnnn			SER nnnnn		Interval HH:MM:SS - HH:MM:SS					Perf. Monitor			
Interval						Pct	DSVBK Steal /s					Core/		
End Time	CPU	Type	PPD	Ent.	DVID	Park	Total	Lvl-00	Lvl-01	Lvl-02	Lvl-03	Lvl-04	Lvl-05	Thread
>>Mean>>	00	CP	Vh	100	0000	0	52.200	49.479	2.721	.000	....	....	....	00/0
>>Mean>>	01	CP	Vh	100	0001	0	133.42	125.96	7.458	.000	....	....	....	01/0
>>Mean>>	02	CP	Vh	100	0002	0	119.10	110.76	8.333	.000	....	....	....	02/0
>>Mean>>	03	CP	Vh	100	0003	0	199.63	184.80	14.821	.000	....	....	....	03/0
>>Mean>>	04	CP	Vh	100	0004	0	170.41	156.60	13.808	.000	....	....	....	04/0
>>Mean>>	05	CP	Vh	100	0005	0	6.950	.225	6.725	.000	....	....	....	05/0
>>Mean>>	06	CP	Vm	79	0006	0	3.588	.267	3.321	.000	....	....	....	06/0
>>Mean>>	07	CP	Vl	0	0007	0	5.559	.313	5.246	.000	....	....	....	07/0
>>Mean>>	08	CP	Vl	0	0008	0	5.651	.163	5.488	.000	....	....	....	08/0
>>Mean>>	09	CP	Vl	0	0009	0	6.133	.179	5.954	.000	....	....	....	09/0
>>Mean>>	0A	CP	Vl	0	000A	0	4.080	.067	4.013	.000	....	....	....	0A/0
>>Mean>>	0B	CP	Vl	0	000B	0	8.346	.313	8.033	.000	....	....	....	0B/0
>>Mean>>	0C	ZIIP	Vm	50	000C	0	.000	.000	.000	.000	....	....	....	0C/0
>>Mean>>	0D	ZIIP	Vm	50	000D	0	.000	.000	.000	.000	....	....	....	0D/0
>>Mean>>	0E	IFL	Vm	57	000E	0	.017	.017	.000	.000	....	....	....	0E/0
>>Mean>>	0F	IFL	Vl	0	000F	0	.166	.158	.008	.000	....	....	....	0F/0
>>Mean>>	10	IFL	Vl	0	0010	0	.538	.000	.538	.000	....	....	....	10/0
>>Mean>>	11	IFL	Vl	0	0011	100	.000	.000	.000	.000	....	....	....	11/0
>>Total>	12	CP	MIX	679	MIX	0	715.05	629.13	85.921	.000	....	....	....	MIX
>>Total>	4	IFL	MIX	57	MIX	100	.721	.175	.546	.000	....	....	....	MIX
>>Total>	2	ZIIP	Vm	100	000D	0	.000	.000	.000	.000	....	....	....	MIX
08:25:41	00	CP	Vh	100	0000	0	56.866	42.733	14.133	.000	....	....	....	00/0
08:25:41	01	CP	Vh	100	0001	0	97.416	83.983	13.433	.000	....	....	....	01/0
08:25:41	02	CP	Vh	100	0002	0	133.22	106.07	27.150	.000	....	....	....	02/0
08:25:41	03	CP	Vh	100	0003	0	176.18	118.15	58.033	.000	....	....	....	03/0
08:25:41	04	CP	Vh	100	0004	0	165.65	135.77	29.883	.000	....	....	....	04/0
08:25:41	05	CP	Vh	100	0005	0	14.067	.200	13.867	.000	....	....	....	05/0
08:25:41	06	CP	Vm	79	0006	0	2.833	.000	2.833	.000	....	....	....	06/0
08:25:41	07	CP	Vl	0	0007	0	.467	.000	.467	.000	....	....	....	07/0
08:25:41	08	CP	Vl	0	0008	0	23.300	1.683	21.617	.000	....	....	....	08/0
08:25:41	09	CP	Vl	0	0009	0	24.534	4.267	20.267	.000	....	....	....	09/0
08:25:41	0A	CP	Vl	0	000A	0	6.083	1.083	5.000	.000	....	....	....	0A/0
08:25:41	0B	CP	Vl	0	000B	0	10.933	.333	10.600	.000	....	....	....	0B/0
08:25:41	0C	ZIIP	Vm	50	000C	0	.000	.000	.000	.000	....	....	....	0C/0
08:25:41	0D	ZIIP	Vm	50	000D	0	.000	.000	.000	.000	....	....	....	0D/0
08:25:41	0E	IFL	Vm	57	000E	0	.150	.117	.033	.000	....	....	....	0E/0
08:25:41	0F	IFL	Vl	0	000F	0	.300	.233	.067	.000	....	....	....	0F/0
08:25:41	10	IFL	Vl	0	0010	0	.017	.000	.017	.000	....	....	....	10/0
08:25:41	11	IFL	Vl	0	0011	100	.000	.000	.000	.000	....	....	....	11/0
Command ==>														
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return									

Figure 197. DSVBK Steals per Processor Log Screen (FCX303)

The information shown is based on CP monitor SYSTEM domain and PROCESSOR domain SAMPLE data.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated. In the case of the >>Mean>> and >>Total> entries, you can find the relevant interval in the header line. The >>Mean>> line shows the overall average values per processor number. The >>Total> line shows the overall average values per processor type when the count of active processors in the overall interval is more than one.

#### CPU

The logical processor ID in hexadecimal format.

The >>Total> line contains a decimal value that represents the total number of logical processors of this type in the partition.

#### **Type**

The processor type. Possible values are:

##### **CP**

Central Processor

##### **ICF**

Internal Coupling Facility

##### **IFL**

Integrated Facility for Linux

##### **ZIIP**

IBM Z Integrated Information Processor

#### **PPD**

The logical processor polarization. Possible values are:

##### **Hz**

Polarization is horizontal.

##### **HzD**

Polarization is horizontal, and the partition is a dedicated partition.

##### **Vh**

Polarization is vertical, with high entitlement.

##### **VhD**

Polarization is vertical, with high entitlement, and the partition is a dedicated partition.

##### **Vm**

Polarization is vertical, with medium entitlement.

##### **VL**

Polarization is vertical, with low entitlement.

##### **MIX**

A mixed polarization, which can be displayed:

- On the >>Mean>> and >>Total> lines if multiple distinct polarization values were encountered during the measured interval (average, from the last reset).
- If a set of processors (on the >>Total> line only) has distinct polarization values.
- On an Interval End Time line if the polarization value for a CPU changes during the interval for which the data is displayed.

#### **Ent.**

The average processor entitlement for the listed CPU as a percentage, where 100% represents one whole physical CPU's worth of processing capacity.

#### **DVID**

A dispatch vector ID that serves the current processor, with hexadecimal representation, or the value MIX on the >>Mean>> and >>Total> lines if distinct dispatch vectors served the CPU. For other lines, the field contains an ID of the latest dispatch vector.

#### **Pct Park Time**

The percentage of the time that the CPU was in a parked state.

#### *DSVBK Steal /s*

The average rate of user block steals for the current processor or group of processors (for Total part). The main group of columns is a set of VMDBK Steal values, grouped by a topological distance of migration.

#### **Total**

The summary rate per second for all distances (levels).

The count of these columns depends on the Maximum Nest value (MNest).

**Lvl-00**

The rate of steals from dispatch vector blocks (DSVBKs) with topologically-equivalent CPUs.

**Lvl-01..Lvl-05**

The rate of steals from DSVBKs with CPUs that are a corresponding topology level away, if it is applicable by current machine architecture.

**Core/Thread**

The core and thread IDs of the current processor, with hexadecimal representation as *cc/t* or a special word MIX for the >>Mean>> and >>Total> lines if distinct cores or threads served the CPU. For other lines, the field contains latest IDs.

**Effect of AVERAGE and INTERIM subcommands:**

None.



## FCX304, Processor Log Screen – PRCLOG

**Performance Toolkit Subcommand:** “PRCLOG” on page 251.

The Processor Log Screen (FCX304) displays logical processor activity by time.

This report should be used in conjunction with “FCX303, DSVBK Steals per Processor Log Screen – DSVSLOG” on page 822.

FXX304	CPU nnnn		SER nnnnn	Interval HH:MM:SS - HH:MM:SS		Perf. Monitor																			
						<--- Percent Busy >---					<-- Rates per Sec. >--					<----- Paging ----->					<Co>	< Di>			
Interval	End Time	CPU	Type	PPD	Ent.	DVID	Pct Park	%Susp	Total	User	Syst	Emul	Inst Siml	DIAG	SI GP	SSCH	<26B /s	PGIN /s	Fast Path %	Page Read /s	<m> /s	<ag> X'9C'	<s /s	Core/Thrd	
>>Mean>	00 CP Vh	100	0000	0	.1	29.1	28.2	.9	27.8	204.4	40.4	789.2	51.1	.0	.0	.0	.0	.0	.0	.0	.8	.0	00/0		
>>Mean>	01 CP Vh	100	0001	0	.1	31.3	30.7	.6	30.3	186.9	14.0	378.1	.6	.0	.0	.0	.0	.0	.0	.4	.0	01/0			
>>Mean>	02 CP Vh	100	0002	0	.1	16.5	15.6	.9	15.0	205.9	13.4	463.8	1.4	.0	.0	.0	.0	.0	.0	.5	.0	02/0			
>>Mean>	03 CP Vh	100	0003	0	.2	15.1	14.0	1.1	13.2	210.7	10.4	452.2	1.7	.0	.0	.0	.0	.0	.0	.6	.0	03/0			
>>Mean>	04 CP Vh	100	0004	0	.1	15.1	14.1	1.0	13.4	202.6	12.0	430.5	1.0	.0	.0	.0	.0	.0	.0	.5	.0	04/0			
>>Mean>	05 CP Vh	100	0005	0	.0	.8	.7	.1	.7	108.5	.1	49.1	.0	.0	.0	.0	.0	.0	.0	.0	.0	05/0			
>>Mean>	06 CP Vm	79	0006	0	.0	.9	.9	.0	.9	20.0	.1	12.3	.0	.0	.0	.0	.0	.0	.0	.0	.0	06/0			
>>Mean>	07 CP V1	0	0007	0	.0	.9	.9	.0	.9	15.8	.7	11.0	.0	.0	.0	.0	.0	.0	.0	.3	.0	07/0			
>>Mean>	08 CP V1	0	0008	0	.0	.5	.5	.0	.5	21.1	.1	9.6	.0	.0	.0	.0	.0	.0	.0	.1	.0	08/0			
>>Mean>	09 CP V1	0	0009	0	.0	.6	.5	.0	.5	7.1	.1	5.1	.0	.0	.0	.0	.0	.0	.0	.4	.0	09/0			
>>Mean>	0A CP V1	0	000A	0	.0	.1	.1	.0	.1	17.3	.0	7.5	.0	.0	.0	.0	.0	.0	.0	.0	.0	0A/0			
>>Mean>	0B CP V1	0	000B	0	.0	.8	.8	.1	.7	154.5	.1	65.1	.0	.0	.0	.0	.0	.0	.0	.0	.0	0B/0			
>>Mean>	0C ZIIP Vm	50	000C	0	.0	.0	.0	.0	.0	.6	.0	1.5	.0	.0	.0	.0	.0	.0	.0	.0	.0	0C/0			
>>Mean>	0D ZIIP Vm	50	000D	0	.0	.0	.0	.0	.0	.0	.0	1.5	.0	.0	.0	.0	.0	.0	.0	.0	.0	0D/0			
>>Mean>	0E IFL Vm	57	000E	0	.0	.3	.2	.1	.1	206.8	.1	56.1	.0	.0	.0	.0	.0	.0	.0	.0	.0	0E/0			
>>Mean>	0F IFL V1	0	000F	0	.0	.1	.1	.0	.0	70.3	.0	19.7	.0	.0	.0	.0	.0	.0	.0	.0	.0	0F/0			
>>Mean>	10 IFL V1	0	0010	0	.0	.1	.0	.0	.0	34.7	.0	10.4	.0	.0	.0	.0	.0	.0	.0	.0	.0	10/0			
>>Mean>	11 IFL V1	0	0011	100	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	11/0			
>>Total>	12 CP MIX	679	MIX	0	.6	111.7	107.0	4.7	104.0	1355	91.3	2674	56.0	.0	.0	.0	.0	.0	.0	3.8	.0	MIX			
>>Total>	4 IFL MIX	57	MIX	100	.0	.4	.3	.1	.2	311.8	.1	86.2	.0	.0	.0	.0	.0	.0	.0	.0	.0	MIX			
>>Total>	2 ZIIP Vm	100	MIX	0	.0	.0	.0	.0	.0	2.1	.0	3.5	.0	.0	.0	.0	.0	.0	.0	.0	.0	MIX			
08:25:41	00 CP Vh	100	0000	0	.1	26.2	25.2	1.0	24.7	168.0	80.2	817.7	51.2	.0	.0	.0	.0	.0	.0	.6	.0	00/0			
08:25:41	01 CP Vh	100	0001	0	.1	31.8	31.2	.6	30.8	126.1	10.9	344.4	.5	.0	.0	.0	.0	.0	.0	.5	.0	01/0			
08:25:41	02 CP Vh	100	0002	0	.1	16.0	15.2	.8	14.6	142.3	50.3	451.1	1.2	.0	.0	.0	.0	.0	.0	.7	.0	02/0			
08:25:41	03 CP Vh	100	0003	0	.2	14.4	13.7	1.2	12.9	178.4	47.6	461.2	.2	.0	.0	.0	.0	.0	.0	.5	.0	03/0			
08:25:41	04 CP Vh	100	0004	0	.1	14.3	13.3	.9	12.7	136.9	13.9	361.0	1.2	.0	.0	.0	.0	.0	.0	.6	.0	04/0			
08:25:41	05 CP Vh	100	0005	0	.0	.9	.9	.1	.8	100.0	.0	42.6	.0	.0	.0	.0	.0	.0	.0	.0	.0	05/0			
08:25:41	06 CP Vm	79	0006	0	.0	.0	.0	.0	.0	.1	.3	.0	.0	.0	.0	.0	.0	.0	.0	.1	.0	06/0			
08:25:41	07 CP V1	0	0007	0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	07/0			
08:25:41	08 CP V1	0	0008	0	.3	3.3	3.3	.2	3.2	311.3	.1	156.9	.0	.0	.0	.0	.0	.0	.0	.0	.0	08/0			
08:25:41	09 CP V1	0	0009	0	.3	3.2	3.1	.1	3.1	163.8	.2	92.5	.0	.0	.0	.0	.0	.0	.0	.0	.0	09/0			
08:25:41	0A CP V1	0	000A	0	.0	.8	.7	.1	.6	236.1	.1	122.5	.0	.0	.0	.0	.0	.0	.0	.0	.0	0A/0			
08:25:41	0B CP V1	0	000B	0	.0	.9	.8	.1	.8	122.0	.1	55.2	.0	.0	.0	.0	.0	.0	.0	.0	.0	0B/0			
08:25:41	0C ZIIP Vm	50	000C	0	.0	.0	.0	.0	.0	.0	.0	1.1	.0	.0	.0	.0	.0	.0	.0	.0	.0	0C/0			
08:25:41	0D ZIIP Vm	50	000D	0	.0	.0	.0	.0	.0	.0	.0	1.2	.0	.0	.0	.0	.0	.0	.0	.0	.0	0D/0			
08:25:41	0E IFL Vm	57	000E	0	.0	.1	.1	.0	.1	88.9	.0	28.1	.0	.0	.0	.0	.0	.0	.0	.0	.0	0E/0			
08:25:41	0F IFL V1	0	000F	0	.0	.6	.4	.2	.2	414.3	.1	105.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	0F/0			
08:25:41	10 IFL V1	0	0010	0	.0	.0	.0	.0	.0	.0	.0	.5	.0	.0	.0	.0	.0	.0	.0	.0	.0	10/0			
08:25:41	11 IFL V1	0	0011	100	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	11/0			
Command ==>																									
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return																									

Figure 198. Layout of Processor Log Screen (FCX304)

The information shown is based on CP monitor SYSTEM domain and PROCESSOR domain SAMPLE data.

### Field Descriptions:

## Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the `>>Mean>>` and `>>Total>` entries, you can find the relevant interval in the header line. The `>>Mean>>` line shows the overall average values per processor number. The `>>Total>` line shows the overall average values per processor type when the count of active processors in the overall interval is more than one.

## CPU

The logical processor ID in hexadecimal format.

The `>>Total>` line contains a decimal value that represents the total number of logical processors of this type in the partition.

## Type

The processor type. Possible values are:

**CP**

## Central Processor

**ICF**

### Internal Coupling Facility

**IFL**

Integrated Facility for Linux

**ZIIP**

IBM Z Integrated Information Processor

**PPD**

The logical processor polarization. Possible values are:

**Hz**

Polarization is horizontal.

**HzD**

Polarization is horizontal, and the partition is a dedicated partition.

**Vh**

Polarization is vertical, with high entitlement.

**VhD**

Polarization is vertical, with high entitlement, and the partition is a dedicated partition.

**Vm**

Polarization is vertical, with medium entitlement.

**VI**

Polarization is vertical, with low entitlement.

**MIX**

A mixed polarization, which can be displayed:

- On the >>Mean>> and >>Total> lines if multiple distinct polarization values were encountered during the measured interval (average, from the last reset).
- If a set of processors (on the >>Total> line only) has distinct polarization values.
- On an Interval End Time line if the polarization value for a CPU changes during the interval for which the data is displayed.

**Ent.**

The average processor entitlement for the listed CPU as a percentage, where 100% represents one whole physical CPU's worth of processing capacity.

**DVID**

A dispatch vector ID that serves the current processor, with hexadecimal representation, or the value MIX on the >>Mean>> and >>Total> lines if distinct dispatch vectors served the CPU. For other lines, the field contains an ID of the latest dispatch vector.

**Pct Park Time**

The percentage of the time that the CPU was in a parked state.

**%Susp**

The average percentage of elapsed time that logical processors were "suspended", that is, were neither running, nor in a wait state, nor in a parked state.

*Percent Busy:*

**Total**

The percentage of total elapsed time that the processor was busy.

**User**

The percentage of total elapsed time that the processor time was attributed to specific users, including emulation time and CP time.

**Syst**

The percentage of total elapsed time that the processor was charged to the system (CP time not attributed to specific users).

**Emul**

The percentage of total elapsed time that the processor was in emulation mode, doing work for users.

Note that all of the previous percentages are based on *elapsed time*. For a system running in an LPAR, this could mean that it is running at maximum capacity even though none of the values have reached 100%. See “FCX100, CPU Load and Transactions – CPU” on page 371 for “busy” values, which are based on the sum of “busy” and voluntary wait time, instead of on elapsed time.

*Rates per Sec.:*

**Inst Siml**

The number of instructions simulated per second.

**DIAG**

The number of Diagnose instructions that guests executed per second.

**SIGP**

The rate at which this CPU is the target of SIGP External Call instructions.

**SSCH**

The number of Start Subchannel or Resume Subchannel instructions that CP executed per second.

*Paging:*

**Below 2GB/s**

The number of pages moved below the 2GB line for translation, per second. The field is available only for z/VM systems running in 64-bit mode.

**PGIN/s**

The number of PGINs per second (page moves from expanded to central storage).

**Fast Path %**

The percentage of PGINs that could be handled by fastpath logic.

**Page Reads/s**

The number of pages read from auxiliary storage, including directory I/O, nucleus I/O, and warmstart I/O, but no spool I/O.

*Msgs/s:*

**Comm Msgs/s**

The number of IUCV and VMCF data transfers per second (successful and unsuccessful) for this processor.

*Diag:*

**Diag X'9C'/s**

The number of times Diagnose X'9C' was issued by CP per second for this processor.

**Core/Thrd**

The core and thread IDs of the current processor, in hexadecimal format (*cc/t*), or the value MIX on the >>Mean>> and >>Total> lines if distinct cores or threads served the CPU. For other lines, this field contains the latest IDs.

**Effect of AVERAGE and INTERIM subcommands:** None.

# FCX305, Logical Partition Activity Menu Screen – LPARMENU

**Performance Toolkit Subcommand:** “LPARMENU” on page 224.

The Logical Partition Activity Menu Screen (FCX305) displays a selection menu of LPAR reports.

FCX305	CPU nnnn	SER nnnnn	LPAR act menu	Perf. Monitor
Logical Partition Activity				
S	Command	Description		
-	LPAR	Common Logical Partition state		
-	LSHARACT	The partitions entitlement state		
Command ==> _				
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd F12=Return

Figure 199. Logical Partition Activity Menu Screen (FCX305)

The information shown is based on CP monitor MONITOR domain SAMPLE data.

**Field Descriptions:**

**S**  
Cursor selection field. The field will be highlighted and made cursor-selectable (set as "unprotected") only if the corresponding LPAR activity data screen is available.

**Command**  
The screen name, equal to the command that can be selected for display. The field will be highlighted if the corresponding LPAR activity data screen is available.

**Description**  
Description of the performance display.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX306, Logical Partition Share Screen – LSHARACT

**Performance Toolkit Subcommand:** “LSHARACT” on page 225.

The Logical Partition Share Screen (FCX306) displays a summary report of LPAR weights.

```
FCX306      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS    Perf. Monitor

LPAR Data, Collected in Partition FCFT

Core counts:  CP ZAAP  IFL  ICF  ZIIP
Dedicated      8    0    0    0    0
Shared physical 10    2   16    1    5
Shared logical  79    0   35    0    0 Unused physical core(s) detected

-----
Core  Partition  Core  Load  LPAR
Type Name      Count Max Weight Entlment Cap AbsCap GrpCapNm GrpCap <CoreTotal,%> Core
                                     Busy Excess Conf
... PLB1      ...    ...    0    ...    ...    ...    ...    ...    ...
CP  ECPX3      3    300    10    31.3 No    ...    SAMPLE  1200.0    5.8    .0 o
CP  EEXT1      4    400    10    31.3 No    ...    SAMPLE  1200.0    1.2    .0 o
CP  EPAT       10   1000    10    31.3 No    ...    SAMPLE  1200.0    5.5    .0 o
CP  EPLX1      6    600    60   187.5 No    ...    ...    ...    326.3   138.8 o
CP  EPLX2      8    800    45   140.6 No    ...    ...    ...    256.0   115.4 o
CP  EPLX3      6    600    45   140.6 No    ...    ...    ...    233.3    92.7 o
CP  EPRF1      4    400    DED   400.0 No    ...    ...    ...    399.8    .0 -
CP  EPRF2      4    400    DED   400.0 No    ...    ...    ...    .0    .0 -
CP  ESTL1      7    700    50   156.3 No    ...    ...    ...    1.5    .0 o
CP  EST1      8    800    10    31.3 No    ...    ...    ...    8.7    .0 o
CP  EST2      6    600    10    31.3 No    ...    ...    ...    1.3    .0 o
CP  EVIC       2    200    10    31.3 No    ...    ...    ...    .0    .0 o
CP  FCFT       8    800    40   125.0 No    ...    ...    ...   117.4    .0 o
CP  K4         6    600    10    31.3 No    ...    ...    ...    5.4    .0 o
CP  PHOS       5    500    10    31.3 No    ...    ...    ...    .7    .0 o
IFL EEXT2     16   1600    10   200.0 No    ...    ...    ...    1.1    .0 o
IFL EPLX1      3    300    60  1200.0 No    ...    ...    ...    2.5    .0 u
IFL EST3     16   1600    10   200.0 No    ...    ...    ...    .0    .0 o

Command ==> _
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

Figure 200. Logical Partition Share Screen (FCX306)

The information shown is based on CP monitor PROCESSOR domain EVENT data.

The sample shows a screen sorted by core type (indicated by the underscore characters above the Core Type heading), which is the initial sort sequence. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the SORT subcommand with the string shown in the lowest line of the column heading as sort argument (see “SORT” on page 295) or by positioning the cursor on the header of the field that is to be sorted and then pressing Enter.

### Field Descriptions:

*Core counts:*

This section contains information about physical cores in the CEC that are configured and online, by type.

Counts for physical cores that are dedicated to LPARs and for shared physical cores are displayed separately.

This section also contains, for comparison, a total count of configured shared logical cores for all partitions.

The alarm phrase `Unused physical core(s) detected` is displayed at the end of the shared-logical row when the number of shared logical cores of any type is less than the number of shared physical cores of the corresponding type.

### Core Type

The logical core type. For partitions with no assigned cores, dots are displayed.

**Partition Name**

The name given to the logical partition. It will be highlighted if the partition is active, that is, if at least one logical core is assigned to the partition.

**Core Count**

The number of defined logical cores of the specified type for the current LPAR.

**Load Max**

The potential (or maximum) core load, expressed as a percentage, for cores of marked type for the partition.

For inactive partitions, this field contains dots.

**LPAR Weight**

The weight defined for this core type in the partition. Weight values range from 1 to 999, with a default of 10. The string DED is displayed when the partition is dedicated. A value of 0 means that the current partition was activated with no core.

**Entlment**

The partition's entitlement, expressed as a percentage, where 100% represents one whole physical core's worth of processing capacity. It is the amount of core power a shared partition is guaranteed that PR/SM will let it consume, according to its weight compared to the weights of the other shared partitions, and according to the number of nondedicated physical cores of that type.

**Cap**

The PR/SM initial capping setting for the partition.

**YES**

An initial cap is set. The partition is capped to its entitlement.

**NO**

An initial cap is not set.

**AbsCap**

The absolute capping value PR/SM is enforcing for cores of this type. A value of 100 represents one whole physical core's worth of power.

If the partition does not have an absolute cap, you will see dots in this field.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

**GrpCapNm**

The group capping name assigned to this partition.

If the partition does not have a group cap, you will see dots in this field.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

**GrpCap**

The group capping value assigned to this partition.

If the partition does not have a group cap, you will see dots in this field.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

**CoreTotal Busy**

The partition's actual consumption of core power. A value of 100% indicates consumption of one physical core's worth of power.

**CoreTotal Excess**

The partition's core consumption beyond its entitlement. A value of 100% indicates consumption of one physical core's worth of power.

**Core Conf**

A flag indicating whether the partition's number of logical cores is too few, too many, or correct, with respect to its entitlement. Possible values are:

- o A possible overconfiguration error. The partition has too many logical cores defined for it, compared to its entitlement. This partition's ability to run all of its logical cores to 100% busy depends on other partitions underconsuming their core type entitlements.
- u An underconfiguration error. The partition has too few logical cores defined for it, compared to its entitlement.
- The LPAR configuration for that core type looks correct.
- The LPAR has dedicated cores for the specified core type, or the current partition was activated with no cores.

**Usage note:**

The term *core* describes a unit that contains one or more processors. All processors of a core have the same processor type. For machine types that do not have the multithreading (MT) facility, the machine has physical processors and the partitions have logical processors. For machine types that have the MT facility, the machine has physical cores, each of which contains two physical processors. A partition that has not activated the MT facility continues to have logical processors. A partition that has activated the MT facility has logical cores, each of which contains one or more logical processors.

**Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages are displayed for all load fields.

# FCX307, Logical Partition Logs Menu Screen – LPARLOGM

**Performance Toolkit Subcommand:** “LPARLOGM” on page 223.

The Logical Partition Logs Menu Screen (FCX307) displays a selection menu of LPAR logs by time.




Figure 201. Logical Partition Logs Menu Screen (FCX307)

The information shown is based on CP monitor MONITOR domain SAMPLE data.

**Field Descriptions:**

**S**

Cursor selection field. The field will be highlighted and made cursor-selectable (set as "unprotected") only if the corresponding LPAR log is available.

**Command**

The screen name, equal to the command that can be selected for display. The field will be highlighted if the corresponding LPAR log screen is available.

**Description**

Description of the performance display.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.



# FCX308, CPU Pool Configuration Screen – CPLCONF

**Performance Toolkit Subcommand:** "CPLCONF" on page 154. The CPU Pool Configuration Screen (FCX308):

FCX308	CPU		nnnn	SER	nnnnn	Status	HH:MM:SS	Perf. Monitor		
-----	CPU	.	.	.	Mem-	.	.	.	.	
PoolName	Type	Aff	Limit		bers	<-- List of Members ----->				
LINUXP1	IFL	ON	100 %		13	LINUX001	LINUX002	LINUX003	LINUX004	LINUX012
						LINUX005	LINUX006	LINUX007	LINUX008	LINUX013
						LINUX009	LINUX010	LINUX011		
TESTVMS	IFL	ON	10 %		5	TST1	TST2	TST5	TST4	TST3
FIREWALL	IFL	ON	999.0 CPU		7	SWHITE	GNOME2	GNOME3	GNOME4	GNOME5
						GNOME1	GNOME6			
LINUXP2	IFL	ON	8.0 CPU		0	...				
LINUXP3	IFL	ON	8.0 CPU		0	*deleted*				
Command ==> _										
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return					

Figure 202. Layout of CPU Pool Configuration Screen (FCX308)

**Field Descriptions:**

**PoolName**

The name of the CPU pool.

**CPU Type**

The CPU type of the CPU pool. Valid values are:

**CP**

Central Processor

**IFL**

Integrated Facility for Linux Processor

**Aff**

The CPU affinity setting for the core type limited by this CPU pool. Valid values are: ON (for CPUAffinity ON) or SUP (for CPUAffinity SUPPRESSED).

**Limit**

The CPU pool maximum share from the DEFINE or SET CPUPOOL command. This is the CAPACITY value (with a suffix of "CPU" or "Cores") or the LIMITHARD value (with a suffix of "%").

When multithreading (MT) is not enabled, each core contains only one CPU.

When MT is enabled, each core contains one or more CPUs (or "threads").

**Members**

The current number of users in the CPU pool.

**List of Members**

The list of users that were assigned to the CPU pool. The contents of the "List of Members" column might differ from the count of members, due to the disabling of some sample or event domains in monitor data flow or if some user data was not saved in the trend/summary file (see the description of the FCONX TRENDREC and FCONX SUMREC file contents - the RECORDS, USER, and USERID keywords).

If the number of users listed in this column is less than the number that is specified in the Members column, an ellipsis ("...") appears at the end of the incomplete list.

If the displayed CPU pool is dropped from the system during the selected time interval, the phrase "\*deleted\*" appears in the column.

The information shown is initially provided by a MONITOR domain record when Performance Toolkit connects to the \*MONITOR service. If the program has been set up to collect the monitor data from another machine's disk file, the program will attempt to obtain the monitor domain data from the first monitor sample in the disk file; the data cannot be shown if the required information is not found in that first sample.

This screen updates at each sample interval based on Monitor Processor Sample records when at least one CPU pool is defined. To make sure the List of Members is correct, Monitor User Sample domain should also be enabled. This screen can be displayed or printed only if Monitor Processor Sample data collection is enabled and at least one CPU pool is defined at some point since Performance Toolkit was started.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** All CPU pools, including deleted ones since the last appropriate reset, are shown.

## FCX309, CPU Pool Activity Screen – CPLACT

**Performance Toolkit Subcommand:** [“CPLACT” on page 153](#). The CPU Pool Activity Screen (FCX309):

FCX309	CPU nnnn		SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor
-----	CPU	Mem-				
PoolName	Type	bers	Limit	%CPU	Limit	Count
LINUXP1	IFL	13.0	100 %	3200	4600	
TESTVMS	IFL	5.0	10 %	0.017	0	
FIREWALL	IFL	7.0	999.0 CPU	23.19	212	
LINUXP2	IFL	0.7	8.0 CPU	0.000	0	
LINUXP3	IFL	0.1	8.0 CPU	9.100	0	

Command ==> \_  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 203. Layout of CPU Pool Activity Screen (FCX309)

### Field Descriptions:

#### PoolName

The name of the CPU pool.

#### CPU Type

The CPU type of the CPU pool. Valid values are:

##### CP

Central Processor

##### IFL

Integrated Facility for Linux Processor

#### Members

The average count of members (users) of the CPU pool in the specified interval.

#### Limit

The CPU pool maximum share from the DEFINE or SET CPUPOOL command. This is the CAPACITY value (with a suffix of "CPU" or "Cores") or the LIMITHARD value (with a suffix of "%").

When multithreading (MT) is not enabled, each core contains only one CPU.

When MT is enabled, each core contains one or more CPUs (or "threads").

#### %CPU

The total percentage of time that CPU pool members consume. A value of 100 represents one physical CPU's worth of power.

When SMT is disabled, this value is based on raw CPU time.

When SMT is enabled, this value is based on:

- Raw CPU time, if CP is running z/VM 6.3 without CP APAR VM65680
- Prorated core CPU time, if CP is running z/VM 6.3 with CP APAR VM65680 or is running a CP release after z/VM 6.3.

#### Limit Count

The number of times per minute that the CPU pool has been limited during the selected interval.

This screen can be displayed or printed only if Monitor Processor Sample data collection is enabled and at least one CPU pool with an active member is defined. In other words, the information on this screen is obtained from a Monitor Processor Sample record, which is produced only when at least one CPU pool exists. Because it is measuring CPU pool activity, there needs to be at least one CPU pool with an active member for there to be any activity data and two sample records are needed to calculate the data.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages are shown for all data fields. All CPU pools with the real activity, including deleted ones since the last appropriate reset, are shown. All CPU pool activity data during the selected period contribute to the averages shown. All of the activity fields (rates and percentages) are based on each actual CPU pool activity time within the overall period, not the total elapsed time for the entire period.

# FCX310, PCI Function Menu Screen – PCIMENU

**Performance Toolkit Subcommand:** “PCIMENU” on page 248. The PCI Function Menu Screen (FCX310) shows a selection menu of all available PCI function performance reports.

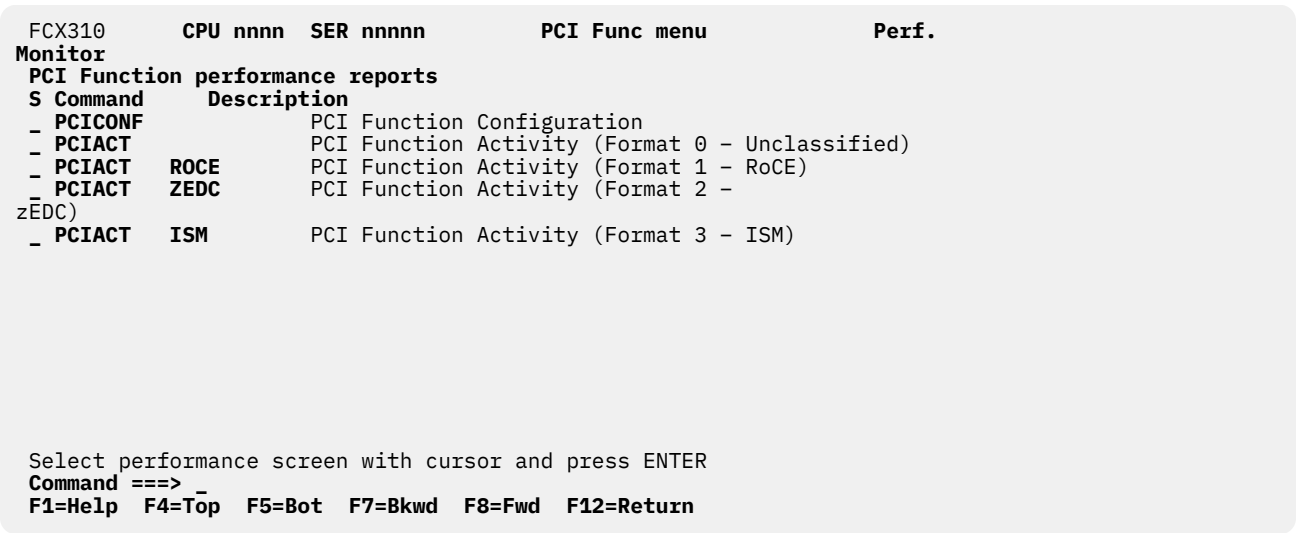


Figure 204. Layout of PCI Function Menu Screen (FCX310)

The information shown is based on CP monitor MONITOR domain and I/O domain PCIF class SAMPLE/ EVENT data.

**Field Descriptions:**

**S**  
Cursor selection field. The field will be highlighted and made cursor-selectable (set as 'unprotected') only if the corresponding PCI Function performance screen is available.

**Command**  
The screen name, equal to the command that can be selected for display. The field will be highlighted if the corresponding PCI Function performance screen is available.

**Description**  
Description of the performance display.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX311, PCI Function Configuration Screen – PCICONF

**Performance Toolkit Subcommand:** [“PCICONF” on page 246](#). The PCI Function Configuration Screen (FCX311):

FCX311	CPU	nnnn	SER	nnnnn	Status	HH:MM:SS	Perf.	Monitor
-----	.	.	.	.	.	.	.	.
						<-DMA parms->		
RPCIF	Owner	VPCIF	PCHID	VFN	State	HT	GT	Size Fmt Type
00000020	.....	.....	017C	1	Disabled	..	..	... 1 10GbE RoCE
00000021	.....	.....	01FC	2	Disabled	..	..	... 1 10GbE RoCE
00000030	.....	.....	0104	10	Disabled	..	..	... 2 zEDC Express
00000031	.....	.....	0184	11	Disabled	..	..	... 2 zEDC Express
00000040	.....	.....	027C	1	Disabled	..	..	... 1 10GbE RoCE
00000041	LINUX009	00000041	027C	2	Enabled	04	04	512M 1 10GbE RoCE
00000045	.....	.....	02FC	1	Disabled	..	..	... 1 10GbE RoCE
00000046	LINUX00A	00000046	02FC	2	Enabled	04	04	512M 1 10GbE RoCE

Command ==>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 205. Layout of PCI Function Configuration Screen (FCX311)

The information shown is based on CP monitor MONITOR domain and I/O domain PCIF class EVENT data. For correct data display, the I/O domain PCIF class EVENT must be enabled.

The sample shows a screen with PCI functions configuration sorted in ascending the real PCI function identifier sequence (indicated by the underscore characters above the RPCIF heading), the initial sequence. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see [“SORT” on page 295](#)) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

### Field Descriptions:

#### RPCIF

Real PCI Function identifier (hexadecimal).

#### Owner

User ID to whom function is attached.

#### VPCIF

Virtual PCI Function identifier (hexadecimal).

#### PCHID

Physical Channel Identifier associated with this PCI Function.

#### VFN

Virtual Function Number associated with this PCI Function.

#### State

Current PCI Function state. Possible states are:

##### Standby

PCI Function not available to the configuration

##### Disabled

PCI Function available for use by the configuration but not enabled

##### Enabled

PCI Function available for use by the configuration and enabled

##### Error

Error detected by z/PCI facility

**Permerr**

Error recovery for PCI Function unsuccessful

**Blocked**

PCI Function blocked

*DMA parameters:*

**DMA parms HT**

Indicator of the type of host translation tables being used (valid when DMA registered).

**DMA parms GT**

Indicator of the type of guest translation tables being used (valid when DMA registered).

**DMA parms Size**

Size of the DMA address space (valid when DMA registered).

**Fmt**

Format number of the measurement block used with this PCIe adapter. It determines the layout of the PCI Function activity report (PCIACT).

**Type**

Symbolic name of this PCIe adapter.

The information shown is initially provided by a MONITOR domain record when PERFKIT connects to the \*MONITOR service. If the program has been set up to collect the monitor data from another machine's disk file, the program will attempt to obtain the monitor domain data from the first monitor sample in the disk file; the data cannot be shown if the required information is not found in that first sample.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** All PCI Functions including deleted ones since the last reset will be shown. In the CURRENT mode, the deleted PCI Functions are never shown.

## FCX312, PCI Function Activity Screen, Format 0 – PCIACT

**Performance Toolkit Subcommand:** [“PCIACT” on page 245](#). The PCI Function Activity Screen, Format 0 (FCX312):

```

FCX312      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS      Perf. Monitor
PCI Function Activity (Format 0 - Unclassified)

----- . . . . . . . . . .

RPCIF      Owner      VPCIF      Pinned  Shadow  Activ  <- Successful operations/s -> <- Byte/s->
000000040  .....  .....  48447K   2 Disabled  1981  2718 14773  4339 3619 3619      0      0
000000041  .....  .....  48224K   2 Disabled  1817  2922 14938  3520 24787 24787      0      0

Command ==> _
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return

```

Figure 206. Layout of PCI Function Activity Screen, Format 0 (FCX312)

The information shown is based on CP monitor MONITOR domain and I/O domain PCIF class SAMPLE/ EVENT data. For correct data display, the I/O domain PCIF class SAMPLE and EVENT must be enabled.

The sample shows a screen with PCI functions activity sorted in ascending the real PCI function identifier sequence (indicated by the underscore characters above the RPCIF heading), the initial sequence. Other possible sort sequences are indicated by dots above the corresponding columns. They can be set using the 'SORT' subcommand with the string shown in the lowest line of the column heading as sort argument (see [“SORT” on page 295](#)) or by positioning the cursor on the header of the field which is to be sorted and then pressing ENTER.

The following field descriptions apply to both the PCIACT and PCILOG displays, with the exception of the "RPCIF" field which will be replaced by the field "Interval End Time" on the PCILOG display.

### Field Descriptions:

## RPCIF

Real PCI Function identifier (hexadecimal).

**Owner**

User ID to whom function is attached.

**VPCIF**

Virtual PCI Function identifier (hexadecimal).

## Pinned

Amount of host storage currently pinned.

## Shadow Tables

Count of shadow tables in use since last DMA register (valid when DMA registered).

## State

Current PCI Function state. Possible states are:

## Standby

PCI Function not available to the configuration

## Disabled

PCI Function available for use by the configuration but not enabled



**Enabled**

PCI Function available for use by the configuration and enabled

**Error**

Error detected by z/PCI facility

**Permerr**

Error recovery for PCI Function unsuccessful

**Blocked**

PCI Function blocked

**Activ Time**

Actual PCI Function activity time within selected interval, in seconds. It is based on the timing measurement provided by the hardware and is not necessarily synchronized with the TOD clock.

*Successful operations/s:*

**Reads**

Count of successful reads from the PCI Function memory or configuration space, per second.

**Writs**

Count of successful writes to the PCI Function memory or configuration space, per second.

**BWrits**

Count of successful block writes to the PCI Function memory or configuration space, per second.

**ATRfs**

Count of successful address-translation refreshes, per second.

**MpRqs**

Count of guest RDMA mapping requests, per second.

**Byte/s Read**

Number of bytes transferred to the PCI Function from main memory, per second. This data is currently not available, so the values will be set to zero.

**Byte/s Write**

Number of bytes transferred from the PCI Function to main memory, per second. This data is currently not available, so the values will be set to zero.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** Overall or interim period averages will be shown for all data fields. All PCI Functions with the real activity, including deleted or disabled ones since the last reset will be shown.

All of a PCIF's activity sessions during the selected period contribute to the averages shown. Note also that all of the activity fields (rates) are based on each actual PCIF's activity time within the overall period, not the total elapsed time for the whole period.

## FCX313, PCI Function Activity Log Screen, Format 0 – PCILOG

**Performance Toolkit Subcommand:** “PCILOG” on page 247. The PCI Function Activity Log Screen, Format 0 (FCX313):

```
FCX313      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor

PCI Function Activity Log for PCIFUNC 00000040 (Format 0 - Unclassified)

Interval    End Time  Owner    VPCIF    Pinned    Shadow    State    Activ    <- Successful operations/s -> <- Byte/s ->
Time                               Tables                               Time Reads Writs BWrits ATRfs MpRqs Read Write
>>Mean>> --- --- 48447K 2 --- 1981 2718 14773 4339 3619 3619 0 0
07:25:28 LINUX009 00000040 48596K 2 Enabled 57.08 52172 40 0 75 75 0 0
07:26:28 LINUX009 00000040 48596K 2 Enabled 60.16 4 32 0 0 0 0 0 0
07:27:28 LINUX009 00000040 48596K 2 Enabled 59.84 4 34 0 0 0 0 0 0
07:28:28 LINUX009 00000040 48596K 2 Enabled 60.16 4 32 0 0 0 0 0 0
07:29:28 LINUX009 00000040 48628K 2 Enabled 59.84 4 8095 2221 1920 1966 0 0
07:30:28 LINUX009 00000040 48660K 2 Enabled 60.16 4 44552 12881 10714 10682 0 0
07:31:28 LINUX009 00000040 48596K 2 Enabled 59.84 4 45721 12935 10742 10774 0 0
07:32:28 LINUX009 00000040 48668K 2 Enabled 60.16 4 42513 11950 9870 9839 0 0
07:33:28 LINUX009 00000040 48692K 2 Enabled 59.84 4 23131 7161 5995 5979 0 0
07:34:28 LINUX009 00000040 48692K 2 Enabled 60.10 4 32 0 0 0 0 0 0
07:35:28 LINUX009 00000040 48692K 2 Enabled 59.90 4 32 0 0 0 0 0 0
07:36:28 LINUX009 00000040 48692K 2 Enabled 60.10 4 32 0 0 0 0 0 0
07:37:28 ..... 48692K 2 Disabled 31.20 4 31 0 0 0 0 0 0
07:38:28 ** No activity for selected interval **
07:39:28 LINUX009 00000040 40488K 2 Enabled 29.76 80567 23 0 124 124 0 0
07:40:28 LINUX009 00000040 48596K 2 Enabled 60.10 6 44 0 10 10 0 0
07:41:28 LINUX009 00000040 48596K 2 Enabled 59.90 4 34 0 0 0 0 0 0
07:42:28 LINUX009 00000040 48628K 2 Enabled 60.10 4 1673 512 405 433 0 0
07:43:28 LINUX009 00000040 48692K 2 Enabled 59.90 4 43212 12977 10794 10821 0 0
07:44:28 LINUX009 00000040 48724K 2 Enabled 60.10 4 47218 13052 10831 10808 0 0
07:45:28 LINUX009 00000040 48692K 2 Enabled 60.16 4 43679 12049 9948 9917 0 0
07:46:28 LINUX009 00000040 48692K 2 Enabled 59.84 4 44302 13220 11035 11068 0 0
07:47:28 LINUX009 00000040 48720K 2 Enabled 60.16 4 44475 13205 11004 10971 0 0
07:48:28 LINUX009 00000040 48660K 2 Enabled 59.84 4 38310 11969 10050 10083 0 0
07:49:28 LINUX009 00000040 48660K 2 Enabled 60.16 4 42366 13509 11316 11284 0 0
07:50:28 LINUX009 00000040 48692K 2 Enabled 59.84 4 17662 5568 4669 4666 0 0
07:51:28 LINUX009 00000040 48692K 2 Enabled 60.16 4 32 0 0 0 0 0 0
07:52:28 LINUX009 00000040 48692K 2 Enabled 59.84 4 32 0 0 0 0 0 0
07:53:28 LINUX009 00000040 48692K 2 Enabled 60.16 4 32 0 0 0 0 0 0
07:54:28 LINUX009 00000040 48692K 2 Enabled 59.84 4 32 0 0 0 0 0 0
07:55:28 LINUX009 00000040 48692K 2 Enabled 60.16 4 32 0 0 0 0 0 0
07:56:28 LINUX009 00000040 48692K 2 Enabled 59.84 4 32 0 0 0 0 0 0
07:57:28 LINUX009 00000040 48692K 2 Enabled 60.10 4 32 0 0 0 0 0 0
07:58:28 LINUX009 00000040 48692K 2 Enabled 59.90 4 32 0 0 0 0 0 0
07:59:28 LINUX009 00000040 48692K 2 Enabled 60.10 4 32 0 0 0 0 0 0
08:00:28 ..... 48692K 2 Disabled 2.211 2 18 0 0 0 0 0 0
08:04:28 ** No activity for selected interval **

Command ===>
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return
```

Figure 207. Layout of PCI Function Activity Log Screen, Format 0 (FCX313)

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval for which the values were calculated.

In the case of the ">>Mean>>" entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'RPCIF' field, all data columns are identical to the ones of the 'PCI Function Activity Screen (PCIACT)'. Please see [“FCX312, PCI Function Activity Screen, Format 0 – PCIACT”](#) on page 840 for their description.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX314, Pinned Storage Log Screen – PINLOG

**Performance Toolkit Subcommand:** "PINLOG" on page 250. The Pinned Storage Log Screen (FCX314):

FCX314	CPU	nnnn	SER	nnnnn	Interval	HH:MM:SS	-	HH:MM:SS	Perf. Monitor
<b>&lt;----- Pinned Storage Statistics -----&gt;</b>									
<b>Interval</b>	<b>&lt;-Config--&gt;</b>	<b>&lt;--Class 0--&gt;</b>	<b>&lt;--Class 1--&gt;</b>	<b>&lt;-Pin requests/s--&gt;</b>					
<b>End Time</b>	<b>Warn%</b>	<b>Fail%</b>	<b>&lt;2GB</b>	<b>&gt;2GB</b>	<b>&lt;2GB</b>	<b>&gt;2GB</b>	<b>Warnings</b>	<b>Failures</b>	
<b>&gt;&gt;Mean&gt;&gt;</b>	---	---	32K	55702K	0	28K	.000	.000	
07:10:28	50	80	32K	1216K	0	24K	.000	.000	
07:15:28	50	80	32K	1216K	0	24K	.000	.000	
07:20:28	50	80	32K	1216K	0	24K	.000	.000	
07:25:28	50	80	32K	10974K	0	26K	.000	.000	
07:30:28	50	80	32K	97257K	0	26K	.000	.000	
07:35:28	50	80	32K	99010K	0	32K	.000	.000	
07:40:28	50	80	32K	46837K	0	24K	.000	.000	
07:45:28	50	80	32K	99034K	0	34K	.000	.000	
07:50:28	50	80	32K	99105K	0	40K	.000	.000	
07:55:28	50	80	32K	99080K	0	24K	.000	.000	
08:00:28	50	80	32K	79518K	0	24K	.000	.000	
08:05:28	50	80	32K	1272K	0	26K	.000	.000	
<b>Command ===&gt;</b>									
<b>F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return</b>									

Figure 208. Layout of Pinned Storage Log Screen (FCX314)

The information shown is based on CP monitor SYSTEM domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the ">>Mean>>" entry, you can find the relevant interval in the header line.

#### Config Warn%

Percent of the total pinned page count compared to the DPA that causes a warning message when a pin request is made.

If this parameter is not defined in the SYSTEM CONFIG file, dashes are displayed on the regular data lines. On the 'Mean' line, dashes are always displayed.

#### Config Fail%

Percent of the total pinned page count compared to the DPA that causes a pin request to fail.

If this parameter is not defined in the SYSTEM CONFIG file, dashes are displayed on the regular data lines. On the 'Mean' line, dashes are always displayed.

#### Pinned Storage Statistics:

##### Class 0 <2GB

Size of pinned storage class 0 below 2G.

##### Class 0 >2GB

Size of pinned storage class 0 above 2G.

##### Class 1 <2GB

Size of pinned storage class 1 below 2G.

##### Class 1 >2GB

Size of pinned storage class 1 above 2G.

#### Pin requests/s Warnings

Count of the times the warning percent was exceeded when a pin request was made, per second.

**Pin requests/s Failures**

Count of the times a pin request failed because the fail percent was exceeded, per second.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

## FCX315, Multiprocessor User Wait States Log Screen – USTMPLOG

**Performance Toolkit Subcommand:** “USTMPLOG” on page 354.

The Multiprocessor User Wait States Log Screen (FCX315) shows a "by time" log of the selected user's virtual CPU wait state statistics. Note that the specified user must have been previously selected for benchmarking.

FCX315	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS										Perf. Monitor										
Multiprocessor User Wait States Log for User TCP50																							
Interval	<vCPU->															<-%Samples showing->					Nr of		
End Time	ID	Type	%ACT	%RUN	%CPU	%LDG	%PGW	%IOW	%SIM	%TIW	%CFW	%IOA	%PGA	%LIM	%OTH	Q0	Q1	Q2	Q3	E0-3	Users		
>>Mean>>	00	CP	6	0	1	0	0	0	0	99	0	0	0	0	0	100	0	0	0	0	0		
>>Mean>>	01	CP	0	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...		
>>Mean>>	02	CP	0	0	100	0	0	0	0	0	0	0	0	0	0	100	0	0	0	0	0		
>>Mean>>	03	CP	0	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...		
>>Mean>>	04	CP	0	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...		
>>Mean>>	05	CP	0	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...		
18:25:38	00	CP	0	0	1	0	0	0	0	99	0	0	0	0	0	100	0	0	0	0	0		
18:25:38	01	CP	0	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...		
18:25:38	02	CP	0	0	100	0	0	0	0	0	0	0	0	0	0	100	0	0	0	0	0		
18:25:38	03	CP	0	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...		
18:25:38	04	CP	0	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...		
18:25:38	05	CP	0	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...		
18:26:38	00	CP	0	0	1	0	0	0	0	99	0	0	0	0	0	100	0	0	0	0	0		
18:26:38	01	CP	0	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...		
18:26:38	02	CP	0	0	100	0	0	0	0	0	0	0	0	0	0	100	0	0	0	0	0		
18:26:38	03	CP	0	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...		
18:26:38	04	CP	0	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...		
18:26:38	05	CP	0	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...		
Command ==> _																							
F1=Help			F4=Top			F5=Bot			F7=Bkwd			F8=Fwd			F12=Return								

Command ==> \_  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 209. Layout of Multiprocessor User Wait States Log Screen (FCX315)

Each line in this display contains data for each virtual processor of the benchmarked user ID, not for the entire user ID.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### vCPU ID

Identifies the virtual processor address, in hexadecimal.

#### vCPU Type

Indicates the type of virtual processor:

##### CP

Central Processor

##### ICF

Internal Coupling Facility

##### IFL

Integrated Facility for Linux

##### ZIIP

IBM Z Integrated Information Processor

##### MIX

The MP user has different PU types.

This field contains a value at the end of the interval.

**%ACT**

The percentage of samples during which the user was found to be active, consuming CPU or waiting for some service.

**%RUN**

The percentage of active samples during which the user was found to be running on a logical processor.

**%CPU**

The percentage of active samples during which the user was found to be waiting to run on the logical processor.

**%LDG**

The percentage of active samples during which the user was found in loading state, that is, with a very high page read rate.

**%PGW**

The percentage of active samples during which the user was found in page wait state.

**%IOW**

The percentage of active samples during which the user was found in I/O wait state.

**%SIM**

The percentage of active samples during which the user was found in instruction simulation wait, that is, waiting for CP to complete simulation of a privileged instruction.

**%TIW**

The percentage of active samples during which the user was found in test idle wait, that is, waiting to be dropped from the dispatch list.

**%CFW**

The percentage of active samples during which the user was found in console function wait, that is, waiting for CP to execute one of the console function commands which allow a user to alter the state and configuration of the virtual machine (ATTACH, DETACH, or LINK, for example).

**%IOA**

The percentage of active samples during which the user was found waiting for something other than the above while an I/O operation started by this machine had not yet completed.

**%PGA**

The percentage of page active wait, that is, the virtual machine has loaded a wait PSW while waiting for the completion of a page read operation. This state is to be expected for guest systems running with PAGEX ON.

**%LIM**

The percentage of active samples during which the user was found in the limit list, that is, in the dispatch list waiting for CPU but prevented from running by a maximum share setting.

**%OTH**

The percentage of active samples during which the user was found in other states.

**%Samples showing:**

The columns under this header show the relative frequency with which a machine was found in the different lists and classes:

**Q0**

The percentage of active samples during which the user was found in Q0 (in the dispatch list as a class 0 user).

**Q1**

The percentage of active samples during which the user was found in Q1 (in the dispatch list as a class 1 user).

**Q2**

The percentage of active samples during which the user was found in Q2 (in the dispatch list as a class 2 user).

**Q3**

The percentage of active samples during which the user was found in Q3 (in the dispatch list as a class 3 user).

**E0-3**

The percentage of active samples during which the user was found in the eligible list (any class).

**Nr of Users**

The number of users that contributed to the average shown on this line (values are inserted only for lines with averages; this field is blank for user data lines).

**Effect of AVERAGE and INTERIM subcommands:**

Average values will be shown for all fields. All of a user's sessions during the selected period contribute to the averages shown.

## FCX316, Subpool Storage Log Screen – SUBPLOG

**Performance Toolkit Subcommand:** “SUBPLOG” on page 305. The Subpool Storage Log Screen (FCX316):

FCX316	CPU	nnnn	SER	nnnnn	Interval HH:MM:SS - HH:MM:SS			Perf.	Monitor
<-- Subpool Storage Usage -->									
Interval	Subpool	<--Config-->			<--Fails/s-->				
End Time	Name	Size	Warn%	Current	Warns/s	Subpool	All		
>>Mean>>	IOAT	2M	---	36K	.000	.000	.000		
07:10:28	IOAT	2M	80	0	.000	.000	.000		
07:15:28	IOAT	2M	80	0	.000	.000	.000		
07:20:28	IOAT	2M	80	0	.000	.000	.000		
07:25:28	IOAT	2M	80	6K	.000	.000	.000		
07:30:28	IOAT	2M	80	64K	.000	.000	.000		
07:35:28	IOAT	2M	80	64K	.000	.000	.000		
07:40:28	IOAT	2M	80	32K	.000	.000	.000		
07:45:28	IOAT	2M	80	64K	.000	.000	.000		
07:50:28	IOAT	2M	80	64K	.000	.000	.000		
07:55:28	IOAT	2M	80	64K	.000	.000	.000		
08:00:28	IOAT	2M	80	51K	.000	.000	.000		
08:05:28	IOAT	2M	80	0	.000	.000	.000		
Command ==> _									
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return				

Figure 210. Layout of Subpool Storage Log Screen (FCX316)

The information shown is based on CP monitor SYSTEM domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval or the BYTIME interval for which the values were calculated.

In the case of the ">>Mean>>" entry, you can find the relevant interval in the header line.

#### Subpool Name

The name of storage subpool.

#### Config Size

The size of storage subpool.

#### Config Warn%

Percent of the subpool current usage compared to its size when a warning message is issued. If this parameter is not defined in the SYSTEM CONFIG file, dashes are displayed on the regular data lines. On the 'Mean' line, dashes are always displayed.

#### Subpool Storage Usage:

#### Current

The size of currently used subpool storage.

#### Warns/s

Count of the times a subpool warning message was issued, per second.

#### Fails/s Subpool

Count of the times subpool memory was not available, per second.

#### Fails/s All

Count of the times memory was not available from the subpool or from the available list, per second.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.



## FCX317, Global Networking Object Activity Screen – GLONACT

**Performance Toolkit Subcommand:** “GLONACT” on page 181. The Global Networking Object Activity Screen (FCX317):

```
FCX317      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor
Total number of global networking objects:  3

-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
          <-Outbound/s> <-Inbound/s->
          Bytes Packet Bytes Packet
LanName  SysID  PGrName  Inst  T_Byte T_Pack  R_Byte R_Pack
GLOB1    2NDKARLA  PGRP1    0  107.08 .635   .000   .000
GLOB2    2NDKARLA  PGRP1    1   .000   .000   .000   .000
GLOB3    2NDKARLA  . . . . .  0   .000   .000   .000   .000

See GVSWITCH / VSWITCH / EVSWITCH for other Vswitch reports
Command ==> _
F1=Help  F4=Top   F5=Bot   F7=Bkwd  F8=Fwd   F12=Return
```

Figure 211. Layout of Global Networking Object Activity Screen (FCX317)

The information shown is based on CP monitor VIRTUAL NETWORK domain data.

### Field Descriptions:

#### LanName

The global object name. This field contains the name of a virtual switch that is defined with the GLOBAL attribute.

#### SysID

The system name.

#### PGrName

The name of the link aggregation port group in use for this virtual switch. This field contains the name of a shared port group if one is in use on the global virtual switch. Otherwise, this field contains blanks.

#### Inst

The shared port group instance (0-3) if a shared port group is in use on the global virtual switch. Otherwise, this field contains a 0.

#### Outbound/s Bytes T\_Byte

The cumulative number of bytes customer data sent over the IVL network on behalf of this global object.

#### Outbound/s Packet T\_Pack

The cumulative number of packets of customer data sent over the IVL network on behalf of this global object.

#### Inbound/s Bytes R\_Byte

The cumulative number of bytes of customer data received over the IVL network on behalf of this global object.

#### Inbound/s Packet R\_Pack

The cumulative number of packets of customer data received over the IVL network on behalf of this global object.

### Effect of AVERAGE and INTERIM subcommands:

Overall or interim period averages will be shown for all data fields. All global networking objects with the real activity, including detached or disconnected ones since the last reset, will be shown.

All of a global networking objects activity sessions during the selected period contribute to the averages shown. Note also that all of the activity fields (rates) are based on each actual IVL's network activity time for a global networking object within the overall period, not the total elapsed time for the whole period.

# FCX318, PCI Function Activity Screen, Format 1 – PCIACT ROCE

**Performance Toolkit Subcommand:** “[PCIACT](#)” on page 245. The PCI Function Activity Screen, Format 1 (FCX318):

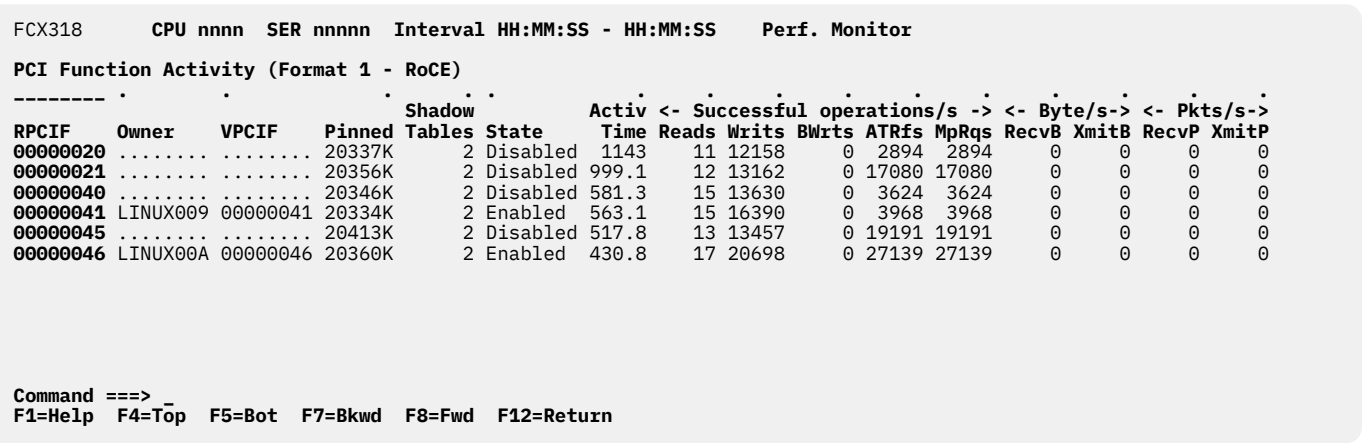


Figure 212. Layout of PCI Function Activity Screen, Format 1 (FCX318)

All data columns are identical to those for “[FCX312, PCI Function Activity Screen, Format 0 – PCIACT](#)” on page 840, except for the following, which are specific to Format 1:

## Format 1-Specific Field Descriptions:

### Byte/s RecvB

Number of bytes received on the external Ethernet interface, per second.

### Byte/s XmitB

Number of bytes transmitted on the external Ethernet interface, per second.

### Pkts/s RecvP

Count of packets received on the external Ethernet interface, per second.

### Pkts/s XmitP

Count of packets transmitted on the external Ethernet interface, per second.

## FCX319, PCI Function Activity Log Screen, Format 1 – PCILOG

Performance Toolkit Subcommand: “PCILOG” on page 247.

The PCI Function Activity Log Screen, Format 1 (FCX319):

```
FCX319      CPU nnnn  SER nnnnn Interval HH:MM:SS - HH:MM:SS    Perf. Monitor

PCI Function Activity Log for PCIFUNC 00000040 (Format 1 - RoCE)

Interval    End Time Owner    VPCIF    Pinned Shadow Tables State    Activ Time    <- Successful operations/s -> <- Byte/s-> <- Pkts/s->
>>Mean>> ---      ---      20337K    2 ---      1143    11 12158    0 2894 2894    0 0 0 0
05:32:25 LINUX009 00000020 20296K    2 Enabled 73.27    57 20 0 32 32    0 0 0 0
05:34:25 LINUX009 00000020 20296K    2 Enabled 119.7    8 16 0 0 0    0 0 0 0
05:36:25 LINUX009 00000020 20296K    2 Enabled 119.7    8 16 0 0 0    0 0 0 0
05:38:25 LINUX009 00000020 20328K    2 Enabled 119.7    8 10853 0 3223 3237 0 0 0 0
05:40:25 LINUX009 00000020 20392K    2 Enabled 119.7    8 32494 0 7463 7662 0 0 0 0
05:42:25 LINUX009 00000020 20360K    2 Enabled 123.9    8 26754 0 6350 6187 0 0 0 0
05:44:25 LINUX009 00000020 20424K    2 Enabled 119.7    8 32667 0 7373 7234 0 0 0 0
05:46:25 LINUX009 00000020 20328K    2 Enabled 119.7    8 12260 0 2966 3061 0 0 0 0
05:48:25 LINUX009 00000020 20328K    2 Enabled 119.7    8 16 0 0 0 0 0 0 0
05:50:25 ..... 20328K    2 Disabled 107.4    8 16 0 0 0 0 0 0 0
06:14:25 ** No activity for selected interval **

Command ==> _
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return
```

Figure 213. Layout of PCI Function Activity Log Screen, Format 1 (FCX319)

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval for which the values were calculated.

In the case of the ">>Mean>>" entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'RPCIF' field, all data columns are identical to those for “FCX318, PCI Function Activity Screen, Format 1 – PCIACT ROCE” on page 851.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX320, PCI Function Activity Screen, Format 2 – PCIACT ZEDC

**Performance Toolkit Subcommand:** “[PCIACT](#)” on page 245. The PCI Function Activity Screen, Format 2 (FCX320):

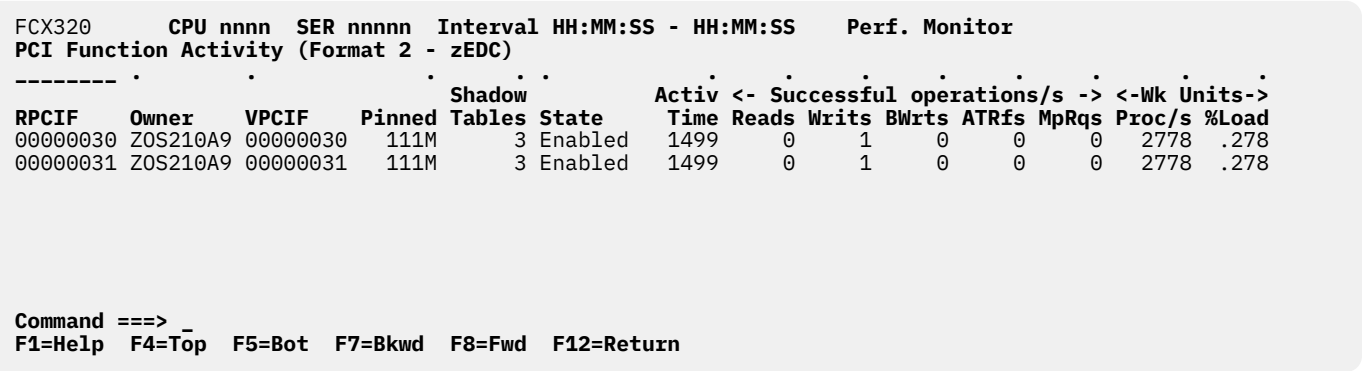


Figure 214. Layout of PCI Function Activity Screen, Format 2 (FCX320)

All data columns are identical to those for “[FCX312, PCI Function Activity Screen, Format 0 – PCIACT](#)” on page 840, except for the following, which are specific to Format 2:

**Format 2-Specific Field Descriptions:**

- Wk Units Proc/s**  
Count of work units processed by the PCI function, per second.
- Wk Units %Load**  
PCI function load percentage during the observed interval.

## FCX321, PCI Function Activity Log Screen, Format 2 – PCILOG

Performance Toolkit Subcommand: “PCILOG” on page 247.

The PCI Function Activity Log Screen, Format 2 (FCX321):

FCX321 CPU nnnn SER nnnnn Interval HH:MM:SS - HH:MM:SS Perf. Monitor													
PCI Function Activity Log for PCIFUNC 00000030 (Format 2 - zEDC)													
Interval	End Time	Owner	VPCIF	Pinned	Shadow Tables	State	Activ Time	Reads	Writes	BWrts	ATRfs	MpRqs	<-Wk Units-> Proc/s %Load
>>Mean>>	---	---	---	111M	3	---	1499	0	1	0	0	0	2778 .278
10:26:30	ZOS210A9	00000030	111M	3	Enabled	57.80	0	1	0	0	0	0	2655 .266
10:27:30	ZOS210A9	00000030	111M	3	Enabled	61.93	0	1	0	0	0	0	2383 .238
10:28:30	ZOS210A9	00000030	111M	3	Enabled	57.80	0	0	0	0	0	0	2614 .261
10:29:30	ZOS210A9	00000030	111M	3	Enabled	61.93	0	1	0	0	0	0	3253 .325
10:30:30	ZOS210A9	00000030	111M	3	Enabled	57.80	0	1	0	0	0	0	2822 .282
10:31:30	ZOS210A9	00000030	111M	3	Enabled	61.93	0	0	0	0	0	0	2448 .245
10:32:30	ZOS210A9	00000030	111M	3	Enabled	57.80	0	1	0	0	0	0	2305 .231
10:33:30	ZOS210A9	00000030	111M	3	Enabled	61.93	0	0	0	0	0	0	2500 .250
10:34:30	ZOS210A9	00000030	111M	3	Enabled	61.93	0	1	0	0	0	0	3145 .315
10:35:30	ZOS210A9	00000030	111M	3	Enabled	57.80	0	0	0	0	0	0	2720 .272
10:36:30	ZOS210A9	00000030	111M	3	Enabled	61.93	0	0	0	0	0	0	2133 .213
10:37:30	ZOS210A9	00000030	111M	3	Enabled	57.80	0	1	0	0	0	0	2696 .270
10:38:30	ZOS210A9	00000030	111M	3	Enabled	61.93	0	1	0	0	0	0	2645 .265
10:39:30	ZOS210A9	00000030	111M	3	Enabled	57.80	0	1	0	0	0	0	2756 .276
10:40:30	ZOS210A9	00000030	111M	3	Enabled	61.93	0	1	0	0	0	0	2551 .255
10:41:30	ZOS210A9	00000030	111M	3	Enabled	57.80	0	1	0	0	0	0	2756 .276
10:42:30	ZOS210A9	00000030	111M	3	Enabled	61.93	0	1	0	0	0	0	3010 .301
10:43:30	ZOS210A9	00000030	111M	3	Enabled	57.80	0	1	0	0	0	0	3129 .313
10:44:30	ZOS210A9	00000030	111M	3	Enabled	61.93	0	1	0	0	0	0	2942 .294
10:45:30	ZOS210A9	00000030	111M	3	Enabled	57.80	0	1	0	0	0	0	2793 .279
10:46:30	ZOS210A9	00000030	111M	3	Enabled	61.93	0	0	0	0	0	0	2555 .256
10:47:30	ZOS210A9	00000030	111M	3	Enabled	57.80	0	1	0	0	0	0	3442 .344
10:48:30	ZOS210A9	00000030	111M	3	Enabled	61.93	0	1	0	0	0	0	3324 .332
10:49:30	ZOS210A9	00000030	111M	3	Enabled	57.80	0	1	0	0	0	0	3166 .317
10:50:30	ZOS210A9	00000030	111M	3	Enabled	61.93	0	1	0	0	0	0	2736 .274

Command ==> \_  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 215. Layout of PCI Function Activity Log Screen, Format 2 (FCX321)

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval for which the values were calculated.

In the case of the ">>Mean>>" entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'RPCIF' field, all data columns are identical to those for “FCX320, PCI Function Activity Screen, Format 2 – PCIACT ZEDC” on page 853.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX322, PCI Function Activity Screen, Format 3 – PCIACT ISM

**Performance Toolkit Subcommand:** “PCIACT” on page 245. The PCI Function Activity Screen, Format 3 (FCX322):

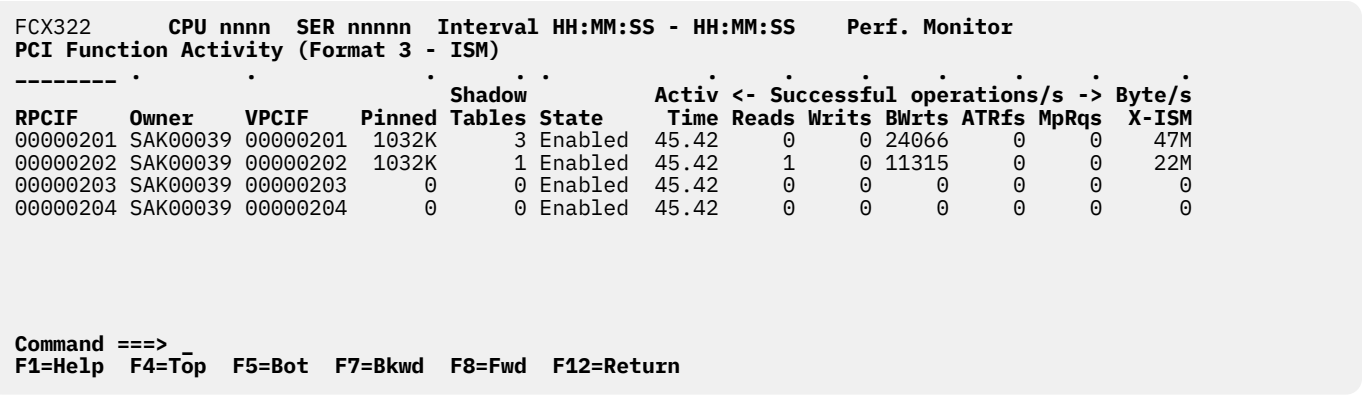


Figure 216. Layout of PCI Function Activity Screen, Format 3 (FCX322)

All data columns are identical to those for “FCX312, PCI Function Activity Screen, Format 0 – PCIACT” on page 840, except for the following, which are specific to Format 3:

**Format 3-Specific Field Descriptions:**

**Byte/s X-ISM**

Number of bytes transmitted via ISM PCI function, per second.

# FCX323, PCI Function Activity Log Screen, Format 3 – PCILOG

**Performance Toolkit Subcommand:** [“PCILOG” on page 247](#). The PCI Function Activity Log Screen, Format 3 (FCX323):

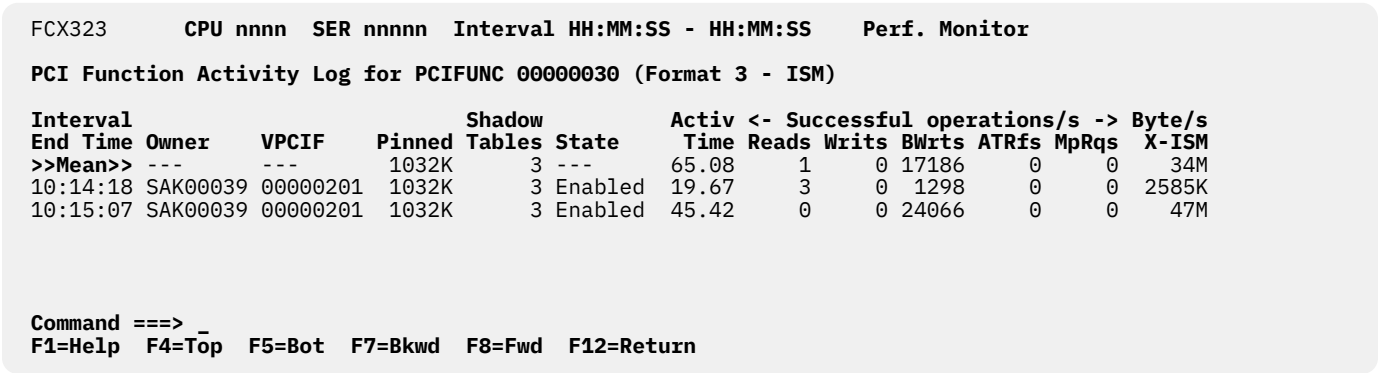


Figure 217. Layout of PCI Function Activity Log Screen, Format 3 (FCX323)

The first entry is reserved for averages since the last RESET subcommand, or since data collection was activated. It is designated by the string >>Mean>> in the 'Interval End Time' column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval for which the values were calculated.

In the case of the ">>Mean>>" entry, you can find the relevant interval in the header line.

Except for the 'Interval End Time' field which replaces the 'RPCIF' field, all data columns are identical to those for [“FCX322, PCI Function Activity Screen, Format 3 – PCIACT ISM” on page 855](#).

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.



# FCX324, CPU Pool Menu Screen – CPLMENU

**Performance Toolkit Subcommand:** “CPLMENU” on page 155. The CPU Pool Menu Screen (FCX324) shows a selection menu of all available CPU pool performance reports.

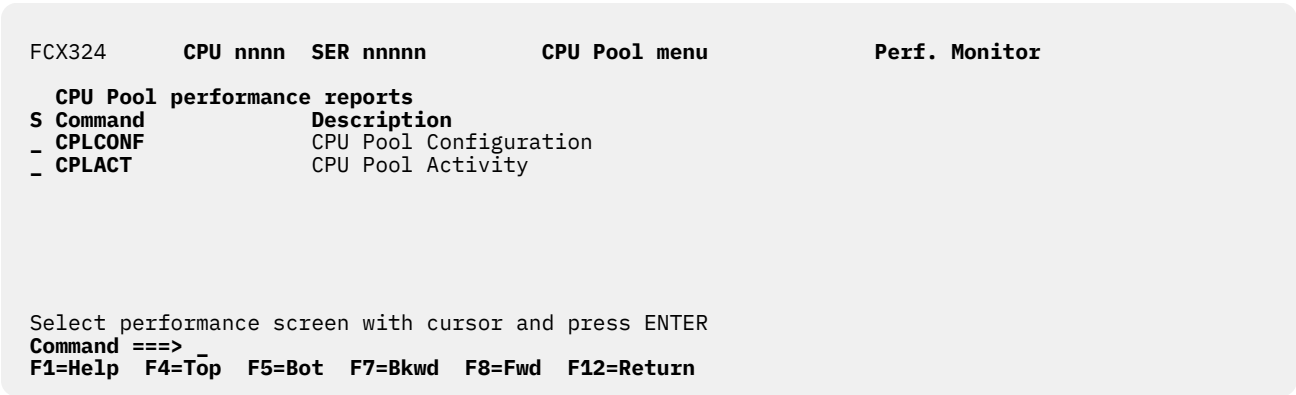


Figure 218. Layout of CPU Pool Menu Screen (FCX324)

**Field Descriptions:**

**S** Cursor selection field. The field will be highlighted and made cursor-selectable (set as 'unprotected') only if the corresponding CPU pool performance screen is available.

**Command** The screen name, equal to the command that can be selected for display. The field will be highlighted if the corresponding CPU pool performance screen is available.

**Description** Description of the performance display.

**Effect of 'AVERAGE' and 'INTERIM' subcommands:** None.

# FCX325, CPU Activity Selection Menu Screen – CPUMENU

**Performance Toolkit Subcommand:** “CPUMENU” on page 163. The CPU Activity Selection Menu Screen (FCX325) shows a selection menu of all available CPU activity reports.

FCX325            CPU nnnn   SER nnnnn            CPU data menu

CPU activity reports	
S Command	Description
— CPU	CPU Load and Transactions
— DSVBKACT	Dispatch Vector Activity
— LOCKACT	Spin Lock Collision Activity
— MTCORELG	SMT Metrics Per Core Log
— MTCTYPLG	SMT Metrics Per Core Type Log

Select performance screen with cursor and hit ENTER  
Command ==>  
F4=Top   F5=Bot   F7=Bkwd   F8=Fwd   F12=Return   F14=Help

Figure 219. Layout of CPU Activity Selection Menu Screen (FCX325)

**Field Descriptions:**

**S**  
Indicates the cursor selection field. This field is highlighted and made cursor-selectable (set as unprotected) only if the corresponding CPU performance screen is available.

**Command**  
The screen name, which is equal to the command that can be selected for display. This field is highlighted if the corresponding CPU performance screen is available.

**Description**  
A description of the CPU performance screen.

**Usage Note:**

The MTCORELG and MTCTYPLG reports are selectable only when multithreading is enabled:

1. For the monitored system when monitoring live data
2. On the system that generated a MONWRITE file that Performance Toolkit for z/VM is processing.

## FCX326, Spin Lock Collision Activity Screen – LOCKACT

**Performance Toolkit Subcommand:** “LOCKACT” on page 218. The Spin Lock Collision Activity Screen (FCX326) shows spin lock information by device.

FCX326	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor								
	<----- Combined ----->				<----- Exclusive ----->				<----- Shared ----->				
LockName	CCol/s	CAvSpn	C%Busy	CCAD/s	ECol/s	EAvSpn	E%Busy	ECAD/s	SCol/s	SAvSpn	S%Busy	SCAD/s	
>>Total>	59.867	4.712	.028	.000	56.833	4.339	.025	.000	3.033	11.692	.004	.000	
SRMSLOCK	25.367	5.984	.015	.000	22.333	5.207	.012	.000	3.033	11.692	.004	.000	
DSV_0000	5.417	13.255	.007	.000	5.417	13.255	.007	.000	.000	...	.000	.000	
DSV_0001	4.067	3.451	.001	.000	4.067	3.451	.001	.000	.000	...	.000	.000	
DSV_0002	4.400	3.080	.001	.000	4.400	3.080	.001	.000	.000	...	.000	.000	
SRMATDLK	4.867	1.096	.001	.000	4.867	1.096	.001	.000	.000	...	.000	.000	
HCPTRQLK	7.283	.606	.000	.000	7.283	.606	.000	.000	.000	...	.000	.000	
DSV_0017	1.600	2.375	.000	.000	1.600	2.375	.000	.000	.000	...	.000	.000	
DSV_0009	1.000	2.050	.000	.000	1.000	2.050	.000	.000	.000	...	.000	.000	
DSV_000D	.883	2.302	.000	.000	.883	2.302	.000	.000	.000	...	.000	.000	
DSV_0015	.833	2.440	.000	.000	.833	2.440	.000	.000	.000	...	.000	.000	
DSV_0016	.683	2.927	.000	.000	.683	2.927	.000	.000	.000	...	.000	.000	
DSV_000B	.650	2.897	.000	.000	.650	2.897	.000	.000	.000	...	.000	.000	
DSV_000C	.567	2.588	.000	.000	.567	2.588	.000	.000	.000	...	.000	.000	
DSV_000F	.583	2.371	.000	.000	.583	2.371	.000	.000	.000	...	.000	.000	
DSV_0008	.417	2.640	.000	.000	.417	2.640	.000	.000	.000	...	.000	.000	
DSV_0010	.367	2.591	.000	.000	.367	2.591	.000	.000	.000	...	.000	.000	
DSV_000A	.317	1.737	.000	.000	.317	1.737	.000	.000	.000	...	.000	.000	
DSV_0011	.183	2.727	.000	.000	.183	2.727	.000	.000	.000	...	.000	.000	
DSV_0012	.117	3.857	.000	.000	.117	3.857	.000	.000	.000	...	.000	.000	
DSV_0013	.117	3.857	.000	.000	.117	3.857	.000	.000	.000	...	.000	.000	
DSV_000E	.150	2.111	.000	.000	.150	2.111	.000	.000	.000	...	.000	.000	
AVZA0001	.000	...	.000	.000	.000	...	.000	.000	.000	...	.000	.000	
AVZA0002	.000	...	.000	.000	.000	...	.000	.000	.000	...	.000	.000	
AVZA0003	.000	...	.000	.000	.000	...	.000	.000	.000	...	.000	.000	
AVZA0004	.000	...	.000	.000	.000	...	.000	.000	.000	...	.000	.000	
AVZA0005	.000	...	.000	.000	.000	...	.000	.000	.000	...	.000	.000	
AVZA0006	.000	...	.000	.000	.000	...	.000	.000	.000	...	.000	.000	
AVZA0007	.000	...	.000	.000	.000	...	.000	.000	.000	...	.000	.000	
AVZA0008	.000	...	.000	.000	.000	...	.000	.000	.000	...	.000	.000	
AVZA0009	.000	...	.000	.000	.000	...	.000	.000	.000	...	.000	.000	
AVZB0000	.000	...	.000	.000	.000	...	.000	.000	.000	...	.000	.000	
AVZB0001	.000	...	.000	.000	.000	...	.000	.000	.000	...	.000	.000	

Command ==>  
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 220. Layout of Spin Lock Collision Activity Screen (FCX326)

### Field Descriptions:

#### LockName

The name of the CP spin lock.

#### Combined CCol/s

The combined spin lock collisions per second.

#### Combined CAvSpn

The average number of microseconds spent spinning per collision, over all collisions.

#### Combined C%Busy

The CPU power spent spinning in all collisions, where a value of 100% means one physical processor's worth of CPU power.

#### Combined CCAD/s

The total number of compare and delay facility (CAD) instructions executed when trying to obtain the lock in shared or exclusive mode, per second.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

When running Performance Toolkit against data from z/VM 6.3 without APAR VM65586 installed, this column contains .0 values.

**Exclusive ECol/s**

The exclusive-mode spin lock collisions per second.

**Exclusive EAvSpn**

The average number of microseconds spent spinning per exclusive-mode collision.

**Exclusive E%Busy**

The CPU power spent spinning in exclusive-mode collisions, where a value of 100% means one physical processor's worth of CPU power.

**Exclusive ECAD/s**

The total number of CAD instructions executed when trying to obtain the lock in exclusive mode, per second.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

When running Performance Toolkit against data from z/VM 6.3 without APAR VM65586 installed, this column contains .0 values.

**Shared SCol/s**

The shared-mode spin lock collisions per second.

**Shared SAVSpn**

The average number of microseconds spent spinning per shared-mode collision.

**Shared S%Busy**

The CPU power spent spinning in shared-mode collisions, where a value of 100% means one physical processor's worth of CPU power.

**Shared SCAD/s**

The total number of CAD instructions executed when trying to obtain the lock in shared mode, per second.

When running Performance Toolkit against z/VM 6.2 or earlier data, you will see dashes in this field.

When running Performance Toolkit against data from z/VM 6.3 without APAR VM65586 installed, this column contains .0 values.

**Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages are displayed for all data fields.

## FCX327, HyperPAV Alias Activity Screen – HPALIAS

**Performance Toolkit Subcommand:** “HPALIAS” on page 196.

The HyperPAV Alias Activity Screen (FCX327) shows pooling information for HyperPAV alias pools.

FCX327	CPU	nnnn	SER	nnnnn	Interval	HH:MM:SS	-	HH:MM:SS	Perf. Monitor	
----	.	<-----Counts----->			<---Rate/s--->		<-----Available Alias			Extremes----->
SSID	Pool	Base	Alias	SYSTEM	Tries	Fails	Low	WhenLow	High	WhenHigh
0100	12	66	48	48	0	0	...	.....	...	.....
035A	10	25	1	1	0	0	...	.....	...	.....
0401	4	50	156	156	60.8	0	137	2017/10/24,12:30:07	156	2017/10/24,12:30:07
0402	0	54	152	152	281.9	0	118	2017/10/24,12:29:07	152	2017/10/24,12:30:07
062B	2	70	156	156	7116.7	0	55	2017/10/24,12:30:07	153	2017/10/24,12:31:07
0633	3	40	156	156	292.9	0	52	2017/10/24,12:30:07	156	2017/10/24,12:30:07
Command ==> _										
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return					

Figure 221. Layout of HyperPAV Alias Activity Screen (FCX327)

### Field Descriptions:

#### SSID

The subsystem ID of the logical control unit (LCU).

#### Pool

The pool number of the HyperPAV pool.

#### Counts:

#### Base

The number of base devices in the LCU associated with this HyperPAV pool.

#### Alias

The number of HyperPAV alias devices in the HyperPAV pool.

#### SYSTEM

The number of HyperPAV alias devices attached to SYSTEM.

#### Rate/s:

#### Tries

The rate, per second, at which CP is trying to get aliases from the pool.

#### Fails

The rate, per second, at which CP is failing to get aliases from the pool.

#### Available Alias Extremes:

#### Low

The lowest number of available SYSTEM-attached HyperPAV aliases observed during the time interval of the report.

#### WhenLow

The date and time when the low value was observed.

#### High

The highest number of available SYSTEM-attached HyperPAV aliases observed during the time interval of the report.

#### WhenHigh

The date and time when the high value was observed.

**Usage note:**

When Performance Toolkit has not processed any Domain 1 Record 20 (MRMTRHPP - HyperPAV Pool Definition) or Domain 6 Record 28 (MRIODHPP - HyperPAV Pool Activity) monitor records before this screen or report was requested, this message appears:

```
FCXxxx494E HPALIAS not available, requires MONITOR I/O SAMPLE data
```

**Effect of AVERAGE and INTERIM subcommands:**

Overall and interim period averages will be shown for all fields except Low, When-Low, High, and When-High, which will display the data identifying the Low and High values observed during the interval and the date and time when the value was observed.

## FCX328, HyperPAV Share Activity Screen – HPSHARE

**Performance Toolkit Subcommand:** “HPSHARE” on page 198.

The HyperPAV Share Activity Screen (FCX328) shows pooling information for HyperPAV alias pools.

FCX328	CPU nnnn		SER nnnnn	Interval HH:MM:SS - HH:MM:SS		Perf. Monitor						
----	<-----Counts----->					<-----Service----->					<Fulfillment>	
SSID	Pool	Base	Alias	SYSTEM	Type	Share	Entlmt	Inuse	Queued	PctUSD	Eligbl	Given
0100	12	66	48	48	MDISK	0	...	...	...	...	0	0
0100	12	66	48	48	PAGING	0	...	...	...	...	0	0
035A	10	25	1	1	MDISK	0	...	...	...	...	0	0
035A	10	25	1	1	PAGING	0	...	...	...	...	0	0
0401	4	50	156	156	MDISK	0	...	0	0	0	0	0
0401	4	50	156	156	PAGING	0	...	2.8	0	0	0	0
0402	0	54	152	152	MDISK	0	...	27.8	.1	0	0	0
0402	0	54	152	152	PAGING	0	...	10.6	0	0	0	0
062B	2	70	156	156	MDISK	0	...	19.3	.0	0	0	0
062B	2	70	156	156	PAGING	0	...	7.6	.1	0	0	0
0633	3	40	156	156	MDISK	0	...	0	0	0	0	0
0633	3	40	156	156	PAGING	0	...	20.0	.1	0	0	0
Command ==> _												
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return							

Figure 222. Layout of HyperPAV Share Activity Screen (FCX328)

### Field Descriptions:

#### SSID

The subsystem ID of the logical control unit (LCU).

#### Pool

The pool number of the HyperPAV pool.

#### Counts:

#### Base

The number of base devices in the LCU associated with this HyperPAV pool.

#### Alias

The number of HyperPAV alias devices in the HyperPAV pool.

#### SYSTEM

The number of HyperPAV alias devices attached to SYSTEM.

#### Type

The type of I/O to which the row applies.

#### Share

The share value for the type of I/O, as set by the SET CU command or the CU statement in the system configuration file.

#### Entlmt

The I/O type's entitlement to aliases, as calculated from the number of SYSTEM-attached HyperPAV aliases and the share settings for the types of I/O.

#### Service:

#### Inuse

The average number of SYSTEM-attached HyperPAV aliases found in use for I/O operations of this type.

#### Queued

The average number of I/O operations of this type found queued on their bases.

**PctUSd**

The percentage of LCU samples finding that this type of I/O is *underserved* by the HyperPAV aliases. In other words, the I/O type is using fewer aliases than it is entitled to and has at least one I/O operation queued.

*Fulfillment:*

**Eligbl**

The rate, per second, at which the pool of available SYSTEM-attached HyperPAV aliases was found empty, an alias became available, and CP was ready to give the alias to an I/O operation of this type in order to fulfill entitlement.

**Given**

The rate, per second, at which the pool of available SYSTEM-attached HyperPAV aliases was found empty, an alias became available, and CP was ready to give the alias to an I/O operation of this type in order to fulfill entitlement and found such an operation queued to which to give the alias.

**Usage Notes:**

When Performance Toolkit has not processed any Domain 1 Record 20 (MRMTRHPP - HyperPAV Pool Definition) or Domain 6 Record 28 (MRIODHPP - HyperPAV Pool Activity) monitor records before this screen or report was requested, this message appears:

```
FCXxxx494E HPSHARE not available, requires MONITOR I/O SAMPLE data
```

**Effect of AVERAGE and INTERIM subcommands:**

Overall and interim period averages will be shown for all fields.



## FCX329, DASD Volume Performance Screen – VOLUME

**Performance Toolkit Subcommand:** “VOLUME” on page 362.

The DASD Volume Performance Screen (FCX329) shows information about the behavior of DASD volumes.

FCX329	CPU	nnnn	SER	nnnnn	Interval	HH:MM:SS	-	HH:MM:SS	Perf.	Monitor								
<-----	Volume	----->	LCU	.	Mdisk	---	.	<-----	Volume	Times	(msec)	----->	Req.	Pct.	<-Serv---			
Addr	Label	WA	SSID	T	Links	I/O	AIOR	Pend	Disc	Conn	Serv	Wait	Resp	CUWt	Qued	Busy	Base	Alias
>>Mean>>	1681	..	...	.	...	3.0	.019	.252	.591	.744	.216	1.19	.036	.047	.03	4	1.35	13.7
Volume Group Data:																		
10TP00L	583	..	...	.	...	8.4	.019	.255	.601	.755	.216	1.21	.034	.048	.08	11	1.37	13.7
Volume Data:																		
C50D	ATP111	CP	0607	H	0	22.4	0	.252	11.6	.685	12.5	8.96	21.4	0	.21	28	12.5	...
C004	ATP032	CP	0602	H	0	22.3	0	.246	10.8	.698	11.7	7.48	19.2	0	.14	26	11.7	...
C701	ATP127	CP	0609	H	0	22.3	0	.251	11.6	.656	12.5	11.6	24.1	0	.27	28	12.5	...
C707	ATP133	CP	0609	H	0	22.3	0	.250	11.4	.708	12.4	10.7	23.1	0	.25	28	12.4	...
CD05	ATP215	CP	060F	H	0	22.3	0	.267	13.4	.694	14.4	13.9	28.3	0	.30	32	14.4	...
CD0A	ATP220	CP	060F	H	0	22.3	0	.257	12.2	.700	13.1	12.1	25.3	0	.27	29	13.1	...
C207	ATP063	CP	0604	H	0	22.2	0	.255	11.7	.734	12.7	9.75	22.5	0	.21	28	12.7	...
C80A	ATP150	CP	060A	H	0	22.2	0	.255	11.6	.690	12.5	8.96	21.5	0	.17	28	12.5	...
C901	ATP155	CP	060B	H	0	22.2	0	.257	11.9	.713	12.9	11.4	24.3	0	.30	29	12.9	...
C90A	ATP164	CP	060B	H	0	22.2	0	.255	11.7	.693	12.6	10.1	22.7	0	.26	28	12.6	...
BE00	ATP000	CP	0600	H	0	22.1	.464	.262	12.3	1.45	14.0	.006	14.0	.769	0	31	13.9	14.4
BF03	ATP017	CP	0601	H	0	22.1	0	.267	13.5	.673	14.4	13.6	28.0	0	.32	32	14.4	...
C30A	ATP080	CP	0605	H	0	22.1	0	.249	11.0	.698	12.0	8.79	20.8	0	.24	26	12.0	...
Command ==> _																		
F1=He1p	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return													

Figure 223. Layout of DASD Volume Performance Screen (FCX329)

Figure 223 on page 865 shows a screen with DASD volumes sorted in descending order of the I/O rate value, as indicated by the underscore characters above the I/O heading. Other possible sort sequences are indicated by dots above the corresponding columns. These sequences can be set using the SORT subcommand with the string shown in the lowest line of the column heading as the sort argument (see “SORT” on page 295) or by positioning the cursor on the heading of the field that is to be sorted and then pressing ENTER.

Some additional lines with volume group averages are shown at the top when volume groups have been defined. (See “FCONTROL DEVGROUP” on page 26.) In Figure 223 on page 865, the volume group name (10TP00L) is displayed in the Addr field, and the number of volumes in the group (583) is displayed in the Label field.

### Field Descriptions:

**Volume:**

#### Addr

The RDEV number of the volume's base RDEV.

The first line, which contains the prefix >>Mean>> instead of a device number, provides the average load statistics for all volumes for which performance data were collected.

#### Label

The volume label.

For the >>Mean>> row or for a volume-group row, this column contains the number of volumes.

#### WA

Indicates where the volume's base RDEV is attached:

**CP**

Indicates that it is in the CP-owned list.

**SY**

Indicates that it is attached to SYSTEM.

**GU**

Indicates that it is attached to a guest.

If the volume is not attached, this field contains dots.

**LCU SSID**

The subsystem ID (SSID) of the logical control unit (LCU).

**T**

Indicates the RDEV type for the volume's base RDEV:

**P**

For a PAV base

**H**

For a HyperPAV base.

If it is neither a PAV base nor a HyperPAV base, this field contains a dot.

**Mdisk Links**

The number of minidisk links to the volume.

**I/O**

The I/O rate. This is the number of SSCH and RSCH operations executed per second to this volume through the base RDEV or any alias RDEV working on behalf of the base during the selected interval.

**AIOR**

The alias I/O ratio. For each I/O operation performed to the base RDEV, this is the number of I/O operations performed through the alias RDEVs.

*Volume Times (msec):*

**Pend**

The "function pending" time. This is the average amount of time that an I/O request remained pending in the I/O subsystem due to path-busy conditions.

This value and the Disc, Conn, Serv, Busy, and Resp values are calculated based on information obtained from the subchannel measurement blocks. Dots are inserted for volumes where the information is not available or is not calculated.

**Disc**

The "volume disconnected" time. This is the average amount of time that the volume remained disconnected from the channel while executing I/O requests. It includes the time spent on executing the SEEK and SET SECTOR orders and also any re-connect misses. High values could indicate overloaded paths or many long SEEKS to data at opposite ends of the volume.

This value also includes the "device-active-only" time where that information is available.

**Conn**

The "volume connected" time. This is the average time that the volume was connected to the channel while executing I/O requests, that is, receiving commands or actually transferring data. It also includes the SEARCH time.

**Serv**

The service time for the volume. This is the sum of the function pending, connected, and disconnected times.

**Wait**

The volume wait time.

**Resp**

The volume response time. This is the sum of the service time and the wait time.

**CUWt**

The control unit queuing time, that is, the additional wait time caused when the target device was already busy with an I/O operation that was started from another system.

**Req. Qued**

The average I/O queue length in the selected interval.

The value is taken from CP monitor high-frequency sample data; it is an indicator of I/O contention on the volume.

The I/O queue length (Req. Qued) is not meaningful for volumes that are dedicated (attached) to guest systems. The guest system will not try to start another I/O operation before an active one has ended, but rather will do its own queuing. Guest system queuing cannot be shown by host performance monitors, so they will have to be analyzed using a performance monitor on the guest system.

The queue length reported here is the number of real I/O operations (channel programs) that CP's real I/O layer is waiting to start to the device or volume. This includes I/O operations performed for all reasons: for minidisk I/O, for paging I/O, and for spooling I/O, for example. For devices in the CP-owned list, additional queueing information can be found on one of these screens:

- [“FCX109, CP-Owned Device Screen – DEVICE CPOWNED” on page 416](#)
- [“FCX170, CPOWNED Device Log Screen – CPOWNLOG” on page 553.](#)

**Pct. Busy**

The percentage of time during which the volume was busy. This value could be greater than 100 if there are alias RDEVs doing work on behalf of the base RDEV.

*Serv:*

**Base**

The average amount of service time per I/O operation through the base RDEV to the volume.

**Alias**

The average amount of service time per I/O operation through an alias RDEV to the volume.

**Defining a volume group:**

To define a volume group, use FC DEVGROUP. Include only non-PAV DASD RDEVs, PAV base RDEVs, and HyperPAV base RDEVs. Any other types of RDEVs are excluded from the calculation.

The appropriate subset of FC DEVGROUP DASD is used for each volume group.

**Effect of AVERAGE and INTERIM subcommands:**

Overall or interim period averages will be shown for all load fields. Real averages will be shown for all fields that display device load figures. The Label, WA, LCU SSID, T, and Mdisk Links fields will always contain current information.

## FCX330, DASD Volume Performance Log Screen – VOLLOG

**Performance Toolkit Subcommand:** “VOLLOG” on page 361.

The DASD Volume Performance Log Screen (FCX330) shows volume-level statistics from “FCX329, DASD Volume Performance Screen – VOLUME” on page 865.

FCX330	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS	Perf. Monitor
DASD Volume Performance Log BE00				
Interval	Volume	LCU	Mdisk	
End Time	Label	WA	SSID T	Links
>>Mean>>				
14:34:16	ATP000	CP	0600	H
14:34:46	ATP000	CP	0600	H
14:35:16	ATP000	CP	0600	H
14:35:46	ATP000	CP	0600	H
14:36:16	ATP000	CP	0600	H
14:36:46	ATP000	CP	0600	H
14:37:16	ATP000	CP	0600	H
14:37:46	ATP000	CP	0600	H
14:38:16	ATP000	CP	0600	H
14:38:46	ATP000	CP	0600	H
14:39:16	ATP000	CP	0600	H
14:39:46	ATP000	CP	0600	H
14:40:16	ATP000	CP	0600	H
14:40:46	ATP000	CP	0600	H
14:41:16	ATP000	CP	0600	H
14:41:46	ATP000	CP	0600	H
14:42:16	ATP000	CP	0600	H
14:42:46	ATP000	CP	0600	H
14:43:16	ATP000	CP	0600	H
14:43:46	ATP000	CP	0600	H
14:44:16	ATP000	CP	0600	H
Command ===> _				
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return				

Figure 224. Layout of DASD Volume Performance Log Screen (FCX330)

The information shown is based on MONITOR domain and I/O domain SAMPLE data.

The first entry is reserved for averages since the last RESET subcommand or since data collection was activated. It is designated by the >>Mean>> entry in the Interval End Time column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval for which the values were calculated. In the case of the >>Mean>>entry, you can find the relevant interval in the header line.

#### Volume Label

The volume label. For the >>Mean>> row or for a volume-group row, this column contains the number of volumes.

#### WA

Indicates where the volume's base RDEV is attached:

##### CP

Indicates that it is in the CP-owned list.

##### SY

Indicates that it is attached to SYSTEM.

##### GU

Indicates that it is attached to a guest.

If the volume is not attached, this field contains dots.

**LCU SSID**

The subsystem ID (SSID) of the logical control unit (LCU).

**T**

Indicates the RDEV type for the volume's base RDEV:

**P**

For a PAV base

**H**

For a HyperPAV base.

If it is neither a PAV base nor a HyperPAV base, this field contains a dot.

**Mdisk Links**

The number of minidisk links to the volume.

**I/O**

The I/O rate. This is the number of SSCH and RSCH operations executed per second to this volume through the base RDEV or any alias RDEV working on behalf of the base during the selected interval.

**AIOR**

The alias I/O ratio. For each I/O operation performed to the base RDEV, this is the number of I/O operations performed through the alias RDEVs.

*Volume Times (msec):*

**Pend**

The "function pending" time. This is the average amount of time that an I/O request remained pending in the I/O subsystem due to path-busy conditions.

This value and the Disc, Conn, Serv, Busy, and Resp values are calculated based on information obtained from the subchannel measurement blocks. Dots are inserted for volumes where the information is not available or is not calculated.

**Disc**

The "volume disconnected" time. This is the average amount of time that the volume remained disconnected from the channel while executing I/O requests. It includes the time spent on executing the SEEK and SET SECTOR orders and also any re-connect misses. High values could indicate overloaded paths or many long SEEKS to data at opposite ends of the volume.

The value also includes the "device-active-only" time where that information is available.

**Conn**

The "volume connected" time. This is the average time that the volume was connected to the channel while executing I/O requests, that is, receiving commands or actually transferring data. It also includes the SEARCH time.

**Serv**

The service time for the volume. This is the sum of the function pending, connected, and disconnected times.

**Wait**

The volume wait time.

**Resp**

The volume response time. This is the sum of the service time and the wait time.

**CUWt**

The control unit queuing time, that is, the additional wait time caused when the target device was already busy with an I/O operation that was started from another system.

**Req. Qued**

The average I/O queue length in the selected interval.

The value is taken from CP monitor high-frequency sample data; it is an indicator of I/O contention on the volume.

The I/O queue length (Req . Qued) is not meaningful for volumes that are dedicated (attached) to guest systems. The guest system will not try to start another I/O operation before an active one has ended, but rather will do its own queuing. Guest system queuing cannot be shown by host performance monitors, so they will have to be analyzed using a performance monitor on the guest system.

The queue length reported here is the number of real I/O operations (channel programs) that CP's real I/O layer is waiting to start to the device or volume. This includes I/O operations performed for all reasons: for minidisk I/O, for paging I/O, and for spooling I/O, for example. For devices in the CP-owned list, additional queueing information can be found on one of these screens:

- [“FCX109, CP-Owned Device Screen – DEVICE CPOWNED” on page 416](#)
- [“FCX170, CPOWNED Device Log Screen – CPOWNLOG” on page 553.](#)

**Pct Busy**

The percentage of time during which the volume was busy. This value could be greater than 100 if there are alias RDEVs doing work on behalf of the base RDEV.

*Serv:*

**Base**

The average amount of service time per I/O operation through the base RDEV to the volume.

**Alias**

The average amount of service time per I/O operation through an alias RDEV to the volume.

**Effect of AVERAGE and INTERIM subcommands:** None.

## FCX331, SMT Metrics Per Core Type Log Screen – MTCTYPLG

**Performance Toolkit Subcommand:** “MTCTYPLG” on page 239.

The SMT Metrics Per Core Type Log Screen (FCX331) shows SMT data for each core type.

FCX331	CPU nnnn	SER nnnnn	Interval	HH:MM:SS - HH:MM:SS	Perf. Monitor		
Interval	Core		Threads	%Core	Avg. Thread	%Capacity	%Max Capacity
End Time	Type	#Cores	per Core	Busy	Density	Factor	Factor
>>Mean>>	CP	26	1	6.6	1.00	100.0	100.0
>>Mean>>	IFL	10	2	99.8	1.51	139.5	177.6
>>Mean>>	ZIIP	2	1	.0	1.00	100.0	100.0
>>Mean>>	ICF	2	1	.0	1.00	100.0	100.0
03:32:23	CP	26	1	6.5	1.00	100.0	100.0
03:32:23	IFL	10	2	99.9	1.50	137.9	175.1
03:32:23	ZIIP	2	1	.0	1.00	100.0	100.0
03:32:23	ICF	2	1	.0	1.00	100.0	100.0
03:33:23	CP	26	1	6.4	1.00	100.0	100.0
03:33:23	IFL	10	2	99.6	1.51	139.3	177.1
03:33:23	ZIIP	2	1	.0	1.00	100.0	100.0
03:33:23	ICF	2	1	.0	1.00	100.0	100.0
03:34:23	CP	26	1	6.8	1.00	100.0	100.0
03:34:23	IFL	10	2	99.9	1.51	139.7	178.0
03:34:23	ZIIP	2	1	.0	1.00	100.0	100.0
03:34:23	ICF	2	1	.0	1.00	100.0	100.0
Command							
===>							
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return		

Figure 225. Layout of SMT Metrics Per Core Type Log Screen (FCX331)

The information shown is based on CP monitor MONITOR domain SAMPLE, SYSTEM domain CONFIGURATION, and PROCESSOR domain EVENT data.

The first entry is reserved for the overall averages since the last RESET subcommand or since data collection was activated. It is designated by the >>Mean>> string in the Interval End Time column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### Core Type

The core type of the logical cores defined for the partition. Possible values are:

##### CP

Central Processor

##### ICF

Internal Coupling Facility

##### IFL

Integrated Facility for Linux

##### ZIIP

IBM Z Integrated Information Processor

**Note:** For the ZIIP cores, the monitor data shows IIP. Performance Toolkit translates this value as ZIIP.

#### #Cores

The number of cores that returned data for this sample interval.

MIX is displayed on the >>Mean>> line if the number of cores of a given type changes during any of the sample intervals. This happens when the system is switched to MT-1 or MT-2 mode or when cores are varied online or offline.

**Note:** Some cores might not return data in some intervals.

#### **Threads per Core**

The number of activated threads on cores of this type.

MIX is displayed on the >>Mean>> line if the number of cores of a given type changes during any of the sample intervals. This happens when the system is switched to MT-1 or MT-2 mode.

#### **%Core Busy**

The percentage of time the logical cores were busy, defined as the percentage of elapsed time during which the logical cores were assigned to physical cores. A value of 100% represents one logical core assigned to one physical core for the entire measurement interval. If the partition has been enabled for the wait-completion assist, this value includes wait time.

#### **Avg. Thread Density**

The average number of threads that were active per core when a core of this type was dispatched and at least one thread of the core was active.

#### **%Capacity Factor**

The ratio of the work rate the cores experienced to the work rate the cores experienced with only one thread busy, expressed as a percentage.

#### **%Max Capacity Factor**

The ratio of the work rate the cores experienced with all threads busy to the work rate the cores experienced with only one thread busy, expressed as a percentage.

#### **Usage Note:**

The SMT Metrics Per Core Type Log report is available only when multithreading is enabled:

1. For the monitored system when monitoring live data
2. On the system that generated a MONWRITE file that Performance Toolkit for z/VM is processing.

#### **Effect of AVERAGE and INTERIM subcommands:**

None.



## FCX332, SMT Metrics Per Core Log Screen – MTCORELG

**Performance Toolkit Subcommand:** “MTCORELG” on page 238.

The SMT Metrics Per Core Log Screen (FCX332) shows SMT data for each core.

FCX332	CPU nnnn	SER nnnnn	Interval	HH:MM:SS	- HH:MM:SS	Perf. Monitor	
Interval End Time	Core	Core Type	Threads per Core	Core	PPD Ent	%Core Busy	Avg. Thread Density
>>Mean>>	00	CP	1	Vh	100	53.2	.50
>>Mean>>	01	CP	1	Vh	100	40.6	.50
>>Mean>>	02	CP	1	Vh	100	21.1	.50
>>Mean>>	03	CP	1	Vh	100	11.7	.50
>>Mean>>	04	CP	1	Vh	100	9.8	.50
>>Mean>>	05	CP	1	Vh	100	12.2	.50
.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.
>>Mean>>	1A	IFL	2	Vh	100	99.9	1.09
>>Mean>>	1B	IFL	2	Vh	100	99.9	1.06
>>Mean>>	1C	IFL	2	Vh	100	99.9	.95
>>Mean>>	1D	IFL	2	Vh	100	99.9	.83
>>Mean>>	1E	IFL	2	Vh	100	99.9	1.16
>>Mean>>	1F	IFL	2	Vh	100	99.9	1.14
03:32:23	00	CP	1	Vh	100	50.8	1.00
03:32:23	01	CP	1	Vh	100	34.6	1.00
03:32:23	02	CP	1	Vh	100	17.5	1.00
03:32:23	03	CP	1	Vh	100	11.8	1.00
03:32:23	04	CP	1	Vh	100	10.9	1.00
03:32:23	05	CP	1	Vh	100	14.9	1.00
.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.
03:32:23	1A	IFL	2	Vh	100	99.9	1.71
03:32:23	1B	IFL	2	Vh	100	99.9	1.65
03:32:23	1C	IFL	2	Vh	100	99.9	1.57
03:32:23	1D	IFL	2	Vh	100	99.9	1.42
03:32:23	1E	IFL	2	Vh	100	100.0	1.69
03:32:23	1F	IFL	2	Vh	100	99.9	1.68
Command ==>							
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return							

Figure 226. Layout of SMT Metrics Per Core Log Screen (FCX332)

The information shown is based on CP monitor MONITOR domain SAMPLE, SYSTEM domain CONFIGURATION, and PROCESSOR domain EVENT data.

The first entry is reserved for the overall averages since the last RESET subcommand or since data collection was activated. It is designated by the >>Mean>> string in the Interval End Time column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### Core

The logical core ID in hexadecimal format.

#### Core Type

The core type of this logical core. Possible values are:

##### CP

Central Processor

##### ICF

Internal Coupling Facility

##### IFL

Integrated Facility for Linux

**ZIIP**

IBM Z Integrated Information Processor

**Note:** For the ZIIP cores, the monitor data shows IIP. Performance Toolkit translates this value as ZIIP.

**Threads per Core**

The number of activated threads on this core.

MIX is displayed on the >>Mean>> line if the number of activated threads changes during any of the sample intervals. This happens when the system is switched to MT-1 or MT-2 mode.

**PPD**

The current polarization of the core. Possible values are:

**Vh**

Polarization is vertical, with high entitlement.

**VhD**

Polarization is vertical, with high entitlement, and the partition is a dedicated partition.

**Vm**

Polarization is vertical, with medium entitlement.

**VI**

Polarization is vertical, with low entitlement.

**MIX**

A mixed polarization, which can be displayed:

- On the >>Mean>> line if multiple distinct polarization values were encountered during the measured interval (average, from the last reset).
- On an Interval End Time line if the polarization value for a core changes during the interval for which the data is displayed.

**Ent**

The average entitlement for the logical core as a percentage, where 100% represents one whole physical core's worth of processing capacity.

**%Core Busy**

The percentage of time the logical core was busy, defined as the percentage of elapsed time during which the logical core was assigned to a physical core. If the partition has been enabled for the wait-completion assist, this value includes wait time.

**Avg. Thread Density**

The average number of threads that were active on this core when at least one thread of this core was active.

**Usage Note:**

The SMT Metrics Per Core Log report is available only when multithreading is enabled:

1. For the monitored system when monitoring live data
2. On the system that generated a MONWRITE file that Performance Toolkit for z/VM is processing.

**Effect of AVERAGE and INTERIM subcommands:**

None.

## FCX333, User Processor Time Screen – USRPRCTM

**Performance Toolkit Subcommand:** “USRPRCTM” on page 349.

The User Processor Time Screen (FCX333) shows three sets of user processor time data that are calculated for each user with activity on virtual IFL processors.

FCX333	CPU nnnn	SER nnnnn	Interval HH:MM:SS - HH:MM:SS				Perf. Monitor				
-----	.	.	.	.	.	.	<---	Virt.	CPUs	----	Nr of
<b>Userid</b>	<b>T-Raw</b>	<b>V-Raw</b>	<b>T-MT1</b>	<b>V-MT1</b>	<b>T-Pro</b>	<b>V-Pro</b>	<b>Type</b>	<b>Aff</b>	<b>Def</b>	<b>Stop</b>	<b>Users</b>
>>Mean>>	26.843	26.822	26.556	26.536	26.349	26.329	---	---	---	---	23
<b>User Class Data:</b>											
LINX	41.109	41.093	40.600	40.584	40.467	40.451	---	---	---	---	2
<b>User Data:</b>											
CFT2NDB	.019	.005	.019	.005	.017	.004	IFL	On	12	0	
CHG00001	.020	.013	.020	.013	.020	.013	IFL	On	1	0	
CHG00002	.002	.001	.002	.001	.002	.001	IFL	On	1	0	
CHG00003	.048	.035	.048	.035	.047	.034	IFL	On	1	0	
CHG00004	.453	.318	.451	.316	.447	.314	IFL	On	1	0	
CHG00005	.010	.007	.010	.006	.010	.006	IFL	On	1	0	
CHG00006	.122	.077	.122	.077	.115	.072	IFL	On	1	0	
CHG00007	.212	.189	.210	.188	.212	.189	IFL	On	1	0	
IFG00001	41.090	41.074	40.603	40.587	40.275	40.259	IFL	On	1	0	
IFG00002	41.127	41.111	40.597	40.581	40.658	40.642	IFL	On	1	0	
IFG00003	41.125	41.110	40.678	40.663	39.710	39.695	IFL	On	1	0	
IFG00004	41.143	41.128	40.670	40.655	39.710	39.695	IFL	On	1	0	
IFG00005	41.116	41.101	40.672	40.657	41.071	41.055	IFL	On	1	0	
IFG00006	41.140	41.125	40.782	40.767	40.358	40.343	IFL	On	1	0	
<b>Command ==&gt;</b>											
<b>F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return</b>											

Figure 227. Layout of User Processor Time Screen (FCX333)

The information shown is based on CP monitor MONITOR domain and USER domain SAMPLE data.

The first entry is reserved for the overall averages since the last RESET subcommand or since data collection was activated. It is designated by the >>Mean>> string in the Interval End Time column.

Some additional lines with user class averages will be shown at the top when user classes have been defined (see “FCONTROL UCLASS” on page 89). The user ID shown will then be the user class name.

### Field Descriptions:

#### Userid

The user identification.

#### T-Raw

The total raw time that has been expended for this user. This includes time that the user spent doing its own work and time that CP spent doing work on behalf of the user.

#### V-Raw

The virtual raw time that this user was running and doing productive work. It does not include any time that CP spends on behalf of the user performing such tasks as instruction simulation or page translation.

#### T-MT1

The total MT-1 equivalent time that has been expended for this user. This includes time that the user spent doing its own work and time that CP spent doing work on behalf of the user.

#### V-MT1

The virtual MT-1 equivalent time that this user was running and doing productive work. It does *not* include any time that CP spends on behalf of the user performing such tasks as instruction simulation or page translation.

**T-Pro**

The total prorated core time that has been expended for this user. This includes time that the user spent doing its own work and time that CP spent doing work on behalf of the user.

**V-Pro**

The virtual prorated core time that this user was running and doing productive work. It does not include any time that CP spends on behalf of that user performing such tasks as instruction simulation or page translation.

*Virt. CPUs:*

The next four fields contain information about virtual processors. On the FCX333 screen (USRPRCTM), these fields contain dashes on the >>Mean>> line and the User Class Data lines.

**Type**

The type of virtual processor group. Valid values are:

**IFL**

Integrated Facility for Linux

**Aff**

Indicates the affinity of the virtual processor type. Valid values are:

**Off**

CPU affinity has not been set on for this user. When CPU affinity is Off, the virtual processor is dispatched on only logical processors of the primary type.

**On**

CPU affinity has been set on for this user using the CP SET CPUAFFINITY command. When CPU affinity is On, this virtual processor is dispatched on a logical processor of the same type.

**Sup**

CPU affinity is suppressed. Although CPU affinity has been set on for this user using the CP SET CPUAFFINITY command, there is currently no matching logical processor on which to run.

**Def**

The number of virtual processors of the specified type in the virtual machine's configuration.

**Stop**

The number of virtual processors that are in a stopped state.

**Nr of Users**

The number of users that contributed to the average shown on this line. Values are inserted only for lines with averages; this field is blank on user data lines.

**Usage Notes:**

1. The User Processor Time report is available only when multithreading is enabled:
  - a. For the monitored system when monitoring live data
  - b. On the system that generated a MONWRITE file that Performance Toolkit for z/VM is processing.
2. The User Processor Time report displays data for only virtual IFL processors because IFL is the only type of processor on which z/VM supports multithreading.

**Effect of AVERAGE and INTERIM subcommands:**

Average and interim values are shown as requested for all time fields. Current data are shown for the Virtual CPU fields.

## FCX334, User Processor Time Log Screen – USRTMLOG

**Performance Toolkit Subcommand:** “USRTMLOG” on page 350.

The User Processor Time Log Screen (FCX334) shows a "by time" log of the selected user's processor time on virtual IFL processors.

FCX334	CPU nnnn		SER nnnnn		Interval HH:MM:SS - HH:MM:SS		Perf. Monitor			
Processor Time Log for user IFG00001										
Interval							<---	Virt.	CPUs	---
End Time	T-Raw	V-Raw	T-MT1	V-MT1	T-Pro	V-Pro	Type	Aff	Def	Stop
>>Mean>>	59.851	59.837	55.224	55.211	46.731	46.720	---	---	---	---
03:31:23	41.090	41.074	40.603	40.587	40.275	40.259	IFL	On	1	0
03:32:23	59.880	59.866	55.322	55.309	48.454	48.443	IFL	On	1	0
03:33:23	59.888	59.875	55.030	55.018	41.550	41.541	IFL	On	1	0
03:34:23	59.884	59.871	54.982	54.970	47.975	47.964	IFL	On	1	0
03:35:23	59.829	59.817	55.035	55.024	36.857	36.850	IFL	On	1	0
03:36:23	59.769	59.756	54.844	54.832	36.646	36.638	IFL	On	1	0
03:37:23	59.738	59.724	55.050	55.037	35.485	35.476	IFL	On	1	0
03:38:23	59.898	59.883	56.925	56.911	55.661	55.647	IFL	On	1	0
03:39:23	59.906	59.891	56.975	56.961	59.852	59.837	IFL	On	1	0
03:40:23	59.895	59.879	52.951	52.937	59.821	59.806	IFL	On	1	0
Command ==>										
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return										

Figure 228. Layout of User Processor Time Log Screen (FCX334)

The information shown is based on CP monitor MONITOR domain and USER domain SAMPLE data.

The first entry is reserved for the overall averages since the last RESET subcommand or since data collection was activated. It is designated by the >>Mean>> string in the Interval End Time column.

### Field Descriptions:

#### Interval End Time

The end time of the monitor interval for which the values were calculated. In the case of the >>Mean>> entry, you can find the relevant interval in the header line.

#### T-Raw

The total raw time that has been expended for this user. This includes time that the user spent doing its own work and time that CP spent doing work on behalf of the user.

#### V-Raw

The virtual raw time that this user was running and doing productive work. It does not include any time that CP spends on behalf of the user performing such tasks as instruction simulation or page translation.

#### T-MT1

The total MT-1 equivalent time that has been expended for this user. This includes time that the user spent doing its own work and time that CP spent doing work on behalf of the user.

#### V-MT1

The virtual MT-1 equivalent time that this user was running and doing productive work. It does *not* include any time that CP spends on behalf of the user performing such tasks as instruction simulation or page translation.

#### T-Pro

The total prorated core time that has been expended for this user. This includes time that the user spent doing its own work and time that CP spent doing work on behalf of the user.

#### V-Pro

The virtual prorated core time that this user was running and doing productive work. It does not include any time that CP spends on behalf of that user performing such tasks as instruction simulation or page translation.

*Virt. CPUs:*

The next four fields contain information about virtual processors.

**Type**

The type of virtual processor group. Valid values are:

**IFL**

Integrated Facility for Linux

**Aff**

Indicates the affinity of the virtual processor type. Valid values are:

**Off**

Indicates that CPU affinity has not been set on for this user.

**On**

Indicates that CPU affinity has been set on for this user using the CP SET CPUAFFINITY command and that this virtual processor is dispatched on a logical processor of the same type.

**Sup**

Indicates that CPU affinity is suppressed. Although CPU affinity has been set on for this user using the CP SET CPUAFFINITY command, there is currently no matching logical processor on which to run.

**Def**

The number of virtual processors of the specified type in the virtual machine's configuration.

**Stop**

The number of virtual processors that are in a stopped state.

**Usage Notes:**

1. The User Processor Time Log report is available only when multithreading is enabled:
  - a. For the monitored system when monitoring live data
  - b. On the system that generated a MONWRITE file that Performance Toolkit for z/VM is processing.
2. The User Processor Time Log report displays data for only virtual IFL processors because IFL is the only type of processor on which z/VM supports multithreading.

**Effect of AVERAGE and INTERIM subcommands:**

None.

---

# Chapter 3. Performance Toolkit Messages

In general, Performance Toolkit messages are displayed in the special message area above the command line on the screen. Some exceptions follow:

- A-type messages generated by the performance monitoring function, which indicate that an operator action is required (due to a high system load, for example). These message lines are inserted into the general output area, where they automatically receive the action message prefix +, which makes them non-scrollable. They can also be processed using the FCONTROL PROCESS command to automatically initiate some action.
- M-type messages, which are generated by Performance Toolkit's performance monitoring routines primarily for informational purposes. These message lines are also inserted into the general output area so that they can be viewed in redisplay mode and processed using the FCONTROL PROCESS command.
- Other messages that are generated while the normal message area is already in use. Such messages are inserted into the general output area, but additional processing by the FCONTROL PROCESS command is not possible.

In basic mode and redisplay mode, the messages shown in the special message area remain on the screen until the next command is entered, even if the screen is updated because new output is received. Messages are removed immediately or are replaced by new messages whenever the screen is updated in performance monitor mode.

The line type to be specified for Performance Toolkit messages in the FCONTROL COLOR and FCONTROL PROCESS commands is ERRMSG.

**Note:** z/VM Performance Data Pump issues messages that have a prefix of FCX and are in the range 3000-3099. Those messages are not documented in *z/VM: Performance Toolkit Reference*. See [Data Pump Messages](#) in *z/VM: Other Components Messages and Codes*.

---

**FCXxxx001E      Performance Toolkit requires  
z/CMS to run - exiting**

**Explanation:**

Performance Toolkit has been invoked under CMS running in 31-bit ESA/390 mode. However, this Performance Toolkit version was designed to exploit the z/Architecture mode available in z/CMS operating system to use virtual memory above the 2G line.

**System action:**

The initialization of Performance Toolkit ends.

**User response:**

Re-IPL ZCMS and rerun Performance Toolkit.

---

**FCXxxx005E      This is not a supported 3270 type  
display**

**Explanation**

The terminal that has been used as console when starting Performance Toolkit does not appear to the program as one of the valid display models.

Valid Performance Toolkit terminals are all display terminals which are supported as a VM LOGON console and for full-screen operation.

Calling Performance Toolkit from non-display type terminals or from unsupported display types or models will lead to the above message.

**System action:**

The initialization of Performance Toolkit ends.

**User response:**

Call Performance Toolkit from a supported terminal model.

---

**FCXxxx010E      Insufficient storage for  
initialization**

**Explanation:**

Performance Toolkit requires some virtual storage for the program itself and might require a considerable amount more to hold its data tables when performance monitoring is activated. The required storage is not available.

**System action:**

The initialization of Performance Toolkit ends.

**User response:**

Increase the storage size of your virtual machine, re-IPL ZCMS, and then restart Performance Toolkit.

---

**FCXxxx012E      PERFKIT called recursively - not  
supported**

## Explanation

A command has been entered to start PERFKIT from within an PERFKIT session. This operating mode is not allowed.

## System action

PERFKIT initialization is aborted.

## User response

None.

---

<b>FCXxxx015E</b>	<b>Invalid parms - enter 'PERFKIT ?' for help</b>
-------------------	---

---

## Explanation

An invalid argument has been passed to PERFKIT.

## System action

PERFKIT initialization ends.

## User response

Refer to "General System Operation" in the *z/VM: Performance Toolkit Guide* for the correct command format.

---

<b>FCXxxx030E</b>	<b>Error nn reading 'FCONX \$PROFILE *'</b>
-------------------	---

---

## Explanation

Performance Toolkit found an error condition when it tried to read the initialization profile from disk. The error code 'nn' is the return code of the CMS 'FSREAD'-macro.

## System action

Performance Toolkit initialization ends.

## User response

See the description of the FSREAD macro, and of its return codes, in the corresponding CMS manual and take the appropriate action.

---

<b>FCXxxx040E</b>	<b>Master file name missing</b>
-------------------	---------------------------------

---

## Explanation

The program was invoked in batch mode but no master file name has been specified. The master file controls operation in batch mode and is required for further processing.

## System action

Batch processing ends.

## User response

See "Operation in Batch Mode" in the *z/VM: Performance Toolkit Guide* for a description of the command format and of the required control files.

---

<b>FCXxxx041E</b>	<b>Master file 'fn ft fm' not found</b>
-------------------	---

---

## Explanation

The master file 'fn ft fm' specified when invoking the program in batch mode was not found.

## System action

Batch processing ends.

## User response

Make sure a master file exists and invoke the program again, specifying the correct master file ID.

---

<b>FCXxxx042E</b>	<b>Monitor data file specification missing</b>
-------------------	--

---

## Explanation

The program was invoked in batch mode without specifying the monitor data file that should have been processed.

## System action

Batch processing ends.

## User response

Invoke the program again and include the argument(s) for specifying the monitor data file.

---

<b>FCXxxx043E</b>	<b>Monitor file 'fn ft fm' not found</b>
-------------------	--

---

## Explanation

The monitor data file 'fn ft fm' specified when invoking the program in batch mode was not found.

## System action

Batch processing ends.

## User response

Invoke the program again, specifying a valid monitor data file.



---

**FCXxxx044E      Tape unit 'xxxx' not found**

---

**Explanation**

The tape unit 'xxxx' specified as monitor data input device when invoking the program in batch mode was not found.

**System action**

Batch processing ends.

**User response**

Invoke the program again, specifying the correct tape device number.

**Explanation**

This message is preceded by message 046E. The RUNFILE contains a log of all prior batch mode actions and will list the previous error messages that have led to this message.

**System action**

Batch processing ends.

**User response**

See the error message(s) in the RUNFILE and correct the problem(s) before invoking batch mode again.

---

**FCXxxx045E      Invalid argument 'ccccccc' in command string**

---

**Explanation**

The argument 'ccccccc' entered when invoking the program in batch mode is not valid, or not valid in that position.

**System action**

Batch processing ends.

**User response**

See "Batch Mode Invocation" in the *z/VM: Performance Toolkit Guide* for a description of the required syntax and reenter the command.

---

**FCXxxx050E      Error nn in FSSTATE for 'mmdy CONLOG'**

---

**Explanation**

The Performance Toolkit found an error condition when it tried to access the console log file with file name 'mmdy'. The error code 'nn' is the return code of the CMS 'FSSTATE'-macro.

**System action**

The Performance Toolkit initialization ends.

**User response**

See the description of the FSSTATE macro, and of its return codes, in the corresponding CMS manual and take the appropriate action.

---

**FCXxxx046E      Input file specification error - cannot continue**

---

**Explanation**

The program has been invoked in batch mode, but a severe specification error has been found in an input file. This message will usually be followed by message 047E.

**System action**

Batch processing ends.

**User response**

Follow the actions recommended for message 047E.

---

**FCXxxx060E      Error nn during xxxxxx xxxxxx - cannot continue**

---

**Explanation**

The Performance Toolkit could not properly establish the connection to the CP \*MSG service. The error code 'nn' is the return code of the 'HNDIUCV SET' or 'CMSIUCV CONNECT' function as defined by 'xxxxxx xxxxxx'.

**System action**

The Performance Toolkit initialization ends.

**User response**

The problem is probably caused by an other CMS program which has already used these macros for establishing its own IUCV environment and then has not reset it. Re-IPL ZCMS and then restart the Performance Toolkit.

---

**FCXxxx047E      See RUNFILE for detailed error description**

---

---

**FCXxxx061E      Invalid IUCV interrupt buffer received**

---

**Explanation**

Performance Toolkit received control from CMS for handling an IUCV \*MSG service interrupt, but the IUCV interrupt buffer contains old data that is not valid. It is possible that another CMS program has been activated under Performance Toolkit and this program also established an IUCV connection, but it used the DECLARE BUFFER function instead of the HNDIUCV SET and CMSIUCV CONNECT macros and so disabled Performance Toolkit's interface.

**System action**

Performance Toolkit will exit after issuing this message.

**User response**

Refrain from using the CMS program that activated the IUCV connection. This is a Performance Toolkit restriction; see "Restrictions for Operation in Basic Command Mode" in the *z/VM: Performance Toolkit Guide*. Re-IPL ZCMS, if necessary, and then restart Performance Toolkit.

---

**FCXxxx062E      Error nn during IUCV receive**

---

**Explanation**

The Performance Toolkit found an error while receiving data via IUCV. The error code 'nn' is the return code (IPRCODE field in IPARML) for the RECEIVE function.

**System action**

The Performance Toolkit will exit after issuing this message.

**User response**

See the description of the IUCV protocol, and of IUCV return codes, in the appropriate manual of your system reference library.

---

**FCXxxx070E      Performance Toolkit not supported by this VM level**

---

**Explanation**

Performance Toolkit has detected that one of the prerequisites for successful initialization is missing.

**System action**

Performance Toolkit will exit after issuing this message.

**User response**

Inform your system programmer if this message is generated during initialization under a supported z/VM release.

---

**FCXxxx091I      This system has not been authorized for using the program**

---

**Explanation**

An attempt was made to start Performance Toolkit on a system where Performance Toolkit has not been licensed or where Systems Management API (SMAPI) is not monitoring system performance.

**System action**

The command terminates.

**User response**

Follow the instructions in the Performance Toolkit Program Directory to obtain a license for the product or follow the directions in *z/VM: Systems Management Application Programming* to enable SMAPI monitoring of system performance.

---

**FCXxxx100I      HMA storage start.size is being used for temporary work area**

---

**Explanation:**

During initialization, Performance Toolkit has detected available High Memory Area (HMA) virtual storage residing above the 2G line and successfully allocated it for internal work area usage. *start* designates the HMA start address and *size* is the HMA size.

**System action:**

Operation continues.

**User response:**

None.

---

**FCXxxx101E      Error nn allocating subpool USERG - HMA start.size won't be used**

---

**Explanation:**

During initialization, Performance Toolkit has detected available High Memory Area (HMA) virtual storage residing above the 2G line and attempted to reserve it via allocating memory in the USERG subpool by CMSSTOR macro. However, the attempt was unsuccessful, and Performance Toolkit will continue to run using the conventional memory below the 2G line

only. The error code *nn* is the return code from the CMSSTOR macro.

**System action:**

Operation continues.

**User response:**

See the description of the CMSSTOR OBTAIN macro and its return codes in *z/VM: CMS Macros and Functions Reference*. Re-IPL ZCMS, making sure no other program allocates any virtual storage in the USERG subpool. Then, reinitialize Performance Toolkit.

---

**FCXxxx110E      Error *nn* writing file 'xxxxxx yyyyyy z'**

---

**Explanation:**

Performance Toolkit found an error while writing the file xxxxxx yyyyyy z. The error code *nn* is the return code of the CMS FSWRITE macro.

**System action:**

Performance Toolkit will stop writing to the Z-disk after finding the first writing error (no console log to be re-displayed).

**User response**

See the description of the FSWRITE macro and its return codes in *z/VM: CMS Macros and Functions Reference* and take the appropriate action.

Likely candidates are:

- RC 12: Disk Z is read-only
- RC 13: Disk Z is full

Re-initialize Performance Toolkit after correcting the problem.

---

**FCXxxx111I      Critical exception message(s) issued**

---

**Explanation**

The Performance Toolkit encountered a severe internal processing problem and issued one of the following exception messages: FCXxxx199A, FCXxxx327A, FCXxxx431E, FCXxxx440E, FCXxxx441E, FCXxxx443E, FCXxxx444E, FCXxxx445E, FCXxxx446E, FCXxxx447E, FCXxxx458A, FCXxxx801A.

These messages usually indicate that the collected and calculated data for some of the Performance Toolkit reports may be inaccurate, incorrect, or missing.

**System action**

Operation continues.

**User response**

If possible, correct the problem as described in the explanation text for these exception messages and re-initialize Performance Toolkit. For message FCXxxx458A, Performance Toolkit will attempt to automatically recover from the encountered error by reestablishing the connection to the \*MONITOR system service.

This notification message FCXxxx111I will stay displayed until the exception messages are explicitly viewed either by entering EXCEPT or BASMODE commands, or just by returning to the basic mode screen. In batch mode of operation, this message will be the last one written to CONLOG and RUNFILE.

---

**FCXxxx112I      Cannot write to disk - correct problem and re-initialize PERFKIT**

---

**Explanation**

When running realtime:

Performance Toolkit could no longer write its log and PROCESS files to disk because of an error condition found previously when trying to write to the disk. This message is preceded by message FCXxxx110E, which also indicates the reason for the original problem.

When running in batch mode:

This could indicate that a disk pointed to by the FCONX MASTER file output is either read-only or full. If the RUNFILE is available, the FCXxxx110E message should give an indication for the original problem.

**System action:**

Performance Toolkit continues displaying this message until the problem has been solved.

**User response:**

Quit Performance Toolkit, take the appropriate action as indicated by the return code of message FCXxxx110E, and then re-initialize Performance Toolkit. If you do not remember the return code in message FCXxxx110E, quit and re-initialize Performance Toolkit first: it will then re-issue this message so that you know what action to take.

---

**FCXxxx113E      System configuration missing| changed - processing terminated**

---

**Explanation**

During TRNDSCAN/MONSCAN/BATCH processing the Performance Toolkit detected an error in the system configuration record, monitor domain 1 record 4 (D1R4). This record contains the vital system counters required for correct monitor data handling and reports creation. The system configuration record either has

information missing or the system configuration was changed:

- **Missing:** at startup, the Performance Toolkit determined that the D1R4 record could not be found in the initial monitor sample. (For example, the MONWRITE utility program with the NOCONFIG option produced mondata files where only the first data chunk had the D1R4 record in it, so in order to process these data files they had to be merged with the first one containing the system configuration data).
- **Changed:** while performing calculations, the Performance Toolkit received a new D1R4 record indicating a z/VM IPL occurred that made further data processing impossible. (The Performance Toolkit does not support a single mondata file spanning the system IPL).

**System action**

Operation terminates.

**User response**

Make sure that the monitor data file contains the D1R4 record and does not span the z/VM IPL.

---

**FCXxxx114I**      **Last valid sample interval ends at yy/mm/dd hh:mm:ss**

---

**Explanation**

During MONSCAN/BATCH processing, Performance Toolkit detected an error condition in the monitor data stream that made further processing impossible. The operation terminated. This message shows the last possible sample interval end time that can be used with the FC MONCOLL RESET or NEXTSAMP subcommand. It should be preceded by the message FCXxxx113E or FCXxxx117E, indicating what caused the problem.

**System action**

Operation terminates.

**User response**

None.

---

**FCXxxx115I**      **Non-standard filename 'cccccccc' cannot switch logs**

---

**Explanation**

The user tried to scroll beyond the bounds of the console log file being viewed. Since the log's file name cccccccc no longer has the original standard file name

format *mmddyy*, the file name of the next or previous day's log cannot be determined, and automatic log switching is not possible.

**System action**

The scrolling request is ignored.

**User response**

Determine the file identification of the console log file and select it explicitly by entering the REDISP command with the file name, file type and file mode of the file to be viewed (see “[REDISP](#)” on page 114 for the command description).

---

**FCXxxx116I**      **No default console log found with filename 'mmddyy'**

---

**Explanation:**

The user tried to scroll beyond the bounds of the console log file being viewed. Performance Toolkit tried to display the next or previous day's console log, with an expected file name of *mmddyy* and a file type of CONLOG *n*, but did not find such a file on any of the accessed disks.

**System action:**

The scrolling request is ignored.

**User response:**

If the console log exists, but was renamed to a non-standard file name, select the log explicitly by entering the REDISP command with the file name, file type, and file mode of the file to be viewed (see “[REDISP](#)” on page 114 in the *z/VM: Performance Toolkit Reference* for the command description).

---

**FCXxxx117E**      **Out-of-sequence record DdRr detected at yy/mm/dd hh:mm:ss**

---

**Explanation**

During MONSCAN/BATCH processing, Performance Toolkit detected a sample record in the monitor data stream (domain *d*, record *r*) with the TOD clock *yy/mm/dd hh:mm:ss*, earlier than the last processed sample record's TOD clock. Most likely, the monitor data file was created by merging multiple data chunks in the wrong order.

**System action**

Operation terminates.

**User response**

Make sure that the composite monitor data file is concatenated correctly.

---

**FCXxxx118I      Excessive parameters specified for  
PIPE mode - ignored****Explanation**

Performance Toolkit was called in MONSCAN or BATCH mode as a PIPE stage. However, some extraneous parameters have been specified on the command line (such as DISK or TAPE) that are unneeded when running Performance Toolkit in PIPE mode. They are ignored.

**System action**

MONSCAN or BATCH processing continues.

**User response**

None.

---

**FCXxxx119E      'fn ft fm' is not a valid monitor data  
file****Explanation:**

Performance Toolkit was invoked in MONSCAN or BATCH mode with the monitor data file *fn ft fm* selected as input, but the specified file format is not correct. Note that a valid file has RECFM=F, LRECL=4096, and starts with the monitor control record.

**System action:**

MONSCAN or BATCH processing terminates.

**User response:**

Make sure a valid monitor data file is specified as input.

---

**FCXxxx120E      Error nn reading file 'xxxxxx yyyyyy  
m'****Explanation:**

Performance Toolkit found an error while reading the file *xxxxxx yyyyyy*. The error code *nn* is the return code of the CMS FSREAD macro.

**System action:**

This depends on file that is being read.

**User response:**

See the description of the FSREAD macro and of its return codes in *z/VM: CMS Macros and Functions Reference* and take the appropriate action.

---

**FCXxxx121E      Return code nn reading tape  
xxxx:cccccc****Explanation**

An error was detected while trying to read the tape on unit *xxxx*. The code *nn* is the return code of the CMS 'RDTAPE' macro. The string *cccccc* provides an

explanation for some of the more common return codes.

**System action**

No monitor sample could be read.

**User response**

See the description of the RDTAPE macro, and of its return codes, in the corresponding CMS manual and take the appropriate action.

---

**FCXxxx122I      Node-ID not found in SYSTEM  
NETID file****Explanation:**

An FC MONCOLL SYSTEMID NODEID command has been entered, but Performance Toolkit could not find an entry in the SYSTEM NETID file that matched the system's CPU ID. The node ID cannot be inserted in the performance monitor header line.

**System action:**

Operation continues.

**User response:**

None required. If desired, provide a correct SYSTEM NETID file, with an entry for the CPU you are working on, and then re-issue the command.

---

**FCXxxx123E      Invalid monitor data record  
detected while reading from PIPE****Explanation**

The Performance Toolkit was invoked in MONSCAN or BATCH mode as a PIPE stage. However, while processing the monitor data from PIPE stream, an invalid record has been read. Note that a valid monitor data stream consists of the records with LRECL=4096, and starts with the monitor control record.

**System action**

MONSCAN or BATCH processing terminates.

**User response**

Make sure a valid monitor data stream is specified as PIPE input.

---

**FCXxxx124E      Error reading monitor data from  
PIPE****Explanation**

The Performance Toolkit was invoked in MONSCAN or BATCH mode as a PIPE stage. However, while reading the monitor data from PIPE stream, a pipeline error has occurred.

### System action

MONSCAN or BATCH processing terminates.

### User response

Make sure a valid monitor data stream is specified as PIPE input, and the pipeline is built correctly.

---

**FCXxxx125E      Error *nn* for FSSTATE of file *xxxxxx*  
*yyyyyy m***

#### Explanation:

Performance Toolkit found an error while looking for file the *xxxxxx yyyyyy* file. This error can occur when trying to read a console log in redisplay mode. The error code *nn* is the return code of the CMS FSSTATE macro.

#### System action:

The file cannot be read and the command will not be executed.

#### User response:

See the description of the FSSTATE macro and its return codes in [z/VM: CMS Macros and Functions Reference](#) and take the appropriate action.

---

**FCXxxx126E      HELP file '*xxxxxx yyyyyy m*' not  
found**

#### Explanation:

While trying to locate the requested HELP information, Performance Toolkit found that the required HELP file *xxxxxx yyyyyy m* is not available, so no HELP information can be displayed.

#### System action:

Operation continues.

#### User response:

Access the minidisk that contains the HELP files before trying to display HELP information again.

---

**FCXxxx127E      No context help available for this  
display**

### Explanation

The HELP function PF-key was pressed to display context help text for the field pointed to by the cursor, but the selected display has not been prepared for showing context help information. This message can be generated when trying to select context help for user defined displays.

### System action

Operation continues.

### User response

Refer to the description of the original data fields.

---

**FCXxxx128E      Error while building window**

### Explanation

An internal software error occurred while the server machine tried to build a window for the requestor in an APPC/VM connection.

### System action

Operation continues, but the window is not built.

### User response

Contact IBM Service.

---

**FCXxxx130E      Invalid FCONTROL subcommand  
'*xxxxxxxx*' entered**

### Explanation

The string *xxxxxxxx* is not recognized as a valid FCONTROL subcommand.

### System action

The command is ignored.

### User response

For valid FCONTROL subcommands, see Chapter 1, “Performance Toolkit Subcommands,” on page 1. Then, enter the correct command.

---

**FCXxxx132E      Command / PF-key not valid for  
current mode**

### Explanation

The command entered, either directly or by pressing a PF-key, is not supported in Performance Toolkit's current operating mode.

### System action

The command is ignored.

### User response

See Chapter 1, “Performance Toolkit Subcommands,” on page 1 for valid command formats in Performance Toolkit's different operating modes, then re-enter the command.

---

**FCXxxx133E      Selected performance display not  
available for *xxx***



### Explanation

The selected performance display cannot be retrieved from a VM system of the type xxx, where xxx is the system type which you have defined in the FCONRMT SYSTEMS file for the remote system.

### System action

The command is ignored.

### User response

See "Remote Performance Data Retrieval and Display" in the *z/VM: Performance Toolkit Guide* for the retrieval commands valid for different systems.

---

<b>FCXxxx140E</b>	<b>PF-key undefined for current mode</b>
-------------------	--

---

### Explanation

The function key you just pressed has not been set or has been set for a different mode.

### System action

The function key interrupt is ignored and operation continues.

### User response

See "FCONTROL PFKEY" on page 66, which explains how to display the current function key settings and how to change them.

---

<b>FCXxxx151E</b>	<b>MSG/ACTMSG number missing</b>
-------------------	----------------------------------

---

### Explanation

A DELETE or REPLY command has been entered without a message line number.

### System action

The command is ignored.

### User response

See "DELETE" on page 104 or "REPLY" on page 115. Then, re-enter the command with the correct format.

---

<b>FCXxxx152E</b>	<b>No pending MSG/ACTMSG with number 'xx' on screen</b>
-------------------	---

---

### Explanation

A DELETE or REPLY command has been entered for message number 'xx'. The number does not correspond to the message number of one of the

message, warning or action message lines that are currently being displayed on the screen.

### System action

The command is ignored.

### User response

Re-enter the command with a valid message number.

---

<b>FCXxxx153E</b>	<b>No CP/CMS command entered with reply</b>
-------------------	---

---

### Explanation

The REPLY command expects a CP/CMS command to be entered following the message line number(s). No such command has been found.

### System action

The command is ignored.

### User response

Use the DELETE command for deleting messages which do not require any additional action.

---

<b>FCXxxx154E</b>	<b>Command is valid only when entered from within FCONX \$PROFILE</b>
-------------------	---

---

### Explanation

The command can be executed only during the Performance Toolkit initialization step and is treated as invalid after this initialization has completed.

### System action

The command is ignored.

### User response

If you want to issue this command, make sure it is included in the Performance Toolkit initialization file, FCONX \$PROFILE.

---

<b>FCXxxx155E</b>	<b>Invalid parameter entered</b>
-------------------	----------------------------------

---

### Explanation:

One of the parameters entered with the last Performance Toolkit subcommand is not valid (it is too long, for example).

### System action:

The command is ignored.

### User response:

For valid command formats, see [Chapter 1, “Performance Toolkit Subcommands,”](#) on page 1 (or enter HELP mode and look up the correct command format). Then, re-enter the command.

---

**FCXxxx156E      Parameter too short**

**Explanation:**  
One of the parameters entered with the last Performance Toolkit subcommand is shorter than the minimum abbreviation.

**System action:**  
The command is ignored.

**User response:**  
See the command description in [Chapter 1, “Performance Toolkit Subcommands,”](#) on page 1. Then, re-enter the command.

---

**FCXxxx157E      Value specified is below minimum**

**Explanation**  
A numeric value entered with the last subcommand is lower than the allowed minimum.

**System action**  
The command is ignored.

**User response:**  
See the command description in [Chapter 1, “Performance Toolkit Subcommands,”](#) on page 1. Then, re-enter the command.

---

**FCXxxx158E      Value specified exceeds maximum**

**Explanation**  
A numeric value entered with the last subcommand is higher than the allowed maximum.

**System action**  
The command is ignored.

**User response:**  
See the command description in [Chapter 1, “Performance Toolkit Subcommands,”](#) on page 1. Then, re-enter the command.

---

**FCXxxx170E      Invalid CP/CMS command entered**

**Explanation:**  
A command has been entered that was not a valid Performance Toolkit subcommand and could not be run as a CP or CMS command. This could be due to incorrect spelling of a subcommand (the command will then be run as a CP or CMS command, following the Performance Toolkit's standard search order).

**System action:**

The command is ignored.  
**User response:**  
Enter the correct command.

---

**FCXxxx175E      More than xx arguments/options for CMS command 'xxxxxxx'**

**Explanation:**  
A command (xxxxxxx) has been entered with more than *nn* arguments or options (including parentheses) that was interpreted as being a CMS command. The maximum number of arguments and options allowed by the Performance Toolkit for CMS commands is *nn*.

**System action:**  
The command is ignored.

**User response:**  
Enter a shorter version of the command or enter the command from native CMS.

---

**FCXxxx180A      Performance Toolkit abending ... {MONDCSS overlay}**

**Explanation**  
Performance Toolkit encountered an error that made further operation impossible. It will try to reset its interfaces and spool a dump to your virtual reader. In the message text, MONDCSS overlay indicates that the monitor segment was overlaid with a new set of data before Performance Toolkit finished processing the previous one. This message is usually accompanied by one of these messages: FCXxxx181, FCXxxx182, or FCXxxx183.  
  
This message is also sent to the maintenance user ID that is defined using the FC MAINTID command.

**System action**  
Performance Toolkit will exit after creating a dump of user storage.

**User response**  
Pass the dump on to your system programmer, along with any other supporting data such as the CP console log or additional information on actions that immediately preceded the ABEND. In the case of the MONDCSS overlay, see the explanation for message FCXxxx456E.

---

**FCXxxx181A      ABEND AT FCXccc +dddd : Program check X'pppp'**

**Explanation**  
Performance Toolkit encountered a program interrupt in the program FCXccc at hexadecimal displacement dddd with the code X'pppp', which made further



operation impossible. This message should be preceded by message FCXxxx180.

This message is also sent to the maintenance user ID that is defined using the FC MAINTID command.

**System action**

Performance Toolkit will exit to CMS after resetting its interfaces and creating a dump of user storage.

**User response**

See the explanation for message FCXxxxx180.

---

<b>FCXxxx182A</b>	<b>ABEND AT xxxxxxxx: Program check X'pppp'</b>
-------------------	---

---

**Explanation**

Performance Toolkit encountered a program interrupt at address xxxxxxxx with the code X'pppp', which made further operation impossible. The address points to some location out of the PERFKIT module. (xxxxxxx refers to the CMS nucleus code or the user program that is running under Performance Toolkit.) This message should be preceded by message FCXxxx180.

This message is also sent to the maintenance user ID that is defined using the FC MAINTID command.

**System action**

Performance Toolkit will exit to the CMS abend recovery routine after resetting its interfaces and creating a dump of user storage.

**User response**

See the explanation for message FCXxxxx180.

---

<b>FCXxxx183A</b>	<b>ABEND AT xxxxxxxx: CMS abend code X'cccc', reason X'rrrrrrr'</b>
-------------------	---

---

**Explanation**

Performance Toolkit encountered a CMS abend at address xxxxxxxx with the abend code X'cccc' and reason code X'rrrrrrr', which made further operation impossible. The address points to some location out of the PERFKIT module. (xxxxxxx refers to the CMS nucleus code or the user program that is running under Performance Toolkit.) This message should be preceded by the message FCXxxx180.

This message is also sent to the maintenance user ID that is defined using the FC MAINTID command.

**System action**

Performance Toolkit will exit to the CMS abend recovery routine after resetting its interfaces and creating a dump of user storage.

**User response**

See the explanation for message FCXxxxx180.

---

<b>FCXxxx190I</b>	<b>String 'xxxxxxx' not found</b>
-------------------	-----------------------------------

---

**Explanation**

The string 'xxxxxxx' specified with the 'locate' command could not be found in the remainder of the console log (search starts at current line + 1 for forward search or current line - 1 for backward search). Only the first eight characters of the string will be inserted into the error message.

**System action**

None.

**User response**

None.

---

<b>FCXxxx192I</b>	<b>Permanent perf. data collection active - re-enter "QUIT" to exit</b>
-------------------	---

---

**Explanation:**

A QUIT command has been entered in basic mode, requesting exit from the Performance Toolkit back to native CMS. Because permanent performance data collection is currently active, and because all data collected so far in the re-display buffer will be lost after quitting, you are asked to enter a second QUIT command to effectively end Performance Toolkit operation.

**System action:**

Operation continues.

**User response:**

Re-enter the QUIT command (or press PF3) to exit Performance Toolkit, or enter any other command if Performance Toolkit operation should not be ended.

---

<b>FCXxxx198A</b>	<b>GPR rr-rr : XXXXXXXX XXXXXXXX XXXXXXXX XXXXXXXX</b>
-------------------	--

---

**Explanation**

The Performance Toolkit encountered an internal processing error (such as a divide exception) and displayed the contents of the general purpose registers at that point. This diagnostic information may be needed for further problem analysis. Four FCXxxx198 messages are written to show all 16

registers. These messages should be preceded by the message FCXxxx199, or a similar message, with the explanation of the actual reason for this problem.

**System action**

Operation continues.

**User response**

Contact IBM Service.

---

<b>FCXxxx199A</b>	<b>Divide exception at FCXccc+dddd:</b> <b>XXXXXXXX XXXXXXXX /</b> <b>YYYYYYYY</b>
-------------------	--

---

**Explanation**

The program FCXccc at hex displacement dddd tried to perform a divide operation against dividend XXXXXXXX XXXXXXXX and divisor YYYYYYYY, but failed: either the divisor was zero or the quotient was too large to fit into one register. Possible reasons for this error are: unusual I/O device conditions (possible hardware failure or device reset), extremely large data counters due to increased CPU or I/O activity, or an abnormality in the software data processing algorithms. This message will be accompanied by the messages FCXxxx198 with the contents of the general purpose registers at the divide exception time.

This message is also sent to the maintenance user ID defined with the FC MAINTID command. Please note that only the first occurrence per day of this message will be sent to that user ID.

**System action**

Operation continues, but some data fields calculated by the program FCXccc might be incorrect.

**User response**

Contact IBM Service.

---

<b>FCXxxx200E</b>	<b>Command valid on 3278-5 screens only</b>
-------------------	---

---

**Explanation**

The 'MODEL' argument of the FCONTROL subcommand is valid on a 3278-5 screen only or on another display terminal that can emulate a 3278-5 (e.g. 3290). The terminal you are working on appears as a different terminal type to the Performance Toolkit.

**System action**

The command is ignored.

**User response**

Use a 3278-5 display terminal for entering the FCONTROL MODEL command.

---

<b>FCXxxx203E</b>	<b>Invalid screen mode 'nn' specified</b>
-------------------	---

---

**Explanation**

You specified an invalid screen model number 'nn' for the FCONTROL MODEL command.

**System action**

The command is ignored.

**User response**

See command description in “FCONTROL MODEL” on [page 50](#) and re-enter the correct command.

---

<b>FCXxxx210E</b>	<b>Invalid PF-key number entered</b>
-------------------	--------------------------------------

---

**Explanation**

Command text cannot be assigned to the PF-key number entered with the FCONTROL PFKEY command, because the number specified is not within the accepted range of 1 - 24.

**System action**

The command is ignored.

**User response**

See command description in “FCONTROL PFKEY” on [page 66](#) and re-enter the correct command.

---

<b>FCXxxx211E</b>	<b>Invalid PF-key function 'xxxxxxx' specified</b>
-------------------	--

---

**Explanation**

A 'FC PFKEY SET nn FUNCTION xxxxxxxx' command has been entered, but the specified function xxxxxxxx is invalid.

**System action**

The command is ignored.

**User response**

See the command description in “FCONTROL PFKEY” on [page 66](#) for valid functions which can be assigned, and re-enter the command.

---

<b>FCXxxx216E</b>	<b>Parameter xxxx contains non-numeric characters</b>
-------------------	---

---

**Explanation:**

The xxxx parameter for a Performance Toolkit subcommand is expected to be numeric, but has been found to contain non-numeric characters.

**System action:**

The command is ignored.

**User response:**

For the correct format of the subcommand, see [Chapter 1, “Performance Toolkit Subcommands,” on page 1](#). Then, re-enter the subcommand with the correct format.

---

**FCXxxx218E      Missing parameter**

**Explanation:**

The Performance Toolkit subcommand has a mandatory argument that was not entered with the command.

**System action:**

The command is ignored.

**User response:**

For the correct format of the subcommand, see [Chapter 1, “Performance Toolkit Subcommands,” on page 1](#). Then, re-enter the subcommand with the correct format.

---

**FCXxxx220E      Command exceeds buffer length of nn**

**Explanation**

The PFKEY command allows only command strings which fit into the corresponding buffer. The maximum length of 'nn' has been exceeded.

**System action**

The command is ignored.

**User response**

Assign a shorter form of the command to the PF-key (use abbreviations).

---

**FCXxxx230E      Invalid field name 'xxxxxxx' entered**

**Explanation:**

The first argument passed must be the designation of a system output message line type or of a Performance Toolkit field, but the string xxxxxxxx is neither.

**System action:**

The command is ignored.

**User response:**

For the correct format of the subcommand, see [Chapter 1, “Performance Toolkit Subcommands,” on](#)

[page 1](#). Then, re-enter the subcommand with the correct format.

---

**FCXxxx231E      Invalid color/highlighting 'xxxxxxx' - set to default**

**Explanation**

The color or highlighting definition xxxxxxxx is not valid.

**System action**

Color and highlighting of the respective field are set to the default values.

**User response**

For valid color and highlighting specifications, see [“FCONTROL COLOR” on page 18](#).

---

**FCXxxx240E      Number nnnn below minimum of mmmm, set to qqqq**

**Explanation:**

The number nnnn entered with a Performance Toolkit subcommand is smaller than the minimum number mmmm that is allowed.

**System action:**

The value qqqq is set.

**User response:**

Check to see if value mmmm is acceptable. If not, see the command description in [Chapter 1, “Performance Toolkit Subcommands,” on page 1](#) for the valid range of numbers. Then, re-enter the command.

---

**FCXxxx241E      Number nnnn exceeds maximum of mmmm, set to qqqq**

**Explanation**

The number nnnn entered with a Performance Toolkit subcommand is higher than the maximum number mmmm that is allowed.

**System action**

The value qqqq is set.

**User response**

Check to see if value mmmm is acceptable. If not, see the command description in [Chapter 1, “Performance Toolkit Subcommands,” on page 1](#) for the valid range of numbers. Then, re-enter the command.

---

**FCXxxx242E      Number nnnn exceeds screen size-1 of mmmm, set to mmmm**

**Explanation:**

The number *nnnn* entered with a Performance Toolkit subcommand exceeds the number *mmmm* of lines available in the scrollable output area minus 1 and is invalid. The number *mmmm* (which is also the maximum number allowed) has been used instead.

**System action:**

The maximum value *mmmm* is set.

**User response:**

Check to see if value *mmmm* is acceptable. If not, see the command description in [Chapter 1, “Performance Toolkit Subcommands,”](#) on page 1 for the valid range of numbers. Then, re-enter the command.

---

**FCXxxx243E**      **Number '*nnnn*' below minimum of '*mmmm*'**

---

**Explanation:**

The number *nnnn* entered with a Performance Toolkit subcommand is smaller than the minimum number *mmmm* that is allowed.

**System action:**

The command is ignored.

**User response:**

For the valid range of numbers, see the command description in [Chapter 1, “Performance Toolkit Subcommands,”](#) on page 1. Then, re-enter the command.

---

**FCXxxx244E**      **Number *nnnn* exceeds maximum of *mmmm***

---

**Explanation:**

The number *nnnn* entered with a Performance Toolkit subcommand is higher than the maximum number *mmmm* that is allowed.

**System action:**

The command is ignored.

**User response:**

For the valid range of numbers, see the command description in [Chapter 1, “Performance Toolkit Subcommands,”](#) on page 1. Then, re-enter the command.

---

**FCXxxx250E**      **Displacement *nnnn* exceeds maximum of 160**

---

**Explanation**

The displacement *nnnn* entered with the FCONTROL PROCESS subcommand is higher than the maximum displacement of 160 that is allowed.

**System action**

The command is ignored.

**User response**

For the correct format of the command, see [“FCONTROL PROCESS”](#) on page 70.

---

**FCXxxx251E**      **Length of string exceeds maximum of 20**

---

**Explanation**

A character string entered with the FCONTROL PROCESS subcommand is longer than the maximum length of 20 that is allowed.

**System action**

The command is ignored.

**User response**

For the correct format of the command, see [“FCONTROL PROCESS”](#) on page 70. Re-enter the command with a shorter character string. Use the AND operator to check for strings with a length of up to 40 characters (2 x 20).

---

**FCXxxx252E**      **Starting quote missing for character string**

---

**Explanation**

A character string (or an asterisk) must be entered with the FCONTROL PROCESS subcommand (positional parameter) and it must be enclosed in single quotation marks. The expected starting quotation mark could not be found.

**System action**

The command is ignored.

**User response**

For the correct format of the command, see [“FCONTROL PROCESS”](#) on page 70. Then, re-enter the command.

---

**FCXxxx253E**      **Ending quote missing after character string**

---

**Explanation**

The character string entered with the FCONTROL PROCESS subcommand must be enclosed in single quotation marks. The expected ending quotation mark could not be found.

### System action

The command is ignored.

### User response

For the correct format of the command, see [“FCONTROL PROCESS” on page 70](#). Then, re-enter the command.

---

**FCXxxx254E** Invalid line type xxxxxxxx entered

### Explanation

The line type xxxxxxxx is not recognized.

### System action

The command is ignored.

### User response

For the correct format of the command, see [“FCONTROL PROCESS” on page 70](#). Then, re-enter the command.

---

**FCXxxx255E** Invalid message type xxxxxxxx specified

### Explanation

The message type xxxxxxxx is not valid for re-routing message lines.

### System action

The command is ignored.

### User response

For the valid message line types, see [“FCONTROL PROCESS” on page 70](#). Then, re-enter the command.

---

**FCXxxx256E** File 'xxxxxxx EXEC \*' not found

### Explanation

The EXEC procedure you specified in the FCONTROL PROCESS or FCONTROL EMERGENC command could not be found and the complete command will be ignored.

### System action

The command is ignored.

### User response

Create file 'xxxxxxx EXEC' on any accessed disk before re-entering the command or, if an EXEC

already exists, re-enter the FCONTROL PROCESS or FCONTROL EMERGENC command with the correct file name.

---

**FCXxxx257E** Second string required for INSERT argument

### Explanation

The second data string is mandatory when using the INSERT argument of the FCONTROL PROCESS subcommand; otherwise, it does not make sense.

### System action

The command is ignored.

### User response

For the correct format of the command, see [“FCONTROL PROCESS” on page 70](#). Then, re-enter the command.

---

**FCXxxx260E** Invalid parameter xxxxxxxx specified

### Explanation:

The parameter xxxxxxxx entered with the last Performance Toolkit subcommand is not valid for the command or in that position.

### System action:

The command is ignored.

### User response

For the correct format of the command, see the command description in [Chapter 1, “Performance Toolkit Subcommands,” on page 1](#).

---

**FCXxxx262E** PROCESS block nnnn not found on ccccc chain

### Explanation

An attempt was made to delete block number nnnn from the ccccc chain of PROCESS blocks, but no block with that number could be found in that chain.

### System action

The command is ignored.

### User response

Verify the block number using the 'Query' argument of the FC PROCESS command and reenter the command with the correct block number.

---

**FCXxxx263E** SETEVENT block nnnn not found

## Explanation

An attempt was made to delete block number *nnnn* from the SETEVENT block chain, but no block with that number could be found.

## System action

The command is ignored.

## User response

Verify the block number using the "Query" argument of the FC SETEVENT command and reenter the command with the correct block number.

---

<b>FCXxxx265E</b>	<b>Invalid reset specification</b> 'cccc...cccc'
-------------------	---

---

## Explanation

An FC MONCOLL RESET command has been entered with an invalid argument *cccc...cccc*.

## System action

The command is ignored.

## User response

See the command description and re-issue the FC MONCOLL RESET command with a correct argument.

---

<b>FCXxxx266I</b>	<b>Unexpected processor <i>xx</i> counters reset detected Unexpected processor <i>xx</i> counters reset detected at <i>hh:mm:ss</i></b>
-------------------	---

---

## Explanation

The Performance Toolkit found a situation where a performance counter for a specified real processor was rolled back and started from zero. In monitor file scan and trend file scan modes, the original time of detection was appended.

## System action

The Performance Toolkit will handle the reset operation for the specified real processor performance counters. Operation continues.

## User response

Enable EVENT PROCESSOR monitor domain if permanent data collection is active, or ignore it otherwise.

---

<b>FCXxxx267I</b>	<b>YYYYYYYYY monitor data lost for XXXXXXXXX between <i>hh:mm:ss</i> and <i>hh:mm:ss</i></b>
-------------------	--

---

## Explanation

The Performance Toolkit found a situation where a performance counter for a specified monitor domain YYYYYYYY for object XXXXXXXX was rolled back or started from zero, and the original time interval of detection was appended.

## System action

Operation continues.

## User response

Enable EVENT and SAMPLE data for the named monitor domain if permanent data collection is active; otherwise ignore the message.

---

<b>FCXxxx300I</b>	<b>Message(s) pending for over <i>nn</i> min.</b>
-------------------	---

---

## Explanation

One or more messages on the screen have been left pending for more than the number '*nn*' of minutes specified with the FCONTROL MSGWAIT command.

## System action

Operation continues.

## User response

Perform the actions required by the messages and then delete them.

---

<b>FCXxxx301I</b>	<b>File xxxxxxxx yyyyyyyy saved on your A-disk</b>
-------------------	--

---

## Explanation

A command has been entered which resulted in a file being written to your A-disk, with a file name of xxxxxxxx and a file type of yyyyyyyy.

## System action

Operation continues.

## User response

None: this is an informational message only.

---

<b>FCXxxx305I</b>	<b>Destination ID xxxxxxxx is not logged on</b>
-------------------	---

---

### Explanation

The virtual machine that should receive re-routed output lines is not logged on and will not be able to see the lines.

### System action

Operation continues.

### User response

See that machine 'xxxxxxx' is logged on if it is an important part of the operating environment.

---

<b>FCXxxx306E</b>	<b>xxxxxxx is your own ID - cannot re-route to it!</b>
-------------------	--

### Explanation

Re-routing to your own ID could produce a loop and is forbidden.

### System action

The command is ignored.

### User response

Re-enter correct format of the FCONTROL PROCESS command.

---

<b>FCXxxx310I</b>	<b>Monitor data collection not activated</b>
-------------------	--

### Explanation

A 'MONITOR' command has been entered to switch into performance monitor mode, but CP monitor data collection has not previously been activated. This means that none of the automatically updated general performance data screens will be available as they all depend on CP monitor data collection.

### System action

Operation continues.

### User response

Activate collection of CP monitor data, using the appropriate 'FC MONCOLL ... ' command.

---

<b>FCXxxx311E</b>	<b>Command valid for users with priv. classes C/E only</b>
-------------------	--

### Explanation

DIAGNOSE instruction codes X'04' and X'90' are needed to allow performance data collection. This

requires CP command privilege class E on a standard z/VM system. Performance monitoring cannot be activated before your machine has been authorized for the necessary command privilege class.

### System action

The command is ignored.

### User response

Have your machine authorized for executing DIAGNOSEs X'04' and X'90' by adding the required CP command privilege class (usually class E) to the directory entry of your machine. Log off, log back on, restart Performance Toolkit, and then reissue the command.

**Caution: This authorization should be considered a security exposure if the use of the machine is not restricted to operating or system support personnel that already have similar privileges.**

---

<b>FCXxxx312E</b>	<b>Invalid limit 'nnnn' specified</b>
-------------------	---------------------------------------

### Explanation

The value entered with an 'FC LIMIT' command was higher than the acceptable maximum.

### System action

The command is ignored.

### User response

Re-enter the FCONTROL LIMIT command with a valid threshold.

---

<b>FCXxxx313E</b>	<b>Argument 'xxxxxxx' not valid for this CP level</b>
-------------------	---

### Explanation

The argument xxxxxxx entered with the last comma is meaningful only for specific CP levels. Your system is not on one of these levels. See the command description for the actual level required.

### System action

The command is ignored.

### User response

None.

---

<b>FCXxxx314E</b>	<b>Argument 'xxxx' valid for users with priv. class B only</b>
-------------------	--

## Explanation

The argument *xxxx* entered is accepted only when a certain CP command can be executed (e.g. 'Q TDISK'), which requires privilege class 'B' on a standard VM system. See the command description for the actual CP command required.

## System action

The argument is ignored.

## User response

Let your machine be authorized for the use of the CP command in question.

---

**FCXxxx315A**     *xxxx xxxx nn eeeeee LIMIT mm (Q1=q1 QX=qx)*

---

## Explanation

The current value of performance variable *xxxx xxxx* is *nn*, and this value has been higher (if *eeeeee* is 'exceeds') or lower (if *eeeeee* is 'lower') than the defined threshold *mm* for at least the number of minutes defined with the 'FCONTROL LIMIT' command. In parenthesis you will find either the average number of users in queues 1 and in queues 0/2/3 (values '*q1*' and '*qx*'), or the user with the highest CPU, I/O or paging load, or with the highest number of resident pages, depending on the threshold variable. If there is no highest entry then a string of '#####'s will be shown.

## System action

Operation continues.

## User response

Determine the reason for the uncharacteristic load and take the appropriate action.

---

**FCXxxx317A**     **User xxxxxxxx vvvv nn.n exceeded threshold mm.m for tt min.**

---

## Explanation

The current value of performance variable *vvvv* of user *xxxxxxx* is *nn.n*, and this value has been higher than the defined threshold *mm.m* for about *tt* minutes (the *tt* value represents actually the number of measuring intervals during which the threshold has been exceeded, but this usually corresponds closely to the number of minutes unless frequent re-calculations have been forced in shorter intervals).

## System action

Operation continues.

## User response

Check whether this user's high load impairs other users' performance. If so, and if this particular user's work does not have a high priority, you could lower the virtual machine's relative share in order to reduce its impact on general system performance.

---

**FCXxxx318A**     **Channel xx: nn% busy exceeds limit of mm%**

---

## Explanation

The measured channel busy value of *nn%* on CHPID *xx* exceeds the threshold of *mm%*.

## System action

Operation continues.

## User response

Determine the reason for the high channel load. Try to distribute part of the I/O equipment connected to this channel to another one if the load is consistently high.

---

**FCXxxx319A**     **Device xxxx: queue length .nn exceeds limit of .mm**

---

## Explanation

The average I/O request queue length during the last monitor sample interval was *.nn* for the I/O device *xxxx*; it exceeds the limit *.mm* set for this value. This indicates an I/O bottleneck which should be investigated.

## System action

Operation continues.

## User response

Determine the reason for the high device contention. If the problem occurs frequently, inform your systems programmer so that corrective action can be taken, such as re-organizing minidisks or system SPOOL/PAGE area allocation.

---

**FCXxxx320A**     **Device xxxx: resp. time nnnms exceeds limit mmmms (lll IO/s)**

---



### Explanation

The response time *nnnms* of device *xxxx* has exceeded the threshold *mmms* set with the FC LIMIT command for longer than the specified initial period, and the I/O rate was *lll* during the last monitor interval (the message will be generated only if there were at least 5 I/Os per second to that disk). Bad device response times can cause more general response time problems for the whole system and should be investigated. Among the possible causes are high device contention (as indicated by message 319), or high channel load leading to long disconnect times.

### System action

Operation continues.

### User response

Determine the reason for the high response time. If the problem occurs frequently, inform your systems programmer so that corrective action can be taken.

---

<b>FCXxxx321A</b>	<b>Less than <i>nn</i> paths to device(s) <i>xxxx-yyy</i></b>
-------------------	---

---

### Explanation

The number of active channel paths to device *xxxx*, or to the range of I/O devices with the device numbers *xxxx* to *yyyy*, is lower than the required minimum of *nn* paths.

Assuming that the threshold number *nn* has been correctly defined with the command 'FC MINPATHS ...', the message indicates that one or more of the normally available channel paths are currently offline. Possible causes are hardware malfunction, or a configuration change. This can lead to performance degradation for all I/O activity to the affected I/O devices.

### System action

Operation continues.

### User response

Use the DEVCONF display to determine which channel paths have gone offline and try to fix the problem. Inform your customer engineer about the situation if it looks like a hardware problem.

---

<b>FCXxxx323A</b>	<b>HMA page pool <i>nn%</i> usage exceeds threshold <i>mm%</i></b>
-------------------	--

---

### Explanation

The current HMA page pool usage of *nn%* exceeds the threshold of *mm%*.

### System action

Operation continues.

### User response

Increase the HMA storage size allocated for page pool by CP command DEFINE STORAGE CONFIG, Performance Toolkit subcommand FC HMA PAGE *nn*, or both. Then reinitialize Performance Toolkit.

---

<b>FCXxxx324A</b>	<b>HMA print buffer <i>nn%</i> usage exceeds threshold <i>mm%</i></b>
-------------------	---

---

### Explanation

The current HMA print buffer usage of *nn%* exceeds the threshold of *mm%*.

### System action

Operation continues.

### User response

Increase the HMA storage size allocated for print buffer by CP command DEFINE STORAGE CONFIG, Performance Toolkit subcommand FC HMA PAGE *nn*, or both. Then reinitialize Performance Toolkit.

---

<b>FCXxxx325A</b>	<b>Conventional storage &lt;2G <i>nn%</i> usage exceeds threshold <i>mm%</i></b>
-------------------	--

---

### Explanation

The current conventional virtual storage <2G usage of *nn%* exceeds the threshold of *mm%*.

### System action

Operation continues.

### User response

Increase the conventional virtual storage <2G by CP command DEFINE STORAGE CONFIG or DEFINE STORAGE. Then reinitialize Performance Toolkit.

---

<b>FCXxxx326A</b>	<b>Cache status '<i>cfd</i>' not matching '<i>xyz</i>' for <i>xxxx-yyy</i></b>
-------------------	--

---

### Explanation

Cache status monitoring has been activated for disk(s) with device number(s) *xxxx(-yyy)*, with an expected

cache status of xyz for general caching, DASD fast write and dual copy respectively. The current status information, in the same sequence, is *cf*d and does not match the expected status.

**System action**

Operation continues. The message will be repeated if the unexpected status persists.

**User response**

Issue the required CP commands to bring the current cache status back to the expected one. (Or, if the current status is found to be correct, issue an 'FC MONCACHE' command to update the expected status to the new required setting in order to prevent further misleading messages.)

---

<b>FCXxxx327A</b>	<b>Insufficient storage: ccccccc data may be incomplete</b>
-------------------	---

---

**Explanation**

The program tried to obtain free storage as a work buffer for building the ccccccc display, but no more free storage was available.

This message is also sent to the maintenance user ID defined with the FC MAINTID command.

**System action**

Operation continues, if possible, but data will be incomplete.

**User response**

Increase the virtual storage size of the PERFSVM machine.

---

<b>FCXxxx330E</b>	<b>User diagnose code 'nnn' not available</b>
-------------------	---

---

**Explanation**

The user diagnose code 'nnn' entered cannot be used in your virtual machine, probably because it has not been installed on your z/VM system.

**System action**

The command is ignored.

**User response**

Let your systems programmer install the user diagnose code that has been supplied with the program.

Since the CP console log can also be updated if your machine is allowed to issue the 'MSGNOH' command, you might also see that your virtual machine receives the necessary privilege class. The special user diagnose code is then not required for generating a complete CP console log.

---

<b>FCXxxx332E</b>	<b>Data length nn exceeds maximum data length mm</b>
-------------------	--

---

**Explanation**

This message can be shown in response to an FC USERHDR or FC USERBOTL command. The length *nn* of the data entered is longer than the available space *mm* in Performance Toolkit's header or bottom line, so the supplied string cannot be used in this form.

**System action**

The input will be truncated to the maximum length.

**User response**

Reenter a shorter version of user data or accept the truncated version.

---

<b>FCXxxx333A</b>	<b>Insufficient HMA storage for print buffer - exiting</b>
-------------------	--

---

**Explanation**

During generation of the performance reports, Performance Toolkit has detected the lack of a work print buffer residing in HMA storage.

If the emergency safeguard feature (ESF) has been activated previously, Performance Toolkit will initiate actions processing as defined using the FC EMERGENC ORDER command prior to program termination.

This message is also sent to the maintenance user ID that is defined using the FC MAINTID command.

**System action**

If the ESF is activated, Performance Toolkit will try to perform all required actions and then terminate. If the ESF is not activated, Performance Toolkit will terminate immediately.

**User response**

Increase the HMA storage size allocated for the print buffer using the CP command DEFINE STORAGE CONFIG, the Performance Toolkit command FC HMA PAGE *nn*, or both. Then, reinitialize Performance Toolkit.

---

<b>FCXxxx335E</b>	<b>Exclude table full, user xxxxxxxx not added</b>
-------------------	--

---

### Explanation

The exclude table has already been filled with user-IDs, and the userid xxxxxxxx could not be added to the table.

### System action

User xxxxxxxx and any others following it will not be added to the exclude table.

### User response

Re-define the exclude table size, using the command 'FC FORCEUSR EXCLMAX nn', and specify a number which is sufficient to hold the user-IDs of all machines which are to be excluded from FORCE processing.

---

<b>FCXxxx336I</b>	<b>Caution: No exclude list defined, ALL users will be forced</b>
-------------------	---

---

### Explanation

Since no user-IDs have been excluded from FORCE processing, the activation of loop/idle user monitoring will equally affect all virtual machines on the system, including service machines.

### System action

The command is accepted, and operation continues.

### User response

Make sure whether forcing all machines is acceptable, or specify the FORCE interval as 'O' to disable forcing. Otherwise, enter the required 'FC FORCEUSR EXCLUDE' command(s) for inserting the machines which are not to be forced into the 'exclude' table.

---

<b>FCXxxx337E</b>	<b>Required privilege class not available for command xxxxx xxxxx</b>
-------------------	---

---

### Explanation

User monitoring has been enabled (though the FC FORCEUSR command, the FC USRLIMIT command, or both), and Performance Toolkit tried to run a FORCE or SET SHARE command, as indicated by xxxx xxxx. However, the virtual machine where Performance Toolkit is running does not have the necessary CP command privilege class for running that command.

### System action

- If the FORCE command could not be run: Machines will not be forced. All looping and idle user monitoring is disabled.
- If the SET SHARE command could not be run: Users' relative SHARE will not be automatically adapted when they exceed their resource consumption thresholds, but alert messages will still be generated.

### User response

Authorize the machine for the required CP command privilege class (usually class A, which requires a directory change). The machine will then be able to run the command after the next logon.

---

<b>FCXxxx338A</b>	<b>User xxxxxxxx in yyyy yyyy during last nn minutes</b>
-------------------	--

---

### Explanation

User xxxxxxxx has been found to be in the state indicated by yyyy yyyy for at least nn minutes by the user state monitoring logic. The state yyyy yyyy can be 'CPU loop', 'I/O loop', 'WSS loop' or 'IDLE wait', and the value nn is equal to, or larger than the 'warning' limit specified with the 'FC FORCEUSR' command.

### System action

Operation continues.

### User response

Check whether the machine's status represents a problem, and take the appropriate action.

---

<b>FCXxxx339M</b>	<b>User xxxxxxxx forced after nn minutes yyyy yyyy</b>
-------------------	--

---

### Explanation

User xxxxxxxx has been found to be in the state indicated by yyyy yyyy for at least nn minutes by the user state monitoring logic. The state yyyy yyyy can be 'CPU loop', 'I/O loop', 'WSS loop' or 'IDLE wait'. The value nn is equal to, or larger than the 'FORCE' limit specified with the 'FC FORCEUSR' command, so the machine has been forced.

### System action

The user is forced, and operation continues.

### User response

None.

---

**FCXxxx340S      More than *nn* words in VMTAPE message**

---

**Explanation**

Analysis of VMTAPE messages starts with tokenization of the messages to be analyzed. The routine has been written for tokenizing up to *nn* words only, and this limit has been exceeded.

**System action**

Further analysis of the message is skipped and operation continues.

**User response**

Inform systems programmer. If you use the original VMTAPE message format, inform IBM Service.

---

**FCXxxx343S      Error while analyzing message VMTXXXnnn**

---

**Explanation:**

The format of message VMTXXXnnn is not as expected by the the Performance Toolkit subroutine that tried to analyze it, probably because a different VMTAPE release is being used.

**System action:**

Further analysis of the message is skipped and operation continues.

**User response:**

Inform your system programmer. If you use the original VMTAPE message format, inform the IBM Support Center.

---

**FCXxxx344S      No more MSGQBLOKs available**

---

**Explanation:**

The Performance Toolkit for VMTAPE message analysis routines have been written for handling up to 16 concurrently-active tape units. More than 16 tape units are now active, or cleanup of some previously-used control blocks is not complete.

**System action:**

Further processing of the VMTAPE message is skipped and operation continues.

**User response:**

Inform your system programmer. You could also re-initialize Performance Toolkit if the actual number of active tape units is lower than 16.

---

**FCXxxx350E      User 'xxxxxxx' not logged on**

---

**Explanation**

The user-identification specified with the 'USER xxxxxxxx' or with the 'STORAGE DISPLAY xxxxxxxx'

command cannot be found, i.e. it is not currently logged on, or it is in LOGON or LOGOFF processing so that no valid storage data can be extracted.

**System action**

Monitoring switches to the initial 'CPU'-screen.

**User response**

Re-enter the 'USER' sub-command with a valid user-ID.

---

**FCXxxx351E      Device 'xxxx' not found**

---

**Explanation**

No real device has been generated with the device number 'xxxx' specified with the 'DEVICE xxxx' or 'I/O xxxx' command.

**System action**

Monitoring switches to the initial 'CPU'-screen.

**User response**

Re-enter the command with a valid real device number.

---

**FCXxxx353M      Insufficient storage for work area**

---

**Explanation:**

Some Performance Toolkit functions require additional storage for work areas. The amount of storage needed may be very large, especially for performance monitoring on large systems (depends on the actual configuration and the number of users on the system being monitored). Not all of the requested storage has been obtained.

**System action:**

This depends on the block that was not received. Performance monitoring might not be possible at all, or only part of the monitoring might be active.

**User response:**

Increase the size of your virtual machine, re-IPL ZCMS, and re-initialize Performance Toolkit.

---

**FCXxxx354A      Device 'xxxx' waiting for interrupt**

---

**Explanation**

The disk with real device number 'xxxx' has been waiting for an interrupt for at least one measuring interval. This may indicate a problem.

**System action**

Operation continues.

## User response

Check whether there is really a problem with the disk. (The logic used checks whether an IORBK remained queued even though there was no I/O activity to the disk. This can happen also if suspend/resume logic is used by the channel programs.)

---

**FCXxxx355M      Interrupt(s) missing for device(s)  
                         of user xxxxxxxx**

## Explanation

One or more virtual devices of user xxxxxxxx have been waiting for an interrupt for at least one measuring interval. This probably indicates a problem in the case of DASD or GRAF-type devices. It may happen occasionally during normal operation on tape units, and on a CTC may just indicate that the other side is not active.

## System action

Operation continues.

## User response

Force re-calculation of performance data by pressing the 'ENTER'-key to see whether the problem persists. Use your judgement to determine whether there really is a problem and where the problem is (e.g. other end of a CTC connection).

---

**FCXxxx356A      RESERVE pending for device 'xxxx'  
                         and no I/O activity**

## Explanation

Device xxxx was found to have remained in a 'reserved' state during the last measuring interval, without any I/O activity on the device. This usually indicates a failure in the application which issued the RESERVE command, and it will prevent applications on other systems to access the shared DASD.

## System action

Operation continues.

## User response

Check why the RESERVE remained pending (probably the application which issued it just ABENDED). Re-start the application, if necessary, to clear the RESERVED state.

---

**FCXxxx357I      No measurement block data  
                         available for device**

## Explanation

The detailed I/O device display has been selected for an I/O device. The pointer to the measurement block in the RDEV control block was found to be zero.

## System action

Data based on measurement block data cannot be shown. This includes the I/O rate and device service time data, and also some other information (e.g. device response time) which is partially based on such data.

## User response

None.

---

**FCXxxx358E      Command valid only if permanent  
                         data collection is active**

## Explanation

A command has been entered which cannot be executed unless permanent data collection has previously been activated by entering the command 'FC MONCOLL ON'.

## System action

Operation continues.

## User response

Enter the command 'FC MONCOLL ON' to start permanent performance data collection, then re-enter the command(s) which caused the error message.

---

**FCXxxx359E      Persistent data collect error for  
                         user ccccccc**

## Explanation:

The User Performance Details display was selected for user ccccccc, but a persistent error was found while trying to collect data for the user. Performance data is collected using Diagnose X'04' for copying data from CP control blocks. Because this requires many such diagnose calls, they cannot all be executed at the same time, and it is possible that the contents of the original CP control blocks change while collecting data, because the virtual I/O configuration of the monitored machine changed. The Performance Toolkit will detect some such changes, and will try to recover by restarting the data collection. It will give up trying after the third unsuccessful attempt.

## System action:

The selected user's data cannot be shown.

## User response:

Find out what happened to the monitored machine (an I/O configuration change, for example) and retry monitoring it when its configuration should again be stable.

---

**FCXxxx360E**      **'xxxxxxx' is not a valid monitor sub-command**

---

**Explanation**

The command xxxxxxxx entered is not recognized as a valid Performance Toolkit performance monitor subcommand, nor is it one of the general Performance Toolkit subcommands. No other commands are accepted while in performance monitoring mode.

**System action**

Operation continues.

**User response**

Enter a valid command. For a description of all performance monitor subcommands, see [“Performance Monitor Mode Subcommands” on page 129.](#)

---

**FCXxxx361E**      **Selection not valid for Monitor Scan mode**

---

**Explanation**

A performance monitor screen was selected which is based on data obtained from the system's control blocks. While in monitor scan mode, you are viewing (old) performance data for another period, and possible even for another system, so it makes no sense to retrieve and display *current* performance data for *this* system: the information would not correlate with the CP monitor data being scanned.

**System action**

The selection is ignored.

**User response**

None.

---

**FCXxxx362I**      **Storage utilization monitoring stopped - 255 samples taken**

---

**Explanation**

The Performance Toolkit storage utilization monitoring can only handle up to 255 samples. This limit has been reached, and no further samples will be taken for this user.

**System action**

Data collection for user storage analysis is stopped.

**User response**

If you want to continue monitoring the same user, you can force a reset the accumulated data by starting data collection for an other user. You can then restart monitoring for the original user; any user switch causes the data collected so far to be reset.

---

**FCXxxx364E**      **'xxxxxxx' is not a valid sub-command for this display**

---

**Explanation**

The command xxxxxxxx entered is a performance monitor subcommand which is intended for use with another monitor display only.

**System action**

The command is ignored and operation continues.

**User response**

Enter a valid command.

---

**FCXxxx365E**      **Error while converting PERFLOG data 'xxxx' (field 'cccc')**

---

**Explanation:**

Performance Toolkit found an unexpected data format (xxxx) when trying to convert the field cccc EBCDIC to binary format. This is a Performance Toolkit error, unless you have modified the *mmddy* PERFLOG file yourself which Performance Toolkit tried to convert.

**System action:**

A value of 0 is assumed for the field, and operation continues.

**User response:**

Rename the PERFLOG file to some other filetype to prevent its further use by Performance Toolkit and inform the author of the program.

---

**FCXxxx366I**      **User storage exceeds nn MB, remainder cannot be mapped**

---

**Explanation:**

The Performance Toolkit storage utilization monitoring code is set up for analyzing up to nnMB. Page residency for user storage above this limit cannot be analyzed and displayed.

**System action:**

Operation continues.

**User response:**

None.

---

**FCXxxx367E**      **No mapping possible: storage is much larger than *nn* MB**

**Explanation:**

The Performance Toolkit storage utilization monitoring code is set up for analyzing up to *nn*MB. The virtual machine's storage is much larger than this limit and displaying just the first *nn* MB does not make sense.

**System action:**

User storage monitoring is stopped.

**User response:**

None.

---

**FCXxxx368I**      **Insufficient storage, data may be incomplete**

**Explanation:**

The Performance Toolkit performance monitoring code can require a amount of working storage, especially on systems with many logged-on users. It will try to continue operating even when not all its storage demands have been met, but some of the data might be incomplete (data was not collected for all users on the system, for example).

**System action:**

Operation continues.

**User response:**

Increase the machine's storage size in the directory and then restart the program.

---

**FCXxxx369E**      **64-bit address found: can only handle 31-bit addresses**

**Explanation**

The current implementation of storage utilization monitoring is based on both segment and page tables residing below the 2GB line. However, one of the pointers found in storage appears to be a 64-bit address above 2GB and the code cannot yet handle it.

**System action**

User storage monitoring is stopped.

**User response**

Contact IBM Service.

---

**FCXxxx371I**      **AVERAGE data shown temporarily due to print request**

**Explanation**

An automatic data reset and print request is being executed (defined by 'FC MONCOLL RESET' command). This type of print request implies 'average'

data, and the performance monitor has therefore been switched temporarily to calculate average data before being reset. It will automatically be set to display 'current' data again after the reset.

**System action**

Operation continues.

**User response**

None.

---

**FCXxxx372E**      ***rrrr* cannot be shown - not available on this level of VM**

**Explanation**

The last command requested a screen (*rrrr*) with data that is not available with the current level of z/VM. This particular type of data might be available only when running Performance Toolkit against data from a later release of z/VM.

**System action**

Operation continues.

**User response**

See the description of the screen (*rrrr*) requested. It should state any special prerequisites.

---

**FCXxxx373E**      **Invalid report name *rrrrrrrr* specified**

**Explanation**

The requested report *rrrrrrrr* was not recognized.

**System action**

Operation continues.

**User response**

See the documentation for a description of available performance reports, and of the commands to be used for selecting them.

---

**FCXxxx374E**      **INTERIM period not available for *cccccc ccccccc***

**Explanation**

The specification *cccccc ccccccc* in the REPORTS control file calls for an 'interim' period report. However, this type of report is not available with the current level of the program.

## System action

Operation continues.

## User response

See Appendix C, “Printed Reports,” on page 979 for a list of the available performance reports, and the commands to be used for selecting them.

---

**FCXxxx375E      No data found for current selection**

## Explanation

The history data to be analyzed does not contain records which meet the criteria specified with the last command.

## System action

Operation continues.

## User response

Check the time span covered by the available history data and adapt the command accordingly.

---

**FCXxxx376E      File *fn ft fm*: Invalid format for performance log file**

## Explanation

File *fn ft fm* was selected as input for creating graphics, but the file format is not correct (record length does not match the length expected for a detailed or summary performance log file).

## System action

The command is aborted.

## User response

Check the file identification entered. Make sure the original record length of PERFLOG, RMONLOG and PERFHIST files is not changed when copying the data to an archive disk.

---

**FCXxxx377E      Return code *rc* for ccccccc with UNPACK option**

## Explanation:

As the result of a previous command, Performance Toolkit attempted to unpack a packed performance data log file which was needed as input. However, calling ccccccc for unpacking the file resulted in a return code of *rc*.

## System action:

The command is aborted.

## User response:

See the documentation of program ccccccc for the meaning of the return code and try to fix the problem before re-executing the command.

---

**FCXxxx378E      TERSE MODULE not found, cannot unpack file**

## Explanation

A command required data from a packed performance data log file as input. The program determined that the performance data log file had been packed using the TERSE module, but cannot find this module which it needs for unpacking.

## System action

The command is aborted.

## User response

Execute the necessary LINK and ACCESS commands in the server machine to give it access to the TERSE module, then re-execute the command.

---

**FCXxxx380I      Entry point 'ccccccc' at address X'xxxxxxxx'**

## Explanation

This is the reply to a 'LOCATE' command in performance monitor mode. It indicates the address X'xxxxxxxx' in real storage at which entry point ccccccc is located.

## System action

Operation continues.

## User response

None.

---

**FCXxxx381E      Entry point 'xxxxxxxx' not found**

## Explanation

The entry point xxxxxxxx could not be located by diagnose X'90'. The message will be shown when an entry point specified with a 'LOCATE' command could not be found or during initialization, if one of the CP entry points needed for performance monitoring could not be located.

## System action

Operation continues. If the message was generated during initialization, a flag will be set to disable



monitoring for the remainder of the current session in order to avoid ABENDs.

**User response**

For the 'LOCATE' command, make sure the string xxxxxxxx was entered correctly. If the message was generated during initialization, the probable cause is a CP change (a module split, for example) that renamed an entry point. Inform IBM Service.

---

<b>FCXxxx382I</b>	<b>Monitoring disabled (CP entry point missing)</b>
-------------------	---

---

**Explanation**

This message is usually preceded by message 381 during initialization: a CP entry point which is required for the monitor's data collection process could not be located. All monitoring has been disabled to avoid ABENDs.

**System action**

Operation continues.

**User response**

Resolve the problem indicated by message 381.

---

<b>FCXxxx383E</b>	<b>Invalid HEX specification 'xx' found</b>
-------------------	---

---

**Explanation**

The CHAR argument of the PLOT or PLOTSUM command has been entered with an apostrophe (') in the string specification. This indicates that all characters up to the next apostrophe are to be interpreted as character pairs with the hexadecimal representation of one of the desired plotting characters, but the pair xx either contains non-hex characters, or a value lower than X'40' has been defined (not valid as display character).

**System action**

The invalid hexadecimal specification and any other characters following it are ignored, and the default plotting characters are used instead.

**User response**

Re-enter the command with a valid specification.

---

<b>FCXxxx384E</b>	<b>Time stamp hh:mm not found</b>
-------------------	-----------------------------------

---

**Explanation**

A PLOTxxx or a GRAPHxxx command has been entered with the 'FROM hh:mm' argument The starting time hh:mm specified could not be found in any of the available data records.

**System action**

The oldest values will be plotted (first few entries in redisplay buffer).

**User response**

Re-enter the command with a valid start time.

---

<b>FCXxxx385E</b>	<b>Entry point 'ccccccc' not resident</b>
-------------------	---

---

**Explanation:**

The entry point ccccccc was found to be not resident. The message will be shown when an entry point specified with a LOCATE command is not resident or during initialization if one of the CP entry points needed for performance monitoring is found to be not resident. The latter case indicates a logic error: a CP module was probably made pageable in a new VM release, and your copy of Performance Toolkit is not yet adapted to that change.

**System action:**

Operation continues.

**User response:**

Obtain a new PERFKIT module, adapted to your VM level, if the message was generated during initialization.

---

<b>FCXxxx386E</b>	<b>Error nnnn for 'GLOBAL TXTLIB ADMNLIB ADMGLIB ADMHLIB' command</b>
-------------------	---

---

**Explanation**

The command GLOBAL TXTLIB ADMNLIB ADMGLIB ADMHLIB is executed before loading the graphics interface modules for preparing GDDM output. The error code nnnn is the return code received while executing this command.

**System action**

The CPU display will be shown and operation continues.

**User response**

Determine the reason why the command GLOBAL TXTLIB ADMNLIB ADMGLIB ADMHLIB could not be executed (for example, GDDM not installed or disk

not accessed). Try selecting the graphics screen again when the problem has been resolved.

---

**FCXxxx387E      Error *nnnn* loading 'FCXPGD TEXT'**

---

**Explanation**

The command 'LOAD FCXPGD' is executed while preparing the graphics display. The FCXPGD TEXT file, or another TEXT file that was to be loaded with FCXPGD, could not be loaded properly. The error code *nnnn* is the return code received while executing this command.

**System action**

The CPU display will be shown and operation continues.

**User response**

Determine the reason why this command could not be executed (e.g. FCXPGD and/or FCXPGU TEXT file not on any of the accessed disks). Switch back to basic mode (press PF-key 3) and look for CMS error messages generated while trying to load FCXPGD/FCXPGU. Try selecting the graphics display again when the problem has been resolved.

---

**FCXxxx388E      Error code *nnnn* from *xxxxxxx***

---

**Explanation:**

Module FCXPGD or FCXPGU which prepare and execute GDDM graphics output by calling the appropriate GDDM routines has received an unexpected error code of *nnnn* from GDDM routine *xxxx*.

**System action:**

The CPU display will be shown and operation continues.

**User response:**

Refer to the GDDM messages manual for the meaning of the GDDM error code. The code *nnnn* corresponds to the GDDM message number. Look also for GDDM messages which may have been generated (leave performance monitor mode, you can only see the messages while Performance Toolkit is in basic mode).

---

**FCXxxx391E      CP MONITOR data analysis not activated, data not available**

---

**Explanation:**

A performance monitor subcommand was entered which called for data which be supplied only when the CP MONITOR interface Performance Toolkit has been activated, by means of the FC MONCOLL CPMON command. This interface is not currently active, and the selected data cannot be shown.

**System action:**

None.

**User response:**

Activate data extraction from CP MONITOR data by means of the FC MONCOLL CPMON command, then retry the command that resulted in the error message.

---

**FCXxxx392E      No control unit cache data available**

---

**Explanation**

A 'CACHDBSE' command has been entered and, although CP MONITOR data extraction has been activated, no cache data are available for display. Possible reasons are:

- The CACHDBSE command was entered too soon after starting CPMON data collection, i.e. before the end of the next CP MONITOR data collect interval.
- The I/O domain is not enabled for the cached disks
- No control unit with cache exists in the system's I/O configuration

**System action**

CACHE data cannot be shown, and the general CPU overview screen is shown instead.

**User response**

Make sure the I/O domain has been enabled for all cached disks, and MONITOR SAMPLE is active. Wait, if necessary, for the next CP MONITOR data collect interval.

---

**FCXxxx396E      No shared file system data available**

---

**Explanation**

An 'SFS' command has been entered and, although CP MONITOR data extraction has been activated, no shared file system data are available for display. Possible reasons are:

- The SFS command was entered too soon after starting CPMON data collection, i.e. before the end of the next CP MONITOR data collect interval.
- The APPLDATA domain is not enabled for the shared file system servers
- No shared file system servers are logged on to the system

**System action**

SFS data cannot be shown, and the general CPU overview screen is displayed instead.

### User response

Make sure the APPLDATA domain is enabled for all shared file system servers, and MONITOR SAMPLE is active. Wait, if necessary, for the next CP MONITOR data collect interval.

---

**FCXxxx397E**      **No shared file system data available for user xxxxxxxx**

### Explanation

An 'SFS *userid*' command has been entered and, although CP MONITOR data extraction has been activated, and shared file system data found, no shared file system data are available for the user-ID which you selected. Possible reasons are:

- The APPLDATA domain has been enabled only for some SFS servers, but not all of them
- The selected virtual machine is not working as an SFS server

### System action

SFS data for the selected user cannot be shown, and the general CPU screen is displayed instead.

### User response

Check whether the selected virtual machine is currently working as an SFS server and, if so, make sure the APPLDATA domain is enabled to collect data for this server.

---

**FCXxxx398E**      **No matching performance history files found**

### Explanation

A HISTDATA command was entered, but no files matched the specifications of the command or the default ones.

### System action

The selection is ignored.

### User response

None.

---

**FCXxxx399E**      **No LPAR data available**

### Explanation

An 'LPAR' command has been entered, but no CP monitor data have been collected so far with LPAR information. Possible reasons are:

- The system is running second level under control of another z/VM system
- No new SAMPLE interval has elapsed since activating the Performance Toolkit performance monitoring function

### System action

LPAR data cannot be shown, and the initial menu will be displayed instead.

### User response

Make sure that CP MONITOR data collection is active, and the Performance Toolkit is running in the appropriate system environment.

---

**FCXxxx400E**      **'fn ft fm' not found - request cancelled**

### Explanation

A command has been issued that requires a control file with the file identification *fn ft fm* as input, but the required file could not be found or contained only comments.

### System action

The request is cancelled.

### User response

Ask the owner of the machine to create the control file.

---

**FCXxxx401E**      **You have not been authorized for retrieving data**

### Explanation

Your address (RSCS node-ID and user-ID) could not be found in file 'FCONRMT AUTHORIZ' in the machine where you sent the remote performance data retrieval request. Requests will be executed only when you have been authorized by your address in this file.

### System action

The request is cancelled, no performance data will be returned.

### User response

Contact the owner of the machine and let yourself be authorized for remote performance data retrieval, then re-issue the retrieval request.

---

**FCXxxx402E**      **No perm. perf. data collection - retrieval not possible**

## Explanation

Remote performance data retrieval is possible only from machines where Performance Toolkit has been set up for permanent performance data collection. The machine where you sent your retrieval request is not set up for this.

## System action

The request is cancelled. No performance data will be returned.

## User response

Ask the owner of the machine to let performance data be collected permanently (the FC MONCOLL ON command - see [“FCONTROL MONCOLL” on page 53](#)), then issue the request again.

---

<b>FCXxxx403E</b>	<b>You have not been authorized for executing commands</b>
-------------------	--

---

## Explanation

Your address (RSCS node-ID and user-ID) has not been included in file 'FCONRMT AUTHORIZ' with the 'CMD' argument in the machine where you sent the CP or CMS command. CP and CMS commands will be executed only when you have received special authorization for executing commands.

## System action

The request is cancelled, the command will not be executed.

## User response

Contact the owner of the machine and let yourself be authorized for executing commands, then re-issue the retrieval request.

---

<b>FCXxxx404E</b>	<b>You have not been authorized for submitting user exceptions</b>
-------------------	--

---

## Explanation

Your address (user ID and RSCS node ID) has not been included in the FCONRMT AUTHORIZ file with the EXCPMSG argument in the machine where you sent the FCONRMU exception message. User-generated exception information will be accepted, and included in the consolidated exception log file, only after you have received the required authorization.

## System action

The data sent are not treated as user exception information, but are displayed as an ordinary CP message.

## User response

Contact the owner of the machine to get authorized to submit user-generated exception data.

---

<b>FCXxxx405E</b>	<b>Authorization type missing for user 'uuuu' at 'nnnn'</b>
-------------------	---

---

## Explanation

The entry for the user with userid *uuuu* on system *nnnn* in the authorization file 'FCONRMT AUTHORIZ' does not contain the mandatory specification of the authorization type to be given.

## System action

The invalid record is ignored, and processing continues with the next record in the FCONRMT AUTHORIZ file.

## User response

Adapt the corresponding record in file FCONRMT AUTHORIZ and then re-initialize. See "Implementing Remote Performance" in the *z/VM: Performance Toolkit Guide* for a complete description of the record layout.

---

<b>FCXxxx406E</b>	<b>Invalid authorization 'xxxx' for user uuuu at nnnn</b>
-------------------	---

---

## Explanation

The authorization type *xxxx* specified for the user with userid *uuuu* on system *nnnn* in the authorization file 'FCONRMT AUTHORIZ' is not valid.

## System action

The invalid record is ignored, and processing continues with the next record in the FCONRMT AUTHORIZ file.

## User response

Adapt the corresponding record in file FCONRMT AUTHORIZ and then re-initialize. See "Implementing Remote Performance" in the *z/VM: Performance Toolkit Guide* for a complete description of the record layout.

---

<b>FCXxxx407E</b>	<b>Missing Node-ID and/or User-ID</b>
-------------------	---------------------------------------

---

## Explanation

No node-ID and/or user-ID specification has been found in one of the records of the authorization file 'FCONRMT AUTHORIZ'. Both the node-ID and user-ID specifications, and a specification of the authorization type, are mandatory.

## System action

The invalid record is ignored, and processing continues with the next record in the FCONRMT AUTHORIZ file.

## User response

Adapt the corresponding record in file FCONRMT AUTHORIZ and then re-initialize. See "Implementing Remote Performance" in the *z/VM: Performance Toolkit Guide* for a complete description of the record layout.

---

<b>FCXxxx408E</b>	<b>Invalid Node-ID or User-ID</b> 'xxxxxxx'
-------------------	--

---

## Explanation

The node-ID or user-ID name beginning with the characters xxxxxxxx is invalid (contains invalid characters or is more than eight characters long), and cannot be accepted.

## System action

The invalid record is ignored, and processing continues with the next record in the file FCONRMT AUTHORIZ or FCONRMT SYSTEMS.

## User response

Adapt the corresponding record in the file FCONRMT AUTHORIZ or FCONRMT SYSTEMS and then re-initialize. See "Implementing Remote Performance" in the *z/VM: Performance Toolkit Guide* for a complete description of the record layout.

---

<b>FCXxxx409E</b>	<b>Intermediate server not authorized for S&amp;F operation</b>
-------------------	---

---

## Explanation

An attempt was detected to establish an APPC/VM connection, posing as an intermediate store and forward server, but the virtual machine has not been defined as an S&F server in the FCONRMT AUTHORIZ file.

## System action

Store and forward requests are not allowed.

## User response

Adapt the FCONRMT AUTHORIZ file and include the required S&FSERV definition if the virtual machine is to be set up as a store and forward server for APPC/VM remote data retrieval.

---

<b>FCXxxx410A</b>	<b>Unauth. cccccc request by uuuu at nnnn</b>
-------------------	---

---

## Explanation

User uuuu on VM system nnnn entered a remote command, but is not authorized for cccccc requests.

## System action

The command is not executed.

## User response

Refrain from using unauthorized requests, or contact the administrator to have your entry in the FCONRMT AUTHORIZ file updated and the cccccc authorization added.

---

<b>FCXxxx411E</b>	<b>CP/CMS commands not accepted via Internet interface</b>
-------------------	--

---

## Explanation:

Current HTTP protocol is based on stateless transactions where a user's connection is maintained only for the duration of a single interaction. Some logic has been implemented in Performance Toolkit to allow maintaining something similar to an ongoing session, but it is not reliable enough to really guarantee that the requestor is still the same one which originally provided a logon password. Allowing execution of CP or CMS commands by users whose identity cannot be guaranteed would introduce a severe security exposure, so we inhibit use of such commands.

## System action:

None.

## User response:

None.

---

<b>FCXxxx412E</b>	<b>Invalid selection(s)</b>
-------------------	-----------------------------

---

## Explanation

An invalid selection character has been entered.

## System action

The selection is ignored.

## User response

Enter the correct selection character. See explanation in the message line above the command line for valid selections.

---

**FCXxxx413I**      **Screen definition not found - backlevel module**

## Explanation

While processing remote performance data from another system, the program could not locate the screen definition that describes the display layout for the remote data in its internal table, most likely because performance data were retrieved from a remote Performance Toolkit module that is at a higher maintenance level.

## System action

The remote data are displayed, but scrolling operations and HELP calls might yield unexpected results.

## User response

Have the Performance Toolkit module on the local system upgraded to the same level as the Performance Toolkit module on the remote systems.

---

**FCXxxx414E**      **Command chaining via #-character not allowed in remote mode**

## Explanation:

Logical line-end characters are usually processed within the requesting the Performance Toolkit module, which then passes the logical command on separately for validation/execution. A command string was submitted with an embedded # logical line-end character, indicating either a logic error in the requesting Performance Toolkit module or an attempt to submit such a string from another program, posing as a Performance Toolkit client.

## System action:

The command is not executed.

## User response:

Inform the IBM Support Center if the command string was entered from within a Performance Toolkit client, or refrain from using command chaining if another client program was involved.

---

**FCXxxx415E**      **You have not been authorized for executing CP queries**

## Explanation

A CP QUERY or INDICATE command has been entered in a remote performance data retrieval session, but the user is not authorized to run such commands or the user is connected via the web interface and the web interface has not been given the CPQRY authorization.

## System action

The command is not executed and the connection to the system is severed.

## User response

Contact the system administrator and ask for CPQRY authorization if you have a business need for it. Be aware, however, that a policy may be in place that does not allow web clients to be given such an authorization, due to the nature of the connection or user validation.

---

**FCXxxx416E**      **Invalid ':systemid:' specification found**

## Explanation

An entry was found in the file that looked like an attempt at a :systemid: specification, with a leading ':', but it cannot be valid because either the trailing ':' is missing, or because the 'systemid' is not 1-8 characters long.

## System action

The record is ignored.

## User response

Check the control file with the incorrect statement and fix the invalid entry.

---

**FCXxxx417E**      **:EXCL: argument out of sequence**

## Explanation

An :EXCL: argument was found, but was not the first argument in a new 'group' of systems. **All or none** of the systems in a group must be excluded, a mix is not possible.

## System action

The record is ignored.

## User response

Check the control file with the incorrect statement and fix the invalid entry.

---

**FCXxxx419E**      **Intermediate server not authorized for S&F operation**

---

### Explanation

An attempt to establish an APPC/VM connection posing as an intermediate store and forward server was detected, but the virtual machine has not been defined as an S&F server in the FCONRMT AUTHORIZ file.

### System action

Store and forward requests are not allowed.

### User response

Update the FCONRMT AUTHORIZ file on the client machine to include the required S&FSERV definition.

---

**FCXxxx420E**      **No history data available for node xxxxxxxx**

---

### Explanation

An 'h' selection has been entered for system xxxxxxxx on the general system overview display in remote monitoring mode, but no performance history data has been found for the selected system. History data will be collected only if the 'append' flag for the system has been set to 'Y'.

### System action

The selection is ignored.

### User response

Change the 'append' flag for the system to 'Y' if you want history data to be collected, or enter the 'h' selection for another system where history data is available.

---

**FCXxxx421E**      **No exception data available for node xxxxxxxx**

---

### Explanation

An 'e' selection has been entered for system xxxxxxxx on the general system overview display in remote monitoring mode, but no exception data has been received today for the selected system.

### System action

The selection is ignored.

### User response

None.

---

**FCXxxx422E**      **No exceptions logged**

---

### Explanation

The exception log screen was selected for display, but no exception messages have been saved yet for the current day.

### System action

The selection is ignored.

### User response

None.

---

**FCXxxx423E**      **Cannot show SYSLOAD: Remote data collection not active**

---

### Explanation

A request has been entered to display the system load overview. This display can be shown only on systems that have been set up as central collectors for remote data from other VM systems.

### System action

The selection is ignored.

### User response

None.

---

**FCXxxx430E**      **Error nn reading remote data from card reader**

---

### Explanation

An error has occurred while reading data sent from another machine. *nn* is the return code given by the CMS RDCARD macro.

### System action

The reader file is not read.

### User response

See the appropriate CMS manual for the meaning of the indicated return code.

---

**FCXxxx431E**      **Error nn for xxxx of target disk**

---

## Explanation

A command ('FCONRMT ...' or 'FC MONCOLL CPMON DISK ...') has been entered which will automatically try to LINK and ACCESS the selected target disk. An error has occurred while trying to LINK or ACCESS the disk, as indicated by the string *xxxx.nn* is the return code received when trying to execute the LINK or ACCESS command.

This message is also sent to the maintenance user ID defined with the FC MAINTID command.

## System action

Data on the disk are not accessed.

## User response

See the appropriate CP or CMS documentation, and the error message generated on the basic mode screen, to determine the cause of the error (e.g. not authorized to LINK to the disk).

---

<b>FCXxxx432E</b>	<b>Command not executed, remote data collection active</b>
-------------------	--

---

## Explanation

An FCONRMT command has been entered for accessing a central data collector machine's A-disk, but the machine where the command has been entered is already acting as a central data collector itself. The two modes exclude each other

## System action

Operation continues.

## User response

None.

---

<b>FCXxxx433E</b>	<b>Command not executed, linked to central monitoring disk</b>
-------------------	--

---

## Explanation

An 'FC MONCOLL REMCOLL ..' command has been entered to start operation as central data collector machine, but the machine where the command has been entered is already linked to a central data collector itself. The two modes exclude each other.

## System action

Operation continues.

## User response

None.

---

<b>FCXxxx434E</b>	<b>GRAPHDET command supported only from REDISP mode while in FCONRMT</b>
-------------------	--

---

## Explanation

A 'GRAPHDET ...' command has been entered in remote monitoring mode. This command is supported only when working with performance redisplay data *in storage*, i.e. REDISP data must be displayed when entering the command in remote monitoring mode.

## System action

Operation continues.

## User response

Select REDISP data for display, then re-enter the GRAPHDET command.

---

<b>FCXxxx435E</b>	<b>'FILE' spec. needed unless in REDISP mode or RMONLOG file available</b>
-------------------	--

---

## Explanation

A 'GRAPHxxx ...' command has been entered in remote monitoring mode. The GRAPHSUM and GRAPHVAR commands will use, by default, performance REDISP data if you are viewing a REDISP screen, or the '*node-ID* RMONLOG' file for the current system if this file can be found, but neither is available.

## System action

The command is not executed, and operation continues.

## User response

Use the 'FILE' argument of the GRAPHSUM or GRAPHVAR commands to specify the complete file-ID of the RMONLOG*n*, PERFLOG*n* or PERFHIST file which contains the performance source data.

---

<b>FCXxxx436E</b>	<b>RSCS user-ID for this system not found</b>
-------------------	---

---

## Explanation:

A command has been entered which requires data transfer to a remote system, but Performance Toolkit could not determine the user ID of the RSCS machine.

## System action:



The command is not executed, and operation continues.

**User response:**

Make sure that an RSCS machine has been set up and that it has been defined in a SYSTEM NETID file to which the Performance Toolkit has access.

---

**FCXxxx437E      No data in storage, input file specification required**

**Explanation:**

You tried to create graphical output based on REDISP data in storage, but there are no such data, probably because you are working in a virtual machine where Performance Toolkit has not been set up for permanent data collection.

**System action:**

The selection is ignored.

**User response:**

Specify a valid input data file.

---

**FCXxxx438E      Variable 'vvvvvvvv' not valid for iiiiii input**

**Explanation**

You tried to create graphical output based on iiiiii input, but vvvvvvvv is not a valid variable name for that case.

**System action**

The selection is ignored.

**User response**

Specify a valid variable name.

---

**FCXxxx439E      Insufficient number of perf. variables specified**

**Explanation**

At least one variable must be specified for 'by time' graphics, and at least two variables are required for variable correlation type graphics. The number of variables found was lower. Defaults will be set when displaying HISTLOG/HISTSUM data, but not for any other type of input file.

**System action**

The selection is ignored.

**User response**

Specify at least the minimum number of variables.

---

**FCXxxx440E      Error nnnn loading MONITOR segment cccccc**

**Explanation**

An error was found when trying to load the MONITOR segment cccccc. The value nnnn is the return code from the SEGMENT LOAD macro.

This message is also sent to the maintenance user ID defined with the FC MAINTID command.

**System action**

The CPMON interface could not be activated.

**User response**

See the description of the SEGMENT LOAD macro and check the meaning of the return code received, then take the appropriate action and try to activate the interface again.

---

**FCXxxx441E      Error nnnn during cccc cccc for \*MONITOR service**

**Explanation**

An error was found while trying to establish the connection to the \*MONITOR IUCV service. The value nnnn is the return code from the IUCV function cccc cccc.

This message is also sent to the maintenance user ID defined with the FC MAINTID command.

**System action**

The CPMON interface could not be activated.

**User response**

See the description of the appropriate IUCV macro and check the meaning of the return code received. Take the appropriate action and try to activate the interface again.

---

**FCXxxx442E      CP MONITOR data collection already active - command ignored**

**Explanation**

You have tried to activate the Performance Toolkit's interface for extraction of CP MONITOR data. However, this interface is already active, and no further action is necessary.

**System action**

None.

**User response**

None.

---

**FCXxxx443E      IUCV connection severed by  
                  \*MONITOR, rc = nnnn**

---

**Explanation**

The IUCV connection to the \*MONITOR system service has been severed. The value *nnnn* is the reason code for this action.

This message is also sent to the maintenance user ID defined with the FC MAINTID command.

**System action**

The MONITOR interface is completely reset.

**User response**

Check the meaning of the code given, and look also for informational messages from CP which could explain why the connection was severed. Restart the interface when the problem has been solved.

---

**FCXxxx444E      IUCV reply failed with reason code  
                  nnnn**

---

**Explanation**

An IUCV REPLY which the Toolkit attempted upon receiving data from the \*MONITOR service failed. The value *nnnn* is the reason code given by the IUCV REPLY macro.

This message is also sent to the maintenance user ID defined with the FC MAINTID command.

**System action**

Depending on the exact nature of the problem, the interface may have been completely reset.

**User response**

See the description of the IUCV REPLY macro and check the meaning of the reason code given. Look also for messages from CP which could explain why the REPLY action failed. Restart the interface when the problem has been solved.

---

**FCXxxx445E      Return code nnnn for MONITOR  
                  START command**

---

**Explanation**

Performance Toolkit will attempt to issue a MONITOR START command after establishing an IUCV connection to the \*MONITOR service (after an FC

MONCOLL CPMON DCSS ON command). The value *nnnn* is the return code from the command MONITOR START.

This message is also sent to the maintenance user ID defined with the FC MAINTID command.

**System action:**

The MONITOR interface is completely reset.

**User response:**

Check the meaning of the return code received, and look also for informational messages from CP which could explain why the command did not work. Restart the interface when the problem has been solved.

---

**FCXxxx446E      Incomplete monitor data: nnnnn  
                  too small**

---

**Explanation**

Data from the CP monitor indicates that the information received is incomplete because either the DCSS is too small or the SAMPLE CONFIG size in the DCSS is too small. (Accordingly, *nnnnn* will be either 'DCSS' or 'SAMPLE CONFIG'.)

This message is also sent to the maintenance user ID defined with the FC MAINTID command.

**System action**

Data will be shown as far as possible, but the information is likely to be incomplete. Incorrect data may be shown, and the Performance Toolkit for z/VM could even abend.

**User response**

Inform the responsible systems programmer that either a larger monitor segment needs to be defined for the CP monitor data, or that the SAMPLE CONFIG area size in the monitor segment needs to be increased. You might also consider disabling some of the monitor domains, users or devices that are not needed, or enabling only the domains, users or devices that are needed. (Refer to the MONITOR SAMPLE command in [z/VM: CP Commands and Utilities Reference](#).)

---

**FCXxxx447E      Incomplete monitor data:  
                  abnormal termination**

---

**Explanation**

Data from the CP monitor indicate that the information received is incomplete because data notification has been terminated.

This message is also sent to the maintenance user ID defined with the FC MAINTID command.

## System action

Data will be shown as far as possible, but the information may be incomplete.

## User response

None.

---

<b>FCXxxx448I</b>	<b>Not authorized for MONITOR START: Start from other machine!</b>
-------------------	--

---

### Explanation:

Performance Toolkit tried to execute a MONITOR START command when data collection from the monitor shared segment was invoked. The return code received suggests that your virtual machine has not been authorized to execute this command, hence there is no guarantee that the CP monitor is started now.

### System action:

Operation continues, but performance data will be collected only if CP monitor data collection is active too.

### User response:

Make sure CP monitor data collection has already been activated by some other virtual machine, or do so now.

---

<b>FCXxxx449I</b>	<b>End of data reached</b>
-------------------	----------------------------

---

## Explanation

An 'end-of-file' condition has been reached while scanning a disk file. No further data records are available for scanning.

## System action

The command which resulted in coming up against the end of the disk file cannot be executed.

## User response

None.

---

<b>FCXxxx450E</b>	<b>CP MONITOR collection already active from cccc</b>
-------------------	---

---

## Explanation

You have tried to activate the Toolkit's interface for extraction of CP MONITOR data. However, this interface has already been activated from cccc and cannot be started again.

## System action

The command is ignored.

## User response

Stop the currently active interface first before you try to activate another one.

---

<b>FCXxxx451I</b>	<b>SEEKS EVENT data collection started for device xxxx</b>
-------------------	--

---

## Explanation

A 'SEEKDET xxxx' command was entered to display detailed seeks data for device xxxx. SEEKS data collection for that disk was not previously active and has automatically been started now to allow further analysis.

## System action

Operation continues.

## User response

Remember to disable SEEKS data collection for that device again when you have finished your analysis. SEEKS data collection will *not* be deactivated automatically by the Performance Toolkit.

---

<b>FCXxxx452E</b>	<b>Return code <i>nnnn</i> for MONITOR EVENT ENA SEEKS cmd</b>
-------------------	--

---

## Explanation

A SEEKDET xxxx command was entered to display detailed seeks data for device xxxx. Performance Toolkit received return code *nnnn* when it tried to enable SEEKS data collection for device xxxx. Unless SEEKS data collection had previously been enabled for the selected device from another machine, you will not see seeks information for the disk.

## System action

Operation continues.

## User response

Check the return code and give the Toolkit machine the privilege class required for activating SEEKS data collection, or enable SEEKS data collection from another (authorized) machine.

---

<b>FCXxxx453E</b>	<b>Not authorized for 'LOCATE', user limit checking disabled</b>
-------------------	--

---

## Explanation

The CP LOCATE command is used for USER LIMIT checking. This requires CP command privilege class C or E on a standard z/VM system. User limit checking

cannot be activated before your machine has been authorized for the necessary command privilege class.

**System action**

Operation continues.

**User response**

Have your machine authorized for running the LOCATE command by adding the required CP command privilege class (usually class C or E) to the directory entry of your machine. Log off, log back on, restart Performance Toolkit, and then issue the command again.

Caution: This authorization should be considered a security exposure if the use of the machine is not restricted to operating or system support personnel that have similar privileges already.

---

<b>FCXxxx454E</b>	<b>Internal Performance Toolkit for VM error, user limit checking disabled</b>
-------------------	--

---

**Explanation**

An internal error has been found which does not allow user limit checking to function properly. User limit checking has, therefore, been disabled.

**System action**

Operation continues.

**User response**

Contact IBM service.

---

<b>FCXxxx455E</b>	<b>cccccc process started before end of dddddd</b>
-------------------	--

---

**Explanation**

Process ccccc should normally start after completing process dddddd, so this message indicates an internal problem which is probably linked to excessive system load. This message will usually be accompanied by message FCXxxx459.

**System action**

Operation continues, but data tables may be cleared as indicated by message FCXxxx459.

**User response**

Refer to messages FCXxxx456, FCXxxx457, and FCXxxx458. The data collector machine probably needs a higher priority.

---

<b>FCXxxx456E</b>	<b>Monitor segment overlaid before completion of data collection</b>
-------------------	--

---

**Explanation**

The data collector machine could not complete processing of all monitor data presented in the monitor segment before the segment was overlaid with a new set of data. This problem was probably caused by excessive system load and/or insufficient priority for the data collector machine. The message will usually be accompanied by message FCXxxx459.

This message is also sent to the maintenance user ID defined with the FC MAINTID command.

**System action**

Operation continues, but data tables may be cleared as indicated by message FCXxxx459.

**User response**

Look also for messages FCXxxx457 or FCXxxx458. The data collector machine probably needs a higher priority.

---

<b>FCXxxx457A</b>	<b>uuuuuuuu nn% active; ccccccc higher priority</b>
-------------------	---

---

**Explanation**

The data collector machine uuuuuuuu was found to have been more than 50% active during the last monitor sample interval, where nn indicates the actual value. Since the machine must be able to process all of the available monitor data before CP overlays the segment with a new set of data, high activity values indicate a risk that the machine might no longer be able to do that if the situation deteriorates further. This is indicated by the text in string ccccccc, which will say 'may need' or 'requires' depending on the urgency of the problem, as indicated by the current activity level.

This problem may be caused by either the user running additional resource intensive tasks in the data collector machine or by excessive system load and/or insufficient priority for the data collector machine.

**System action**

Operation continues.

## User response

Do not run resource intensive tasks in the data collector machine. If the message was generated without such additional tasks being run, verify that the machine is being run with absolute share. Increasing the machine's priority may also help eliminate this message.

---

<b>FCXxxx458A</b>	<b>Insufficient resources for reliable data collection</b>
-------------------	--

---

## Explanation

A data collect error was found which is thought to be due to a lack of sufficient resources for the data collector machine so that it could not keep up processing monitor data. The message will usually be preceded by another message giving more details about the actual problem.

This message is also sent to the maintenance user ID defined with the FC MAINTID command.

## System action

Operation continues. Data tables may be cleared as indicated by message FCXxxxx459.

## User response

Do not run resource intensive tasks in the data collector machine. Increase the machine's priority if the message was generated without such additional tasks being run.

---

<b>FCXxxx459I</b>	<b>Data not reliable: all tables cleared</b>
-------------------	--

---

## Explanation

A data collect error was found which may have led to an inconsistent set of data being loaded. To prevent display of incorrect information all internal data tables have been cleared, causing partial loss of monitor data so far collected.

This message should be preceded by one or more of the messages to FCXxxxx458, indicating what caused the problem.

This message is also sent to the maintenance user ID defined with the FC MAINTID command.

## System action

Operation continues.

## User response

Be aware that no averages can be calculated for periods preceding the table clear.

---

<b>FCXxxx460E</b>	<b>Incorrect perf. variable definition in formula</b>
-------------------	---

---

## Explanation

An apparent performance variable definition has been found in a USERVAR formula (with leading single quote '), but either the variable name or the ending quote were found to be missing.

## System action

The command is ignored.

## User response

Check the formula for correct specification of all performance variables and retry the command.

---

<b>FCXxxx461E</b>	<b>Non-matching parentheses in formula</b>
-------------------	--

---

## Explanation

Opening and closing parentheses in a formula must always appear in pairs, but the number of right parentheses in your formula is not equal to the number of left parentheses, i.e. the format is not valid.

## System action

The command is ignored.

## User response

Correct the parentheses in the USERVAR formula and then retry the command.

---

<b>FCXxxx462E</b>	<b>Incorrect argument cccc in formula</b>
-------------------	---

---

## Explanation

An expression cccc has been detected in your formula which is neither a mathematical operator, a fixed number, nor a valid performance variable name.

## System action

The command is ignored.

## User response

Correct the invalid expression and retry the command.

---

<b>FCXxxx463E</b>	<b>Format ... cccc cccc cccc ... is incorrect</b>
-------------------	---

---

### Explanation

While analyzing the formula a sequence of its elements was found which appears invalid. Examples: two variable names without intervening mathematical operator, or two operators without a variable name between.

### System action

The command is ignored.

### User response

Fix the offending part of the formula and retry the command.

---

<b>FCXxxx464E</b>	<b>More than <i>nnnn</i> elements in formula</b>
-------------------	--

---

### Explanation

The routine for analyzing the formula can only handle up to *nnnn* elements (variable names, mathematical operators, parentheses), but the formula contains more elements than this number.

### System action

The command is ignored.

### User response

Try to devise a method for obtaining the desired results using a formula with less elements.

---

<b>FCXxxx465E</b>	<b>User variable <i>cccc</i> not defined</b>
-------------------	--

---

### Explanation

An 'FC USERVAR QUERY *cccc*' command, or a command for setting the description or the units of user variable *cccc* has been entered, but no user variable with this name has been found.

### System action

None.

### User response

Correct the variable name and re-issue the command, or define the variable first by specifying the formula to be used before you attempt to enter its description.

---

<b>FCXxxx466I</b>	<b>No user variables defined</b>
-------------------	----------------------------------

---

### Explanation

An 'FC USERVAR QUERY' command has been entered but no user variables have been defined, so none can be displayed.

### System action

None.

### User response

Define some user variables before you try to display them.

---

<b>FCXxxx467I</b>	<b><i>nnn</i> zero-divide(s) avoided, <i>mmm</i> times maximum exceeded</b>
-------------------	---

---

### Explanation

While executing the formula for a user variable, Performance Toolkit found *nnn* cases where a division by 0 was attempted, but was changed to a division by 1, and *mmm* cases where the result of the formula exceeded 20,000,000, but was set to this number.

### System action

Operation continues. Graphic output is created regardless of the potential errors that might have been introduced.

### User response

The resulting graph might or might not be meaningful, depending on the number of evasive actions taken and on the impact such actions have on the desired result. Verify the impact, or, to try to avoid such actions, adapt your formula.

---

<b>FCXxxx470E</b>	<b>Specification '<i>c</i>' not valid for '<i>function</i>'</b>
-------------------	---

---

### Explanation

The cache status definition character '*c*' entered with an 'FC MONCACHE ...' command is not valid for defining the expected of cache function '*function*', where *function* can be DEVICE (for general caching), DASD-FW (for DASD fast write), or DUALCOPY (for the dual copy function).

### System action

The command is ignored.

## User response

Refer to the command description (see “FCONTROL MONCACHE” on page 51) and correct the command syntax, then reenter the command.

See the FC MONCACHE command for syntax information. Then reenter the command.

---

<b>FCXxxx471E</b>	<b>Device xxxx is not a DASD and cannot be cached</b>
-------------------	---

---

## Explanation

The device number xxxx entered with an 'FC MONCACHE ...' command refers to an I/O device which is not a disk, and which cannot be cached.

## System action

The command is ignored.

## User response

Correct the device number, then re-enter the command.

---

<b>FCXxxx475I</b>	<b>Packed file, being unpacked now. Re-execute cmd. in a few seconds</b>
-------------------	--

---

## Explanation

A command has been entered which needs a performance data log file for input. The file on disk was found to be in packed format and is now being unpacked automatically (into a work file FCONWORK CMSUT1 A, i.e. the original file remains unchanged). The program may need a few seconds for unpacking, the command can then be entered again.

## System action

The file is unpacked.

## User response

Re-enter the original command after a few seconds. Do *not* change its format to use the file name and file type of the temporary work file.

---

<b>FCXxxx476E</b>	<b>File cannot be unpacked now, CMS command active</b>
-------------------	--

---

## Explanation

A command has been entered which needs a performance data log file for input. The file on disk was found to be in packed format, but it cannot be unpacked at the moment because another CMS command is already active in the server machine.

## System action

The command is aborted.

## User response

Wait some time to give the CMS command time to complete, then re-enter the command.

---

<b>FCXxxx479E</b>	<b>Hash {SEARCH INSERT} detected zero key</b>
-------------------	---

---

## Explanation

An internal error was found during a hash SEARCH or INSERT operation. The error was probably caused by a monitor segment overlay. It is also possible that the MONDCSS size needs to be increased. This message is usually accompanied by message FCXxxx456.

## System action

Operation continues.

## User response

Refer to messages FCXxxx456 to FCXxxx458. The data collector machine probably needs a higher priority. Also check whether the MONDCSS size should be increased.

---

<b>FCXxxx480I</b>	<b>No SPOOL files found</b>
-------------------	-----------------------------

---

## Explanation

A SPOOL data display has been selected, but no SPOOL files could be found on the system.

## System action

The selected screen cannot be shown.

## User response

None.

---

<b>FCXxxx481I</b>	<b>No SPOOL files found for user cccccccc</b>
-------------------	---

---

## Explanation

A 'SPOOL userid' command has been entered to display all the SPOOL files of user cccccccc, but none were found.

## System action

The selected screen cannot be shown.

## User response

None.

---

**FCXxxx482E      No VDISKs defined on this system**

## Explanation

A 'VDISKs' command has been entered to display the 'Virtual Disks in Storage' screen, but no monitor data have been collected for VDISKs even though the STORAGE sample domain appears to be enabled. The probable reason is that no virtual disks have been defined.

## System action

The selected screen cannot be shown and the initial menu is displayed instead.

## User response

None.

---

**FCXxxx484E      No Byte File System data found**

## Explanation

One of the Byte File System logs was selected to be displayed or printed, but no Byte File data had been collected.

## System action

The selected screen cannot be shown and the initial menu is displayed instead.

## User response

Select another display.

---

**FCXxxx485I      I/O Priority Queueing Facility not active**

## Explanation

The I/O priority queueing users screen was selected for display or for printing. I/O priority queueing consists of two parts: priority queueing of virtual I/O requests by CP and, where available, priority queueing of real I/O requests by the corresponding hardware facility. The hardware facility is not active.

## System action

None.

## User response

None.

---

**FCXxxx486E****No I/O Priority Queueing Facility users found**

## Explanation

The I/O priority queueing users screen was selected for display or for printing, but none of the virtual machines on this system are currently using the facility.

## System action

The selected screen cannot be shown and the initial menu is displayed.

## User response

None.

---

**FCXxxx487E****No users with any QDIO activity found**

## Explanation

No users were found with any QDIO activity when the UQDIO command was entered.

## System action

The current screen continues to be displayed.

## User response

None.

---

**FCXxxx488E****No QEBSM devices found**

## Explanation

No QEBSM devices were found when the QEBSM command was entered.

## System action

The current screen continues to be displayed.

## User response

None.

---

**FCXxxx490E****Return code *nnn* received for VMCF *ccccccc* function**

## Explanation

An error occurred while executing a VMCF diagnose function (used for the VMCF data retrieval interface). *nnn* is the return code received, and *ccccccc* the function being executed.



### System action

The function is not executed.

### User response

See the documentation of the VMCF facility in your VM library for the message of the return code.

---

<b>FCXxxx491I</b>	<b>Screen incomplete: partial cccccccc data</b>
-------------------	---

---

### Explanation

An attempt was made to select a performance screen for displaying or printing, but the display could not be prepared completely because not all data has been collected. The string *cccccccc* indicates the kind of data that is required for building the desired display.

### System action

Incomplete performance data is shown. Dots will be inserted for fields where the information is not available.

### User response

None.

---

<b>FCXxxx492I</b>	<b><i>screen not available, requires at least two monitor sample intervals</i></b>
-------------------	--

---

### Explanation

An attempt was made to select a performance screen for displaying or printing, but the screen could not be prepared because it required at least two monitor sample intervals for data calculation. In batch mode, Performance Toolkit was set up (using some combination of the FC MONCOLL RESET, FC SET BYTIME, and FC SET INTERIM commands) to print that screen, but the second set of samples had not yet been processed. An example of such a screen is the Privileged Operations Screen (PRIVOP), which consists entirely of counter deltas.

### System action

The display is not selected or printed.

### User response

The actions to try, before selecting the performance screen again, depend on the current Performance Toolkit mode of operation:

- If permanent data collection was active, wait for the second monitor sample interval to arrive.

- If you are running MONSCAN manually, issue the NEXTSAMP subcommand to get additional monitor samples.

In batch mode, this message can be ignored.

---

<b>FCXxxx493E</b>	<b>No multitasking application data available</b>
-------------------	---

---

### Explanation

A performance display based on multitasking application data was selected, but no such data had been collected yet, and the performance display cannot be shown.

### System action

The selected screen cannot be shown and the initial menu is displayed instead.

### User response

Select another display.

---

<b>FCXxxx494E</b>	<b><i>screen not available, requires data_domain data</i></b>
-------------------	---

---

### Explanation

An attempt was made to select a performance screen for display or printing, but the display could not be prepared because the prerequisite data was not available.

When Performance Toolkit has not processed any Domain 1 Record 20 (MRMTRHPP - HyperPAV Pool Definition) or Domain 6 Record 28 (MRIODHPP - HyperPAV Pool Activity) monitor records before the screen or report was requested, one of the following variations of this message appears:

HPALIAS not available, requires MONITOR I/O SAMPLE data
HPSHARE not available, requires MONITOR I/O SAMPLE data

This could happen because Performance Toolkit has just been started and has not finished processing the initial configuration data. Or, it is possible that no Domain 6 Record 28 monitor records are being produced because the MONITOR I/O SAMPLE domain is disabled.

### System action

The display is not selected or printed.

## User response

Activate the required data domain before trying to call that display again. If the problem persists, refer to the corresponding performance screen description for possible additional prerequisites.

---

**FCXxxx495I**      **Screen incomplete, no ccccc/.../cccc data**

---

## Explanation

The selected performance display requires data from several CP monitor domains, but not all of them are currently enabled. The string *cccc/.../cccc* gives you the name of the domain(s) for which no monitor records have been processed.

## System action

Performance data are shown but will be incomplete.

## User response

Activate the required data domain(s).

---

**FCXxxx496E**      **Error *nnn* processing record *mm* of file *fn ft fm***

---

## Explanation

An error was encountered while processing file *fn ft fm*. Record *mm* of this file contains a command (or specifications which are translated internally to a command) which resulted in an error code *nnn*. This error code corresponds to the error message number generated as a result of the faulty command.

## System action

See the description of message number *nnn*.

## User response

Check and fix the faulty specification; see the description of message number *nnn* for more information.

---

**FCXxxx497E**      **No channel-path measurement facility data available**

---

## Explanation

The LPAR channel load display has been selected, but the Channel-Path Measurement Facility data which are required as input for building this display are not available.

## System action

The selection is ignored.

## User response

Refrain from selecting this particular display on systems where the channel-path measurement facility is not available.

---

**FCXxxx498E**      **No extended function data available**

---

## Explanation

A CACHEXT command has been entered to show the cache extended function display, but this display cannot be generated because no monitor data have been collected yet for extended function cache data. Possible reasons are:

- The system's I/O configuration does not include a cached control unit is capable of handling extended function operations.
- Monitor I/O data collection for the disks connected to the cached control unit has not been enabled.

## System action

The selection is ignored.

## User response

Enable data collection for sample data of the I/O domain, if that was the problem.

---

**FCXxxx499E**      **No Linux ccccccc application data available**

---

## Explanation

A Linux performance data screen has been selected whose contents are based on APPLDATA domain monitor data, but no such application data records have been collected. The string *ccccccc* indicates the kind of application data record that is required for building the desired display.

Possible reasons include:

- The Linux systems have not been set up to generate application data records (requires installation of the data gathering code in the Linux, and the 'APPLMON' option in the virtual machine's directory entry).
- Monitor APPLDATA collection has not been enabled for the Linux guest(s).

## System action

The selection is ignored.

## User response

Inform your systems programmer of the problem and make sure both the Linux systems and the CP monitor are set up to collect the application.

---

**FCXxxx500I**      **Performance Toolkit for VM  
releaselevel servicelevel (64-bit)**

### Explanation:

This message indicates which level of the program you are using. The phrase 64-bit indicates that Performance Toolkit is running under z/CMS in z/Architecture mode exploiting 64-bit memory addressing.

### System action:

None.

### User response:

None.

---

**FCXxxx510I**      **New sort sequence set, will  
become active with next update**

## Explanation

This message is shown after changing the sort sequence on a display where the screen output is sorted only after collecting new data.

## System action

None.

## User response

Wait for the next automatic data collect cycle at the end of the current 1-minute interval, or force an extra collect cycle by pressing the ENTER key.

---

**FCXxxx511I**      **xxxxxxx data set, active for this  
display after next update**

## Explanation

An AVERAGE or CURRENT command has been entered, as indicated by the string xxxxxxx, but the performance display being viewed will show the changed data only after collecting new data.

## System action

None.

## User response

Wait for the next automatic data collect cycle at the end of the current 1-minute interval, or force an extra collect cycle by pressing the ENTER key.

---

**FCXxxx512E**      **Cannot add definition, all entries  
in use**

## Explanation

An 'FC DEFSCRN' command has been entered for a user-defined screen for which eight specifications have already been saved, so there is no more space for another definition.

## System action

The command is ignored.

## User response

None.

---

**FCXxxx513I**      **'cccccccc' definitions deleted**

## Explanation

An 'FC DEFLOG' command has been entered for removing the definitions of user defined log cccccccc. The message indicates successful execution of the command.

## System action

Operation continues.

## User response

None.

---

**FCXxxx514E**      **Trailing quote missing or string too  
long**

## Explanation

An 'FC DEFLOG' command has been entered with the LOCATE argument. The string to be located is either too long, or the mandatory trailing quote after the string is missing.

## System action

The definition is ignored.

## User response

Correct the command syntax and re-enter the command.

---

**FCXxxx515E**      **No string found between quotes**

### Explanation

An 'FC DEFLOG' command has been entered with the LOCATE argument. However, no string was found between the two quotes, i.e. the 'LOCATE' argument does not make sense.

### System action

The definition is ignored.

### User response

Re-enter the command with a valid string definition.

---

**FCXxxx516E      No user defined displays found**

### Explanation

An 'UDEFMENU' command has been entered to display the selection menu for user defined displays. No such displays were found, either because none had previously been defined, or because the monitor data required for building them are not available.

### System action

The general performance data selection menu is displayed.

### User response

None.

---

**FCXxxx519E      S&F connect attempt to S&F interface by uuuuuuuu**

### Explanation

User uuuuuuuu tried to connect to the APPC/VM resource used by the S&F interface. The data received indicates that this attempt was made by the S&F interface of user uuuuuuuu. This is an error due to incorrect target resource specifications in user uuuuuuuu's FCONRMT SYSTEMS file for this system, or in the xCOMDIR NAMES file.

### System action

The connection is severed.

### User response

Systems programmer should specify the correct APPC/VM resource name for the system user uuuuuuuu tried to connect to.

---

**FCXxxx520E      Error nnnn during ccccccc ccccccc for APPC/VM**

### Explanation

A return code of nnnn was received when executing the macro function ccccccc ccccccc for APPC/VM processing. For example you may get 'Error 1041 during CMSIUCV CONNECT' if you specified a modename in the communications directory entry for the local resource.

### System action

The APPC/VM connection is not established.

### User response

See the appropriate manual for the meaning of the return code and take appropriate corrective action.

---

**FCXxxx521I      Connecting to server**

### Explanation

An FCONAPPC command has been entered to start an APPC/VM session with a Performance Toolkit performance data collector machine, and the program is now trying to establish the connection.

### System action

The program is trying to establish the connection.

### User response

None.

---

**FCXxxx522I      Connected to resource ccccccc**

### Explanation

An FCONAPPC command has been entered to start an APPC/VM session with a Performance Toolkit performance data collector machine, and the program is now confirming that the connection could be made to the selected global resource ccccccc.

### System action

None.

### User response

None.

---

**FCXxxx523E      Resource 'ccccccc' not available**

### Explanation

An FCONAPPC command has been entered to start an APPC/VM session with a Performance Toolkit performance data collector machine. The resulting

connect request was made for global resource cccccccc, but that resource was found to be not available.

**System action**

The APPC/VM connection cannot be established.

**User response**

Adapt your communications directory if the FCONAPPC command resulted in unexpected APPC/VM connect request, or check why it was not possible to connect to the selected resource.

---

**FCXxxx524E      Not authorized to manage resource 'ccccccc'**

**Explanation:**

An FC MONCOLL VMCF ON command has been executed, resulting in Performance Toolkit trying to establish a connection to the \*IDENT service in order to define itself as the resource manager for global resource cccccccc. Performance Toolkit received a return code indicating that it has not been authorized to manage that resource.

**System action:**

Performance Toolkit cannot define itself as the resource manager for cccccccc.

**User response:**

Include an IUCV \*IDENT cccccccc GLOBAL statement in the directory of the machine. Make sure the directory is activated, then log off, log on, and try again.

---

**FCXxxx525E      Not currently connected to server, resetting**

**Explanation:**

Performance Toolkit can no longer use the APPC/VM connection needed for data retrieval so it will reset the connection and return to basic mode.

**System action:**

The APPC/VM connection is severed.

**User response:**

Try to re-establish the APPC/VM connection by entering the FCONAPPC command again. You should receive an error message indicating what is wrong or you will be connected again.

---

**FCXxxx526I      Not in SEND state, command stacked**

**Explanation**

A command has been entered for transmission to the server machine, but we are currently in receive state, waiting for the result of a previous command to be

returned first. The new command has been stacked and will be transmitted when the previous command's output has arrived.

**System action**

The new command is stacked for later execution.

**User response**

Wait until the previous command's output has arrived.

---

**FCXxxx527I      User cccccccc connected on path nnnn**

**Explanation**

The APPC/VM server machine (the toolkit's performance data collector machine) has received a request on path nnnn from user cccccccc to connect, and it has granted permission to establish this connection.

**System action**

None.

**User response**

None.

---

**FCXxxx528I      User cccccccc has severed path nnnn**

**Explanation**

The APPC/VM server machine (the toolkit's performance data collector machine) has received a 'sever' indication for path nnnn from user cccccccc, and it has severed the path too.

**System action**

None.

**User response**

None.

---

**FCXxxx529I      Path nnnn to cccccccc severed by other side, code xxxx**

**Explanation**

The APPC/VM path pppp to resource cccccccc has been severed by the server or by some intermediate function (AVS, for example). The sever code xxxx determines the reason. Some sever codes and reasons are:

**X'0142'**

The resource cccccccc is not available on the remote system to which you tried to connect, possibly due to one or more of the following:

1. The AVS connection is down.
2. The TSAF connection is down.
3. The target server did not correctly initialize to manage the target resource. Make sure that the correct directory entries (IUCV statement and communication directory entries) were made and that APPC/VM connections are enabled using the FC MONCOLL VMCF ON command. Check its log. Messages generated while trying to initialize the interface should help you find the problem.

**X'0160'**

An error was found while trying to map your LU and user ID combination to a local user ID in the remote AVS machine. It is possible that an AGW ADD USERID statement is missing.

**X'0301'**

An incorrect target LU is specified in the communication directory.

**X'0302'**

There is a remote connection, but the modname in the communication directory is not valid or is missing.

Check the console logs of the AVS machines involved in the connection; they might contain error messages that will help to locate the source of the problem.

**System action**

The current APPC/VM session is reset and basic mode is entered.

**User response**

None.

---

<b>FCXxxx530I</b>	<b>Connected to *IDENT for resource cccccccc</b>
-------------------	--

---

**Explanation:**

The Performance Toolkit data collector machine has established a connection to the \*IDENT system service and has identified itself as manager of the resource cccccccc, as the result of an FC MONCOLL VMCF ON command.

**System action:**

The machine is now identified as the global resource manager for cccccccc.

**User response:**

None.

**FCXxxx531I**


---

**All paths severed for resource cccccccc**

---

**Explanation**

All paths have been severed for global resource cccccccc, probably following an 'FC MONCOLL VMCF OFF' command.

**System action**

All paths have been severed.

**User response**

None.

**FCXxxx532E**


---

**Not authorized to connect to resource 'ccccccc'**

---

**Explanation**

A connection request to global resource cccccccc resulted in a return code indicating that you are not authorized to connect to that resource.

**System action**

The connect request is aborted.

**User response**

Contact your systems administrator and have your virtual machine authorized to connect to resource cccccccc. (The 'IUCV ALLOW' statement is probably missing in the server machine's CP directory).

**FCXxxx533I**


---

**Path nnnn for resource cccccccc severed by \*IDENT**

---

**Explanation**

The IUCV path nnnn to the \*IDENT system service for global resource cccccccc has been severed by \*IDENT.

**System action**

The connection is severed.

**User response**

Determine the reason for the sever and correct the problem.

**FCXxxx534E**


---

**Incorrect request X'xxxxxxxxxxxxxxxx' received**

---

### Explanation

In an APPC/VM session, the server received this string from its partner:

```
X 'xxxxxxxxxxxxxxxx'
```

The request format is incorrect and the request cannot be executed. This could be due to a mismatch between the code levels of the Performance Toolkit module in the server and requester machines.

### System action

The request is ignored.

### User response

Make sure that Performance Toolkit modules with the same version, release, and maintenance level are used on both sides of the requester-server connection.

<b>FCXxxx535I</b>	<b>Connected to resource rrrrrrr on path p, for uuuuuuuu</b>
-------------------	--

### Explanation

This information message informs the user that the S&F server function has established an APPC/VM connection to resource rrrrrrr for user uuuuuuuu, on IUCV path p.

### System action

None.

### User response

None.

<b>FCXxxx536I</b>	<b>Path p to resource rrrrrrr severed</b>
-------------------	---

### Explanation

This information message informs the user that the S&F server function has severed the APPC/VM connection to resource rrrrrrr on IUCV path p.

### System action

None.

### User response

None.

<b>FCXxxx537E</b>	<b>Requestor module backlevel, no S&amp;F support</b>
-------------------	---

### Explanation

A user connected to the APPC/VM resource defined for the store and forward interface, but is using a backlevel Performance Toolkit module that does not yet support S&F sessions.

### System action

The APPC/VM connection is severed.

### User response

Obtain a current Performance Toolkit module and try again.

<b>FCXxxx538E</b>	<b>Target server module backlevel, no S&amp;F support</b>
-------------------	---

### Explanation

The Performance Toolkit module on the target service machine is backlevel and does not yet support store and forward APPC/VM connections.

### System action

The connection is severed.

### User response

Have a current level of the Performance Toolkit module installed on the target system.

<b>FCXxxx539I</b>	<b>No FCONRMT SYSTEMS file; S&amp;F interface not activated</b>
-------------------	---

### Explanation

No FCONRMT SYSTEMS file was found on any of the accessed minidisks. Performance Toolkit depends on information stored in that file (system identifications and matching APPC/VM resource IDs) for setting up its store and forward connections, so the S&F interface could not be activated.

### System action

None.

### User response

Make a valid FCONRMT SYSTEMS file available to the S&F PERFSVM machine.

<b>FCXxxx540E</b>	<b>ccccccc xxxx not benchmarked, log not available</b>
-------------------	--

## Explanation

A command has been entered to display a log for an object, but the selected log is not available because the specified object is not currently benchmarked. The strings `cccccccc xxxx` will be in one of the following formats: 'User *userid*' if a user log was selected, 'Device *devicenum*' if an I/O device log was selected, 'LLink *partner*' if an ISFC Logical link log was selected, or 'Epoint *epoint*' if an ISFC end point was selected.

## System action

The selection is ignored.

## User response

Use the command 'FC BENCHMRK ...' to activate benchmarking for the specified object before trying to view one of the related logs.

---

<b>FCXxxx541E</b>	<b>No data available for this log of cccccccc xxxx</b>
-------------------	--

---

## Explanation

A command has been entered to display a log for a benchmarked object, but the selected log is not available even though the specified object is currently benchmarked. The string "`cccccccc xxxx`" will appear in one of the following formats:

- "User *userid*" if a user log was selected.
- "Device *devicenum*" if an I/O device log was selected.
- "LLink *partner*" if an ISFC logical link log was selected.
- "Epoint *epoint*" if an ISFC end point was selected.

Possible reasons are:

- The selected log depends on data from a monitor domain that is not currently enabled.
- A log was selected to display data that do not apply to the I/O device (CPOWNLOG was selected for a disk that does not contain system areas or cache data was selected for a disk that is not connected to a cached control unit, for example).

## System action

The selection is ignored.

## User response

Enable monitor data creation for the object, if required.

---

<b>FCXxxx542E</b>	<b>No BENCHMRK data available</b>
-------------------	-----------------------------------

---

## Explanation

A command was entered to display a log for a benchmarked object, but no specified object was previously selected for benchmarking. The string `cccccccc xxxx` will be in one of the following formats: "User *userid*" if a user log was selected, "Device *devicenum*" if an I/O device log was selected, "LLink *partner*" if an ISFC logical link log was selected, or "Epoint *epoint*" if an ISFC end point was selected.

## System action

The selection is ignored.

## User response

Use the FC BENCHMRK command to activate benchmarking for a selected object before trying to view one of the related logs.

---

<b>FCXxxx543E</b>	<b>Specified object not benchmarked</b>
-------------------	---

---

## Explanation

An 'FC BENCHMRK ...' command was entered to disable benchmarking for a specific object (user, I/O device, ISFC end point, ISFC logical link, or PCI function), but the specified object was not previously benchmarked.

## System action

The command is ignored.

## User response

Verify the target object identification and reissue the command with the correct argument.

---

<b>FCXxxx544E</b>	<b>Filetype <i>ttttttt</i> not valid for benchmark log file</b>
-------------------	---

---

## Explanation

The string *ttttttt*, entered as part of a command for displaying or printing a benchmark log file, is interpreted as filetype of the file but is not one of the valid filetypes.

## System action

The command is ignored.

## User response

Re-enter the command with a valid benchmark log filetype.



**FCXxxx545E**      **Unexpected record length *nnnn* for log file *fn ft fm***

**Explanation**

A command has been entered for processing benchmark log file *fn ft fm*, but the record length of that file does not correspond to the expected format for the selected filetype.

**System action**

The input file is considered invalid, and the command is not executed.

**User response**

Re-enter the command for a valid benchmark log file.

**FCXxxx546E**      **Invalid data format for benchmark log file**

**Explanation**

A command has been entered for processing a benchmark log file, but the data read shows a format that is inconsistent with the expected benchmark log file layout.

**System action**

The input file is considered invalid, and the command is not executed.

**User response**

Re-enter the command for a valid benchmark log file.

**FCXxxx547E**      **Invalid data format for *ccccccc* FILE**

**Explanation**

A command has been entered for processing a file but the data read shows a format that is inconsistent with the expected data layout, where the string *ccccccc* indicates the supported input file type(s).

**System action**

The input file is considered invalid, and the command is not executed.

**User response**

Re-enter the command for a valid file.

**FCXxxx550E**      **Control file '*fn ft fm*' not found**

**Explanation**

A command has been entered which requires an additional control file *fn ft fm* as input, but the file could not be found on any of the user's accessed disks.

**System action**

The command aborts.

**User response**

Create the required control file and then re-execute the command.

**FCXxxx551I**      **Control record with unexpected IPTYPE *X'xx'***

**Explanation**

While processing a monitor data file a control record was found with an unexpected IPTYPE value of *X'xx'* in the fourth byte of the record The expected value for files created by MONWRITE is *X'09'*. This message will always be followed by message 553 that lists the contents of the first five fullwords.

**System action**

The control record and any monitor data it describes are ignored.

**User response**

Report the problem.

**FCXxxx552I**      **Control record without data control area**

**Explanation**

While processing a monitor data file a control record was found that does not contain the expected data control area pointers in the fourth and fifth message will always be followed by message 553 that lists the contents of the first five fullwords.

**System action**

The control record and any monitor data it describes are ignored.

**User response**

Report the problem.

**FCXxxx553I**      **Record *nnnn*: *X' xxxxxxxx xxxxxxxx xxxxxxxx xxxxxxxx xxxxxxxx '***

### Explanation

This message lists the first five fullwords of a control record that was found to contain unexpected data while processing a monitor data file. The string *nnn* is the number of the record and the strings *xxxxxxx* show the data in hexadecimal. This message is always preceded by messages 551 or 552 that explain which part of the record contains unexpected information.

### System action

See action for the preceding message 551 or 552.

### User response

See action for the preceding message 551 or 552.

---

**FCXxxx554I      No monitor data found**

### Explanation

An attempt has been made to scan a monitor data file by means of the MONSCAN command but no monitor data could be found. MONSCAN can process only monitor data files created by MONWRITE, or by any other tool that produces the same output record format.

### System action

No performance data can be displayed.

### User response

Make sure the MONSCAN arguments refer to a valid monitor data file and re-enter the command.

---

**FCXxxx560E      Less than two samples for vvv, CORREL, screen cannot be built**

### Explanation

Less than two samples were found with valid data for the base variable *vvvv*, so it is not possible to calculate correlation coefficients, and the CORREL display cannot be built.

### System action

The CORREL display is not shown.

### User response

Wait until more data have been collected, or select another, larger history log as input.

---

**FCXxxx561E      Unexpected HISTLOG/HISTSUM file format**

### Explanation

The format of the selected input file is incorrect. The input file is probably not one of the supported performance history files (default filetypes HISTLOGn or HISTSUM).

### System action

The display is not shown.

### User response

Enter the command again, specifying a valid history file as input.

---

**FCXxxx562I      Base variable is a constant**

### Explanation

The selected base variable appears to be a constant.

### System action

The display is shown, but it does not provide useful information.

### User response

Select another base variable.

---

**FCXxxx563I      No HELP text available for user defined variables**

### Explanation

The current release does not allow users to add help text for variables they define, so it is not possible to display it either.

### System action

None.

### User response

None.

---

**FCXxxx564I      No APPC/VM or TCP/IP connections in FCONRMT SYSTEMS file**

### Explanation

During remote performance monitoring facility initialization Performance Toolkit found the FCONRMT SYSTEMS control file, but no systems with an APPC/VM connection have been defined in it.

### System action

Operation continues.

### User response

Define in the file FCONRMT SYSTEMS at least one system with an APPC/VM connection. For more information, see the “Preparing the Performance Toolkit” chapter in *z/VM: Performance Toolkit Guide*.

---

**FCXxxx565I      S&F interface not activated**

### Explanation

This message is issued after the message FCXxxx564I indicating that remote performance monitoring facility has not been enabled.

### System action

Operation continues. The remote performance monitoring facility is not enabled.

### User response

Define in the file FCONRMT SYSTEMS at least one system with an APPC/VM connection. For more information, see the “Preparing the Performance Toolkit” chapter in *z/VM: Performance Toolkit Guide*.

---

**FCXxxx570I      Unexpected IUCV IPTYPE code cc received**

### Explanation

The Web Server interface received an unexpected IUCV IPTYPE code cc.

### System action

None.

### User response

Contact IBM Service.

---

**FCXxxx571I      Connected to TCP/IP server ccccccc on path p**

### Explanation

This information message tells you that the Web Server interface has successfully established an IUCV connection to TCP/IP server ccccccc on IUCV path p.

### System action

None.

### User response

None.

---

**FCXxxx572I      IUCV path severed by ccccccc**

### Explanation

This information message tells you that the IUCV path to virtual machine ccccccc has been severed by the other side.

### System action

Operation cannot continue, and the Web Server interface is deactivated.

### User response

None.

---

**FCXxxx573I      Return code rr, ERRNO ee for LASTERRNO call**

### Explanation

This information message tells you that an error condition has been detected by the Web Server interface.

### System action

The problem could potentially be severe, so the Web Server interface is deactivated.

### User response

Contact IBM Service.

---

**FCXxxx574I      cccc TCP/IP socket number is nn**

### Explanation

This information message provides some information about TCP/IP socket usage. The string cccc can be

- Max., indicating that the highest socket number available the Web Server interface is nn.
- Main, indicating that socket number nn has been defined as the 'main' socket for the Web Server interface.

### System action

None.

### User response

None.

---

**FCXxxx575I**      **sssssss host IP address is iiiiiiii**

### Explanation

This information message displays the host IP address of the interface sssssss.

### System action

None.

### User response

None.

---

**FCXxxx576E**      **Unexpected TCP/IP error nn on socket socket talking to ipaddr**

### Explanation

An error condition *nn* was found during a TCP/IP call on socket *socket* for *ipaddr*.

### System action

Depends on the error found. Operation will usually continue, but you may have to restart the Web Server interface.

### User response

There may be an accompanying FCX693I message which will provide more details. If not, look up the error code in Appendix C of the [z/VM: TCP/IP Programmer's Reference](#) to resolve the problem.

---

**FCXxxx577I**      **cccc for iiiiiiii on socket s**

### Explanation

This information message tells you that function *cccc* was executed for the connection to IP address *iiiiiii* on socket *s*.

### System action

None.

### User response

None.

---

**FCXxxx578E**      **Doesn't match original IP address of iiiiiiii**

### Explanation

This error message follows a message FCXxxx577I for the RECEIVE function. It indicates that the IP

address *iiiiiii* found in the received data buffer is not the same as the IP address when the user originally logged on.

### System action

The Web Server interface is deactivated to force a user action.

### User response

Find out why the IP address changed (this could be a hacking attempt) and restart the Web Server interface.

---

**FCXxxx579E**      **IUCV SEND error, IPAUDIT=X'xxxxxxx'**

### Explanation

An IUCV SEND error was encountered, where xxxxxxx contains the IPAUDIT field contents.

### System action

Depending on the error condition, the Web Server interface may be deactivated.

### User response

Check whether the TCP/IP server machine is still operating and restart it if necessary.

---

**FCXxxx580E**      **IUCV interrupt from incorrect path**

### Explanation

An IUCV interrupt was received on an unexpected path.

### System action

None.

### User response

Contact IBM Service.

---

**FCXxxx581I**      **Unexpected (non-ASCII) data received**

### Explanation

Non-ASCII data were found in an input buffer from the TCP/IP interface. our TCP/IP connection was established for the exclusive use by Web Browsers over the Internet, such non-ASCII data cannot be processed.

## System action

None.

## User response

None.

---

<b>FCXxxx582I</b>	<b>Cannot process specification cccccccccccc</b>
-------------------	--

---

## Explanation

The Web Server interface received an HTTP specification *cccccccccccc*, but does not know what to do about it. Since the interface has been designed to handle only the subset of Web Server functions required to allow performance data retrieval sessions, the message can probably be ignored.

## System action

None.

## User response

Contact IBM Service.

---

<b>FCXxxx583I</b>	<b>Unknown ccccccc specification aaaaaaaaaaaa</b>
-------------------	---

---

## Explanation

The Web Server interface received an argument *aaaaaaaaaaaa* for HTTP specification *cccccccc*, but does not know what to do about it. Since the interface has been designed to handle only the subset of Web Server functions required to allow performance data retrieval sessions, the message can probably be ignored.

## System action

None.

## User response

Contact IBM Service.

---

<b>FCXxxx584E</b>	<b>Missing request specification after GET</b>
-------------------	--

---

## Explanation

An HTTP 'GET' request was received, but it does not specify what data should be returned.

## System action

The home page (logon page) is returned to the requestor.

## User response

Contact IBM Service.

---

<b>FCXxxx585E</b>	<b>Incorrect request format cccccccccccc</b>
-------------------	--

---

## Explanation

An unexpected request *cccccccccccc* was received which we cannot process.

## System action

The home page (logon page) is returned to the requestor.

## User response

Contact IBM Service.

---

<b>FCXxxx586E</b>	<b>cccccccc HTTP version vv after GET</b>
-------------------	---

---

## Explanation

The Web Server interface has been designed for HTTP 1.0. The HTTP version was either missing, invalid (not in *n.n* format) or unsupported (pre 1.0), as indicated by the string *cccccccc*.

## System action

The home page (logon page) is returned to the requestor.

## User response

Contact IBM Service.

---

<b>FCXxxx587I</b>	<b>Userid <i>userid</i> identified on <i>ip-</i> <i>address</i> Userid <i>userid</i> by <i>byuserid</i> identified on <i>ip-address</i></b>
-------------------	---

---

## Explanation

User *userid* (or *userid* by *byuserid*) on IP address *ip-address* just logged on to the Web Server interface, after validation by the DMSPASS routine.

## System action

None.

## User response

None.

---

**FCXxxx588E**      **User login failed from *ip-address*  
*userid* password expired**

## Explanation

User *userid* on IP address *ip-address* tried to log on to the Web Server interface, but that *userid*'s password has expired.

## System action

The home page (logon page) is returned to the requestor.

## User response

Follow your installation procedure for having your *userid* reinstated.

---

**FCXxxx590I**      **sssssss interface activated**

## Explanation

This information message indicates that the TCP/IP interface for function *sssssss* has been correctly activated and is now ready for use.

## System action

None.

## User response

None.

---

**FCXxxx591I**      **sssssss interface deactivated**

## Explanation

This information message indicates that the TCP/IP interface for function *sssssss* has been deactivated.

## System action

None.

## User response

None.

---

**FCXxxx592I**      **User *uuuuuuuuu* at *iiiiiiiiiii* reset:  
Idle for over *nn* min.**

## Explanation

This information message indicates that user *uuuuuuuuu* has been found idle for over *nn* minutes, and that his session has been reset.

## System action

All control blocks are released, i.e. when the user wants to send another request a new logon is required first.

## User response

None.

---

**FCXxxx593E**      **S&F interface not yet active,  
WebServer interface not started**

## Explanation

Data retrieval via the WebServer interface is based on the APPC/VM store and forward logic managed by the S&F interface. The S&F interface has to be set up and activated, therefore, before the WebServer interface can be started.

## System action

The Web Server interface is not started.

## User response

Set up and start the APPC/VM S&F interface, then activate the Web Server interface again.

---

**FCXxxx594E**      **Password missing, or invalid  
format**

## Explanation

No password string could be found in the data returned from the logon page by the Web Browser, or it has an invalid format.

## System action

The home page (logon page) is returned to the requestor.

## User response

Re-enter *userid* and password.

---

**FCXxxx595E**      **UserId missing, or invalid format**

## Explanation

No userid string could be found in the data returned from the logon page by the Web Browser, or it has an invalid format.

## System action

The home page (logon page) is returned to the requestor.

## User response

Re-enter userid and password.

---

**FCXxxx596E**      **Unexpected forms data received**

## Explanation

The data returned by the Web Browser has an unexpected format and cannot be correctly analyzed. This could happen if a very old Web Browser is used, or if a very new Web Browser introduces a different forms format that the interface is not yet prepared for.

## System action

The home page (logon page) is returned to the requestor.

## User response

Contact IBM Service.

---

**FCXxxx597E**      **File '*fn ft fm*' not found**

## Explanation

The userid/password file with fileid *fn ft fm* could not be found on any of the accessed disks. Userid and password validation is not possible, therefore, and logon to the Web Server interface is inhibited.

## System action

The home page (logon page) is returned to the requestor.

## User response

Contact your systems programmer to have the password file set up, then try logging on again.

---

**FCXxxx598E**      **Missing data in POST request**

## Explanation

A WWW 'POST' request was received, but without the expected data after the HTTP header part. This should

be a problem of the WWW browser used, or of a proxy server that did not pass the data on.

## System action

The home page (logon page) is returned to the requestor.

## User response

Inform your network people about the probable proxy problem.

---

**FCXxxx599I**      **Max. number (*nnn*) of connections reached (*ipaddr*)**

## Explanation

An attempt was made to establish a new TCP/IP connection to the Performance Toolkit web interface from IP address *ipaddr*, but the maximum allowed number of connections was already active.

## System action

Operation continues.

## User response

Try again later to establish the connection. You can also try increasing the maximum number of connections by using the FC MONCOLL WEBSERV MAXCONN command. If the problem persists, restart the Performance Toolkit web interface using the FC MONCOLL WEBSERV command.

---

**FCXxxx601E**      **cccccccc trace already active**

## Explanation

An 'FC TRACE cccccccc ON' command has been entered, but the selected trace was already active.

## System action

Operation continues.

## User response

None, trace is already active.

---

**FCXxxx602E**      **cccccccc trace is not active**

## Explanation

An 'FC TRACE cccccccc OFF' command has been entered, but the selected trace was not active.

## System action

Operation continues.

## User response

None, trace is already deactivated.

---

**FCXxxx603I**      **cccccccc trace activated**

## Explanation

An 'FC TRACE cccccccc ON' command has been entered, and the message indicates that the selected trace is now activated.

## System action

Operation continues.

## User response

None.

---

**FCXxxx604I**      **cccccccc trace deactivated**

## Explanation

An 'FC TRACE cccccccc OFF' command has been entered, and the message indicates that the selected is now deactivated.

## System action

Operation continues.

## User response

None.

---

**FCXxxx605I**      **User *username* at *ip-address*  
logged off**

## Explanation

This message indicates that user *username* has clicked the "Logoff" button and terminated the web session.

## System action

All control blocks are released. That is, when the user wants to send another request a new logon is required first.

## User response

None.

---

**FCXxxx610E**      **Different path used - please logon  
again**

## Explanation

A GET request was received by the web interface, but the request claims to be a continuation of a session previously active from another IP address. This can happen, for instance, when a user decides to use a different proxy server for setting up Internet connections. A new logon is forced to preserve some security.

## System action

The home page (logon page) is returned to the requestor.

## User response

Provide your user ID and password again, and restart your session.

---

**FCXxxx612E**      **No TCP/IP data available**

## Explanation

The TCP/IP performance data menu was selected for display, but no TCP/IP monitor records were received, so far. TCP/IP monitor record generation requires:

- TCP/IP 3.1 or later,
- setting up the TCP/IP server machine(s) so that monitor data will be and
- enabling CP monitor SAMPLE and EVENT APPLDATA domain for the TCP/IP server.

## System action

The selection is not accepted.

## User response

Let your systems programmer set up the TCP/IP server machines so that they can generate APPLDATA for CP monitor.

---

**FCXxxx613E**      **Missing cccccccc config. data -  
APPLDATA EVENT based**

## Explanation:

A TCP/IP server configuration display was selected. However, some of the configuration information cannot be shown because it was never received from CP monitor. Since configuration data is generated by TCP/IP servers only during initialization, or when CP monitor is started, such data will usually be missing whenever Performance Toolkit data collection is activated after starting CP monitor.

## System action:

Most of the configuration data cannot be shown.



**User response:**

None.

---

**FCXxxx614A      SELECT pending for more than *nn* minutes**

---

**Explanation**

The web interface sets up its SELECT requests to the TCP/IP server machine so that it should end after one minute even if there is no client request pending. The message tells us that the normal wait time has been exceeded, probably indicating some sort of problem with the TCP/IP server machine.

**System action**

None.

**User response**

Check whether the TCP/IP stack machine is still active and restart it if necessary.

---

**FCXxxx615I      Attempting restart of *ccccccc* interface**

---

**Explanation**

The *ccccccc* interface was deactivated due to some problem – for example, because the TCP/IP server machine was deactivated or restarted. *ccccccc* will be either WebServer or TCPRequest, depending on which TCP/IP interface is being restarted. Since the deactivation was not the result of a user command, the program will attempt to reactivate the interface.

**System action**

Up to five initial restart attempts will be made in one minute intervals; if not yet successful further single attempts will be made each full hour.

**User response**

Make sure the TCP/IP server machine is restarted, if necessary.

---

**FCXxxx616E      Not authorized to use Diag. X'84'**

---

**Explanation:**

The web interface was started to use CP passwords for validating user ID and password pairs, but the machine where the Performance Toolkit is running is not authorized to use DIAGNOSE code X'84', which is a prerequisite for validating CP passwords.

**System action:**

Users are not allowed to connect to the web interface and performance data cannot be retrieved via web browsers.

**User response:**

Obtain the required authorization for DIAGNOSE code X'84' for the PERFSVM machine or indicate some other method for user ID validation when activating the web interface.

---

**FCXxxx617A      Cannot ACCEPT connection, all *nn* sockets in use**

---

**Explanation**

A user tried to connect to the web interface, but all *nn* sockets defined for the interface were already in use and the additional connection could not be accepted.

**System action**

No new connections can be accepted until some of the existing ones have been terminated.

**User response**

Try again later. Inform your system programmer of the problem if it persists.

---

**FCXxxx620E      Selection not valid for trend file scan mode**

---

**Explanation**

You tried to set the performance data display period to 'current' while in trend file scan mode, but only 'average' data can be shown in this mode.

**System action**

The command is ignored.

**User response**

Reduce the interval for creating trend file records if you need performance data for shorter periods.

---

**FCXxxx621E      Graphics not supported for this file type**

---

**Explanation**

You tried to create graphics from a file that is not supported as input for graphics in the current level of the program.

**System action**

The command is ignored.

## User response

Try building your graphics based on one of the supported input file types.

---

**FCXxxx622E**      **No data found for variable cccccc**

## Explanation

You tried to display a variable for which no valid value could be found, because no matching trend file record is available, or because the value could not be calculated (e.g. due to zero-divide).

## System action

The value is not displayed.

## User response

Make sure you select only data from trend records that are available.

---

**FCXxxx623E**      **Command too long: omit some variables/specifications**

## Explanation

You were working with the graphics selection menu and specified so many variables and/or other specifications that the resulting command string became too long (exceeding the max. length of a command line).

## System action

The graphics selection menu is shown again.

## User response

Omit one or more variables, or some other specification(s), from your selection and try again.

---

**FCXxxx630E**      **No Reusable Server Kernel data available**

## Explanation

You tried to select the Reusable Server Kernel displays selection menu but no such menu could be built because no RSK data were collected.

## System action

The RSK selection menu cannot be shown.

## User response

Make sure that the RSK server is set up to produce APPLDATA monitor records, and that the current CP

monitor settings include the collection of APPLDATA records for the RSK server(s).

---

**FCXxxx631E**      **No Linux data available**

## Explanation

You tried to select the Linux performance data selection menu LINUXMEN, but no Linux internal performance information is currently available because the system has not been set up to produce such data.

## System action

The Linux performance data selection menu is not shown.

## User response

Inform your systems programmer of the problem and make sure that the prerequisite setup is made. See "Monitoring Linux Internal Performance Data" in the *z/VM: Performance Toolkit Guide* for more information.

---

**FCXxxx635E**      **Invalid IP address found: 'ccccccc'**

## Explanation

An invalid IP address has been found in the FCONX LINUXUSR file, where ccccccc is the string that was found to be not valid as an IP address.

## System action

The incorrect statement is skipped.

## User response

Check the control file and fix the Linux system entry with the incorrect IP address.

---

**FCXxxx636I**      **No Linux systems defined**

## Explanation:

A Linux-related performance display has been selected but no Linux systems are currently defined. Performance Toolkit requires that any Linux system that you want to retrieve data from be first defined in control file FCONX LINUXUSR, and that the TCP/IP interface needed for data retrieval be activated.

## System action:

None.

## User response:

Have the prerequisites checked/implemented, then re-issue the command.

---

**FCXxxx637E**      **TCP/IP request interface not active**

### Explanation

A Linux related performance display has been selected but the prerequisite TCP/IP interface is not currently active. Linux performance data are retrieved via TCP/IP connection, and no such data can be made available without first activating the request interface.

### System action

The request is ignored.

### User response

Have the TCP/IP request interface activated, then re-issue the command.

---

**FCXxxx638E**      **No free sockets on TCP/IP request interface**

### Explanation

The TCP/IP request interface has run out of sockets and cannot open yet another connection. The current implementation provides for a maximum of 49 such connections; seeing this message proves that this number is either insufficient, or that sockets were not freed after use due to some software problem.

### System action

None.

### User response

Contact IBM Service.

---

**FCXxxx639E**      **No IP address defined for Linux system ssssssss**

### Explanation

A Linux performance data screen has been selected for system ssssssss, but this system has not been defined and, since its IP address is thus not known, no data can be retrieved.

### System action

None.

### User response

Have your systems programmer include the Linux system and its IP address in file FCONX LINUXUSR.

---

**FCXxxx640E**      **RC = rc for system at ipaddr, more\_info**

### Explanation

A Linux performance data screen has been selected for system *system* at IP address *ipaddr*, but the connection could not be established (TCP/IP CONNECT request ended with return code *rc*). *more\_info* is a further explanation of *rc* (if it is available). Probable cause: either the target system and/or DDS interface are not active, or there is a problem with the TCP/IP network.

### System action

The connection is reset.

### User response

Have your systems programmer check the status of the Linux system and its DDS interface.

---

**FCXxxx641E**      **Unexpected call: cccccccc**

### Explanation

An unexpected TCP/IP data call cccccccc has been received by the Linux performance data interface. This is a probable logic error.

### System action

The current connection is reset.

### User response

Contact IBM Service.

---

**FCXxxx642I**      **Connecting to ssssssss at iiii**

### Explanation:

A Linux performance data display for a new Linux system ssssssss has been selected. Performance Toolkit is trying to open a connection to its IP address iiii.

### System action:

Operation continues.

### User response:

None.

---

**FCXxxx643I**      **Retrieving performance metrics**

### Explanation:

A Linux performance data display for a new Linux system has been selected. The TCP/IP connection could be established, and Performance Toolkit is now trying to retrieve performance metrics from the system.

**System action:**  
Operation continues.

**User response:**  
None.

---

**FCXxxx644E**      **Unexpected data format received from ssssssss**

---

### Explanation

The data returned by Linux system ssssssss does not have the expected XML format.

### System action

The connection is reset.

### User response

Make sure that the IP address specified in file FCONX LINUXUSR for system ssssssss is correct, and that a recent level of the DDS interface is active on the system.

---

**FCXxxx645E**      **cccc...cccc**

---

### Explanation

An error message issued by the Linux DDS interface has been found in the returned data. The string cccc...cccc is a copy of (the beginning of) that string. One possible cause is the selection of Linux performance data for a period where that information has not been collected by the Linux DDS code.

### System action

Operation continues.

### User response

Make sure you are not trying to retrieve data that are not collected by the Linux system.

---

**FCXxxx646E**      **Error rc trying to obtain socket for ssssssss**

---

### Explanation:

A Linux performance display for system ssssssss has been selected, and Performance Toolkit tried to obtain a socket and establish a TCP/IP connection to that system. However, the SOCKET request failed with return code rc and the connection could not be established.

### System action:

Data cannot be shown currently for the selected system.

### User response:

Inform your system programmer.

---

**FCXxxx647E**      **cccccccc interface has already been activated**

---

### Explanation

A command has been entered to activate the TCP/IP cccccccc interface, but the interface appears to have been already activated.

### System action

The command is ignored.

### User response

If the interface is not responding then deactivate it first before trying to activate it again.

---

**FCXxxx648I**      **REQUEST TO cccccccc CANCELLED DUE TO EXCESSIVE WAIT**

---

### Explanation

A request for data retrieval was sent to cccccccc but has not been answered for a long time, so the request has been cancelled.

### System action

The retrieval request is cancelled.

### User response

Make sure the RMF PM DDS interface in the target Linux system is active, and that it can be reached via TCP/IP, then try again.

---

**FCXxxx656E**      **'cccccccc' statement in master file not recognized**

---

### Explanation

The program has been invoked for batch processing and the master control files contains the invalid specification 'cccccccc'.

### System action

Operation continues.

### User response

See the "MASTER File" section of the "Operation in Batch Mode" chapter in the *z/VM: Performance Toolkit Guide* for a description of valid master file specifications and correct the statement.

---

**FCXxxx657E**      **Control file fn ft fm not found**

---

## Explanation

The control file *fn ft fm* specified in a command or control file has not been found. This file either does not exist or is not currently accessed.

## System action

Batch operation ends.

## User response

Make sure the control file is available.

---

<b>FCXxxx658I</b>	<b>No valid report specifications found</b>
-------------------	---

---

## Explanation

The program was invoked for operation in batch mode, but no valid performance report specifications were found in the REPORTS control file.

## System action

Only the default MONDATA report will be generated.

## User response

Prepare or update the REPORTS control file if you expect more than the default MONDATA report to be created

---

<b>FCXxxx661I</b>	<b>No valid report selections processed - generating default MONDATA</b>
-------------------	--

---

## Explanation

The program was invoked for operation in batch mode, but no reports were generated because no valid performance report specifications were found in the REPORTS control file.

## System action

Only the default MONDATA report is generated.

## User response

If you want more than the default MONDATA report to be created, update the REPORTS control file.

---

<b>FCXxxx676I</b>	<b>File <i>fn ft fm</i> reloaded.</b>
-------------------	---------------------------------------

---

## Explanation

A 'FCONTROL RELOAD' command has successfully completed.

## System action

The reload is completed and processing continues. The definitions in the new file are now in effect.

## User response

None.

---

<b>FCXxxx677E</b>	<b>Not authorized for S&amp;FSERV on Resource <i>resource</i></b>
-------------------	---

---

## Explanation

An attempt to establish an APPC/VM connection to resource *resource* posing as an intermediate store and forward server was detected, but the virtual machine has not been defined as an S&F server in the FCONRMT AUTHORIZ file.

## System action

Store and forward requests are not allowed.

## User response

Update the FCONRMT AUTHORIZ file on the client machine to include the required S&FSERV definition.

---

<b>FCXxxx678E</b>	<b>Error processing file <i>fn ft fm</i></b>
-------------------	--

---

## Explanation

One or more errors were encountered while performing an 'FC RELOAD' subcommand.

## System action

The reload is completed, using valid entries from the file.

## User response

Review the console log for error messages associated with the file being reloaded.

---

<b>FCXxxx679E</b>	<b>User login failed from <i>ip-address</i> for <i>userid</i> - ESM not available.</b>
-------------------	--

---

## Explanation

An attempt to logon to *userid* from *ip-address* resulted in return *rc* from DMSPASS, indicating that the External Security Manager (ESM) was not available. This should be accompanied by message FCXxxx680 with more information.

## System action

The home page (logon page) is returned to the requestor.

## User response

Try again later. If the problem persists, contact your support team to have them address the problem.

---

**FCXxxx680I**      *information*

## Explanation

This message contains log data information from the ESM, after the DMSPASS call to authenticate a user who is trying to log on from the web. This message should be accompanied by message FCXxxx679E or FCXxxx692E.

## System action

The user is not able to log on.

## User response

None. This is additional information for your support personnel.

---

**FCXxxx681I**      **Invalid parameter 'IDTEST' ignored**

## Explanation

IDTEST is no longer a valid parameter for the FC MONCOLL WEBSERV command. It has been ignored.

## System action

Processing continues using the user authentication available on the system.

## User response

Remove the IDTEST parameter and other associated parameters.

---

**FCXxxx682E**      **Not authorized to issue DIAGx88, WEBSERV not enabled**

## Explanation

While trying to start the WEBSERV interface, the virtual machine running Performance Toolkit was not authorized to issue DIAGNOSE X'88' for user and password authentication.

## System action

The WEBSERV interface is not started.

## User response

Have the Performance Toolkit machine authorized for password checking.

---

**FCXxxx683I**      **Unexpected remsend data from userid at nodeid ignored**

## Explanation

Performance Toolkit received data sent from another node (*nodeid*) using the FC MONCOLL REMSEND command. The FCONRMT SYSTEMS file on this node does not have an entry for this other node.

## System action

The data is ignored and processing continues.

## User response

Remove the REMSEND from the sending node's FCONX \$PROFILE or add information about the sending node to the FCONRMT SYSTEMS file.

---

**FCXxxx684E**      **File 'FCXBANNER ft fm' is too large**

## Explanation

Performance Toolkit cannot process the FCXBANNER HTML or FCXBANNER GIFBIN banner file because its size exceeds the allowed limit (approximately 11.5K).

## System action

Until this file is replaced, it will not be used for banner setup during web interface initialization.

## User response

Replace the banner file with a smaller file. To start the web server interface again, issue FC MONCOLL WEBSERV OFF and then FC MONCOLL WEBSERV ON.

---

**FCXxxx685E**      **ERROR nn: File 'FCXBANNER HTML fm' has invalid structure**

## Explanation

Performance Toolkit found an error while processing the FCXBANNER HTML web banner file. The reason code *nn* has one of the following explanations:

- 1**      The file has no <body> tag.
- 2**      There is no src= attribute on the <img> tag.

3  
There is a premature end of parameter for the  
src= keyword.

System action

Until the error is fixed, the banner file will not be used  
for banner setup during web interface initialization.

User response

Update the file as needed. To start the web server  
interface again, issue FC MONCOLL WEBSERV OFF and  
then FC MONCOLL WEBSERV ON.

---

FCXxxx686I	File 'FCXBANNR <i>ft fm</i> ' is used for Web banner processing
------------	--

---

Explanation

Performance Toolkit will use the FCXBANNR HTML  
or FCXBANNR GIFBIN banner file for building the  
customized web interface banner screen. You will  
always receive one message for FCXBANNR HTML  
and then optionally one more message for FCXBANNR  
GIFBIN.

System action

The file will be used for banner setup during web  
interface initialization.

User response

None.

---

FCXxxx687I	ERROR <i>nn</i> : 'Insufficient Storage, default TREND or SUMMARY settings used'
------------	--

---

Explanation

Performance Toolkit didn't have enough storage to  
process the TRENDREC or SUMREC file.

System action

The default settings for the TRENDREC or SUMREC file  
will be used.

User response

Increase the virtual storage for the Performance  
Toolkit machine.

---

FCXxxx688I	ERROR <i>nn</i> : 'Error reading file, default TREND or SUMMARY settings used'
------------	--

---

Explanation

Performance Toolkit encountered a problem reading  
the TRENDREC or SUMREC file.

System action

The default settings for the TRENDREC or SUMREC file  
will be used.

User response

Verify that the TRENDREC or SUMREC file is usable.

---

FCXxxx689A	Insufficient storage for PERFKIT to continue - exiting
------------	---

---

Explanation

Performance Toolkit requires a certain amount of  
virtual storage for the program itself, and might require  
a considerable amount of additional storage to hold its  
data tables when performance monitoring is activated.  
The required storage is *not* available.

If the emergency safeguard feature (ESF) has  
been activated previously, Performance Toolkit will  
initiate actions processing as defined with the  
FC EMERGENC ORDER command (creating printed  
reports, saving trend/summary files, executing a user  
EXEC procedure) prior to program termination.

This message is also sent to the maintenance user ID  
defined using the FC MAINTID command.

System action

If the ESF is activated, Performance Toolkit will try  
to perform all required actions, and then terminate.  
If ESF is not activated, Performance Toolkit will  
terminate immediately.

User response

Increase the storage size of your virtual machine, re-  
IPL ZCMS, and then restart Performance Toolkit.

---

FCXxxx690E	User login failed from <i>ip-address</i> for <i>userid</i> Read Error
------------	--

---

Explanation

User *userid* on IP address *ip-address* tried to log on  
to the Web Server interface, but the External Security  
Manager (ESM) encountered a read error while trying  
to access the user directory.

## System action

The home page (logon page) is returned to the requestor.

## User response

Contact your support team to have them address the problem.

---

<b>FCXxxx691E</b>	<b>User login failed from <i>ip-address</i> for <i>userid</i> CP is backlevel</b>
-------------------	---

---

## Explanation

User *userid* tried to log on to Performance Toolkit through the web from IP address *ip-address*, but the external security manager (ESM) could not verify the user ID and password because CP was not at the required level.

## System action

The home page (or logon page) is returned to the requester.

## User response

Contact your support team to have them address the problem.

---

<b>FCXxxx692E</b>	<b>User login failed from <i>ip-address</i> RC = <i>rc</i> from DMSPASS</b>
-------------------	---

---

## Explanation

A user attempted to log on to Performance Toolkit through the web from IP address *ip-address*, but the external security manager (ESM) could not verify the user ID and password and returned an unexpected return code. There might be an accompanying FCXxxx680I message with additional information.

## System action

The home page (or logon page) is returned to the requester.

## User response

Contact your support team to have them address the problem.

---

<b>FCXxxx693I</b>	<b><i>operation</i> - <i>explanation</i></b>
-------------------	--

---

## Explanation

An error condition was found during a TCP/IP call, as described in message FCXxxx576E. *operation* tells

what operation was being attempted, followed by the explanation of the error code.

## System action

Depends on the error found. Operation will usually continue, but you may have to restart the Web Server interface.

## User response

Have the error condition corrected and try again.

---

<b>FCXxxx694E</b>	<b>User login failed from <i>ip-address</i> <i>userid</i> id or pw not valid</b>
-------------------	--

---

## Explanation

User *userid* on IP address *ip-address* tried to log on to the Web Server interface, but the *userid*/password combination is not valid.

## System action

The home page (logon page) is returned to the requestor.

## User response

Determine the correct *userid* and password and try again.

---

<b>FCXxxx695E</b>	<b>User login failed from <i>ip-address</i> <i>userid</i> logonby <i>byuserid</i></b>
-------------------	---

---

## Explanation

A password provided during the 'Logon by' procedure is incorrect.

## System action

The home page (logon page) is returned to the requestor.

## User response

Re-enter the user ID(s) and password.

---

<b>FCXxxx698E</b>	<b>You are not authorized to act as <i>userid</i></b>
-------------------	---

---

## Explanation

A user attempted to log on to the *userid* user ID, but the 'Logon by' authorization was not allowed.



### System action

The home page (logon page) is returned to the requestor.

### User response

Re-enter the user ID(s) and password.

---

**FCXxxx696E      Your password has expired.**

---

### Explanation

You tried to log on with a password that has expired.

### System action

The home page (logon page) is returned to the requestor.

### User response

Follow your installation procedure for having your userid reinstated.

---

**FCXxxx697E      The userid or password you entered is not valid**

---

### Explanation

You tried to log on with a userid/password combination that is not valid.

### System action

The home page (logon page) is returned to the requestor.

### User response

Determine the correct userid and password and try again.

---

**FCXxxx699E      Error reading the user directory**

---

### Explanation

The external security manager (ESM) encountered a read error while trying to access the user directory. Therefore, Performance Toolkit cannot authenticate your user ID at this time.

### System action

The home page (or logon page) is returned to the requester.

### User response

Contact your support team to have them address the problem. The Performance Toolkit console might provide additional information.

---

**FCXxxx700E      External Security Manager not available. Try later.**

---

### Explanation

You attempted to log on to Performance Toolkit, but the external security manager (ESM) was not available.

### System action

The home page (or logon page) is returned to the requester.

### User response

Try again later. If the problem persists, contact your support team to have them address the problem. The Performance Toolkit console might provide additional information.

---

**FCXxxx701E      PERFKIT Server is not authorized to verify user IDs**

---

### Explanation

The external security manager (ESM) did not allow Performance Toolkit to verify the user ID and password.

### System action

The home page (or logon page) is returned to the requester.

### User response

Contact your support team to have them address the problem. The Performance Toolkit console might provide additional information.

---

**FCXxxx702E      Unexpected error while authenticating your userid**

---

### Explanation

Performance Toolkit received an unexpected return code from DMSPASS while trying to verify your user ID and password.

### System action

The home page (or logon page) is returned to the requester.

## User response

Contact your support team to have them address the problem. The Performance Toolkit console might provide additional information.

---

<b>FCXxxx703E</b>	<b>User login failed from <i>ip-address</i> for <i>userid</i> No Diag88 authorized.</b>
-------------------	---

---

## Explanation

User *userid* attempted to log on to Performance Toolkit through the web from IP address *ip-address*, but the external security manager (ESM) did not allow Performance Toolkit to use DIAGNOSE X'88' to verify the user ID and password.

## System action

The home page (or logon page) is returned to the requester.

## User response

Contact your support team to have them address the problem.

---

<b>FCXxxx704E</b>	<b>Unable to verify your userid/ password CP is back level</b>
-------------------	--

---

## Explanation

The External Security Manager (ESM) was unable to provide the verification of the userid and password because CP was not at the required level.

## System action

The home page (logon page) is returned to the requestor.

## User response

Contact your support team to have them address the problem.

---

<b>FCXxxx705I</b>	<b>Emergency Safeguard Feature activated</b>
-------------------	--

---

## Explanation

In response to an FC EMERGENC ON command, Performance Toolkit's emergency safeguard feature is enabled for detection of emergency events such as a z/VM shutdown or a shortage of virtual memory. Under these conditions, all actions defined with the FC EMERGENC ORDER command will be performed. The main purpose of these actions will be to save

any unsaved collected data. For more information, see [“FCCONTROL EMERGENC” on page 29](#).

## System action

Operation continues. Performance Toolkit will detect the emergency events and take actions as defined with the FC EMERGENC ORDER command.

## User response

None.

---

<b>FCXxxx706I</b>	<b>Emergency Safeguard Feature deactivated</b>
-------------------	--

---

## Explanation

In response to an FC EMERGENC OFF command, Performance Toolkit's emergency safeguard feature is disabled for detection of emergency events. No actions will be taken if an emergency event occurs.

## System action

Operation continues. Performance Toolkit will ignore any subsequent emergency events.

## User response

None.

---

<b>FCXxxx707A</b>	<b>Signal Shutdown event detected - exiting</b>
-------------------	---

---

## Explanation

Performance Toolkit's emergency safeguard feature detected a z/VM shutdown and started performing the required emergency actions as defined with the FC EMERGENC ORDER command (creating printed reports, saving trend/summary files, executing a user EXEC procedure) prior to program termination.

This message is also sent to the maintenance user ID defined with the FC MAINTID command.

## System action

Performance Toolkit will perform all required actions and load a disabled wait PSW with the code X'FFF'.

## User response

None.

---

<b>FCXxxx708E</b>	<b>Error '<i>errcode</i>' activating Emergency Safeguard Feature</b>
-------------------	--

---

## Explanation

An FC EMERGENC ON command was entered, but Performance Toolkit could not activate the emergency safeguard feature because an internal processing error occurred. The error code *errcode* designates the reason for this error.

## System action

Operation continues. The emergency safeguard feature is not activated.

## User response

Contact IBM service.

---

<b>FCXxxx709E</b>	<b>Error '<i>errcode</i>' processing Signal event</b>
-------------------	---

---

## Explanation

The signal event external interrupt was received, but Performance Toolkit could not process it correctly because an unexpected internal processing error occurred. The error code *errcode* designates the reason of this error.

## System action

Operation continues.

## User response

Contact IBM service.

---

<b>FCXxxx710A</b>	<b>Emergency Safeguard: <i>action</i> processing {started   ended} Emergency Safeguard: all processing complete</b>
-------------------	---

---

## Explanation

The first message is issued when *action* processing is started or ended, as part of the emergency safeguard feature's response to the detection of an emergency event. Possible actions and their meanings are as follows:

### PRINT

Create print reports defined with the FC MONCOLL RESET *hh:mm:ssP* command.

### INTERIM

Create print reports defined with the FC MONCOLL RESET *hh:mm:ssI* command.

### TREND

Process and save trend records defined with the FC MONCOLL RESET *hh:mm:ssT* command.

## SUMMARY

Process and save summary records defined with the FC MONCOLL RESET *hh:mm:ssS* command.

## REMPRINT

Create print reports defined with the FC SETEVENT ... PRINTRMT command.

## EXECPROC(*name*)

Call the user-defined *name* EXEC procedure.

The second message is issued when all required emergency actions have been processed.

For more information, see [“FCONTROL EMERGENC” on page 29](#).

## System action

Operation continues.

## User response

None.

---

<b>FCXxxx711E</b>	<b>Emergency Safeguard Feature already active - command ignored</b>
-------------------	---

---

## Explanation

An FC EMERGENC ON command has been entered, but Performance Toolkit's emergency safeguard feature has already been activated.

## System action

The command is ignored.

## User response

None.

---

<b>FCXxxx712E</b>	<b>Invalid action '<i>action</i>' specified</b>
-------------------	---

---

## Explanation

An FC EMERGENC ORDER <...> *action* <...> command has been entered with an invalid action.

## System action

Only actions preceding the first invalid action (if any) are accepted. The rest that follow are ignored.

## User response

Check the syntax of the FC EMERGENC ORDER command and re-enter it with correct action(s).

---

<b>FCXxxx713E</b>	<b>Duplicate action '<i>action</i>' specified</b>
-------------------	---

---

## Explanation

An FC EMERGENC ORDER <...> *action* <...> command has been entered with an action specified more than once.

## System action

Only actions preceding the first duplicate occurrence of the invalid action are accepted. The rest that follow are ignored.

## User response

Re-enter the FC EMERGENC ORDER command, specifying each required action only once.

---

<b>FCXxxx714E</b>	<b>No CMS EXEC procedure name defined</b>
-------------------	---

---

## Explanation

An FC EMERGENC ORDER <...> EXECPROC <...> command was entered, but the CMS EXEC procedure name had not been previously specified with the FC EMERGENC EXECPROC command.

## System action

Only actions preceding the EXECPROC action (if any) are accepted. The rest that follow are ignored.

## User response

Define the CMS EXEC procedure name with the FC EMERGENC EXECPROC command first, and then re-enter the FC EMERGENC ORDER <...> EXECPROC <...> command.

---

<b>FCXxxx715E</b>	<b>CMS EXEC procedure name missing or invalid</b>
-------------------	---

---

## Explanation

An FC EMERGENC EXECPROC command was entered, but the EXEC procedure name was not specified or is too long (more than 8 characters).

## System action

The command is ignored.

## User response

Re-enter the FC EMERGENC EXECPROC command, specifying the name of existing CMS EXEC procedure.

---

<b>FCXxxx716I</b>	<b>ESF enabled, but no emergency actions have been defined yet</b>
-------------------	--

---

## Explanation

In response to an FC EMERGENC ON command, Performance Toolkit is enabled for detection of emergency events, but because no emergency actions have been defined, no actions will be taken if such events occur.

## System action

Operation continues.

## User response

Define the appropriate emergency actions with the FC EMERGENC ORDER command.

---

<b>FCXxxx717I</b>	<b>No emergency actions have been defined - ESF exiting</b>
-------------------	---

---

## Explanation

The emergency safeguard feature has detected an emergency event, but no emergency actions were previously defined with the FC EMERGENC ORDER command.

## System action

Operation continues.

## User response

None at present. But for future emergency events, you should define the appropriate emergency actions with the FC EMERGENC ORDER command.

---

<b>FCXxxx718I</b>	<b>ESF is now in suspended state - no emergency actions allowed</b>
-------------------	---

---

## Explanation

Although the emergency safeguard feature was previously enabled, permanent data collection appears to be turned off. No actions will be taken upon detection of any emergency event.

## System action

If this message was received because permanent data collection was turned off by the FC MONCOLL DCSS/ DISK OFF command, then operation continues with the emergency safeguard feature in a suspended state.

If permanent data collection was turned off by some type of emergency event, then operation terminates.

## User response

To activate the emergency safeguard feature, permanent data collection must be turned on using the FC MONCOLL DCSS/DISK ON command.

---

**FCXxxx719I**      **action processing skipped - not scheduled till the end of this day**

## Explanation

During processing of an emergency event, Performance Toolkit tried to perform a specific action (*action*), but all scheduled actions of that type have already been processed for the day. This is possible if *action* is INTERIM, PRINT, SUMMARY, or TREND, which are scheduled using the FC MONCOLL RESET command, or PRINTRMT, which is scheduled using the FC SETEVENT command.

## System action

Operation continues.

## User response

None.

---

**FCXxxx720I**      **rrr TREND|SUMMARY record(s) X'FCnn' not written - data not available**

## Explanation

During processing of *rrr* TREND/SUMMARY record(s) X'FCnn', the Performance Toolkit determined that data required for building these records is not available. Possible reasons for that are the following:

- The corresponding CP sample monitor domain was not enabled during the whole time interval for which the history record was to be written
- There is insufficient storage to allocate work buffer for these records.

*rrr* is the name of the TREND/SUMMARY record as defined in the file FCONX TRENDREC of FCONX SUMREC by the 'RECORDS' statement, and *nn* is its number.

## System action

Operation continues.

## User response

None.

---

**FCXxxx721I**      **screen\_name not available, requires data\_domain data for CPUPOOLS**

## Explanation

An attempt was made to select the *screen\_name* performance screen for display or printing, but the display could not be prepared or the screen could not be printed due to one or both of these conditions:

1. The prerequisite monitor data in *data\_domain* was not available
2. There was no CPU pool activity during the time interval.

The CPU Pool Configuration Screen requires Monitor Processor Sample data to be enabled and at least one CPU pool to be defined at some point while Performance Toolkit is running.

The CPU Pool Activity Screen requires Monitor Processor Sample data to be enabled and at least one CPU pool with an active member in it.

## System action

The selected screen is not displayed or printed.

## User response

Verify that the Processor Sample domain is activated. If this domain is not activated, enable it using the CP MONITOR command. Then, issue the CP QUERY CPUPOOL command on the system to determine whether any CPU pools exist. If there are no CPU pools defined, there is no data to display for the CPU Pool Configuration Screen or the CPU Pool Activity Screen. If there is no CPU pool with an active member in it, there is no data for the CPU Pool Activity Screen. If the CP QUERY MONITOR and QUERY CPUPOOL commands indicate that there should be data, contact IBM support to report the problem.

---

**FCXxxx768E**      **SEGOUT segment segname is already in use by user: user**

## Explanation

Only one user at on your system can write into the SEGOUT segment. You have attempted to load *segname* in write mode in a second user's virtual machine.

## System action

PERFKIT operation continues, but formatted data collection terminates.

## User response

Have the first user stop using *segname* or define another DCSS to be used as a second SEGOUT segment.

---

**FCXxxx770I**      **Arithmetic value overflow for attribute: *attribute***

## Explanation

An internal processing error occurred in the formatted data collectors. The values are not logged in the SEGOUT segment.

## System action

PERFKIT operation and formatted data collection continues.

## User response

Contact your IBM representative if the problem persists.

---

**FCXxxx771I**      **SEGOUT data collection is off**

## Explanation

This the response to the FC MONCOLL SEGOUT OFF subcommand, indicating that the formatted output collectors are turned off.

## System action

Processing continues.

## User response

None.

---

**FCXxxx772I**      **SEGOUT data collection is active.  
Using segment : *segname***

## Explanation

This the response to the FC MONCOLL SEGOUT ON *segname* subcommand, indicating that the formatted output collectors are turned on and using segment *segname*.

## System action

Processing continues.

## User response

None.

---

**FCXxxx773E**      **SEGOUT data collection is off due to an error, Reason: *reason\_code***

## Explanation

An internal error occurred during formatted data collection processing.

Reason Code	Explanation
5	The Performance Toolkit is not set to Current mode. It is either in Interim mode or Average mode.

## System action

PERFKIT operation continues, but formatted data collection terminates.

## User response

Ensure that the Performance Toolkit is set to Current mode.

---

**FCXxxx774E**      **Insufficient space in SEGOUT segment. Larger segment needed**

## Explanation

The formatted data collectors do not have a large enough segment in which to collect data.

## System action

PERFKIT operation continues, but formatted data collection terminates.

## User response

Use the FCXSEGSZ EXEC to re-estimate the size of the segment. The amount of space in the segment needed by the formatted data collectors depends on the number of users, the number of devices, and so on being monitored in your z/VM system.

---

**FCXxxx775E**      **Derived offset less than zero, SEGOUT segment underflow**

## Explanation

An internal processing error occurred in the formatted data collectors. Initialization of the SEGOUT segment has failed.

## System action

PERFKIT operation continues, but formatted data collection terminates.

## User response

Correct any problems with the SEGOUT segment. Contact your IBM representative if the problem persists.

---

**FCXxxx777I      SEGOUT is not enabled**

---

## Explanation

The segment used by the formatted data collectors is not enabled.

## System action

PERFKIT operation continues, but formatted data collection terminates.

## User response

Issue the FC MONCOLL SEGOUT ON to start collecting data in the segment.

---

**FCXxxx778I      SEGOUT data collection ended normally**

---

## Explanation

The formatted data collectors have stopped processing data without errors.

## System action

PERFKIT operation continues, but formatted data collection terminates.

## User response

A FC MONCOLL SEGOUT OFF was issued to stop collecting data in the segment. Issue FC MONCOLL SEGOUT ON to restart the data collection.

---

**FCXxxx779E      Error locating DCSS, unable to find segment: *segname***

---

## Explanation

The specified DCSS, *segname*, does not exist.

## System action

PERFKIT operation continues, but formatted data collection terminates.

## User response

Use a valid SEGOUT DCSS on the FC MONCOLL SEGOUT command.

---

**FCXxxx780E      Error on DCSS load, unable to load segment: *segname* Storage overlay**

---

## Explanation

There was an error while loading segment *segname*. The CMS SEGMENT LOAD command encountered a problem trying to load the *segname* segment in the PERFKIT virtual machine.

## System action

PERFKIT operation continues, but formatted data collection terminates.

## User response

Issue the command "CMS SEGMENT LOAD *segname*" from the PERFKIT command. Check the error message from that command and correct the problem stated. The most likely cause of this error is that the SEGOUT *segname* is defined at a storage location which overlaps another DCSS (e.g. MONDCSS). This is error message DMSxxx343E from SEGMENT LOAD. If this is the case, you must redefine the *segname* DCSS in a different storage location. After doing that, you can then restart the formatted data collectors.

---

**FCXxxx781E      Error collecting *data* data in shared memory**

---

## Explanation

An internal processing error occurred in the formatted data collectors while collecting data of the type *data*.

## System action

PERFKIT operation continues, but formatted data collection terminates.

## User response

Correct any problems with the SEGOUT segment. Contact your IBM representative if the problem persists.

---

**FCXxxx784E      Command option: *option* is not valid**

---

## Explanation

The option *option* is either not valid for the command or is incorrectly specified.

## System action

PERFKIT operation continues, but formatted data collection terminates.

## User response

Check and correct the command syntax and reissue the command

---

<b>FCXxxx787E</b>	<b>SEGOUT segment name: <i>segname</i> is not valid</b>
-------------------	---

---

## Explanation

The segment name defined for use by the formatted data collectors is not valid.

## System action

PERFKIT operation continues, but formatted data collection terminates.

## User response

Issue the command again with a valid segment name.

---

<b>FCXxxx788I</b>	<b>Interval time value: INTRVLTM unavailable from monitor, RC=<i>rc</i></b>
-------------------	---

---

## Explanation

The CP MONITOR did not return a valid interval time. The return code is the error code returned by the CP MONITOR.

## System action

PERFKIT operation continues, but formatted data collection terminates.

## User response

Correct any problems with the CP MONITOR based on the return code.

---

<b>FCXxxx789E</b>	<b>Error inserting LPAR record in shared memory, FCXOMXSM RC=8</b>
-------------------	--

---

## Explanation

There is no LPAR performance data available from PERFKIT. If z/VM is running in logically partitioned mode on the hardware (first level), this error is likely a result of CP MONITOR not being enabled to monitor processor data. If z/VM is running as a guest on a z/VM system (second level), then this error is a result of no LPAR data being available to a z/VM guest.

## System action

PERFKIT operation continues, but formatted data collection terminates.

## User response

If running first level, enable processor monitoring in CP MONITOR.

---

<b>FCXxxx790E</b>	<b>DCSS initialization failed. Error code: <i>error_code</i> RC=<i>rc</i></b>
-------------------	---

---

## Explanation

The segment defined for use by the formatted data collectors did not initialize correctly. This error can be caused by any of the following:

### Error

#### Code

#### Reason

INIT_DP1	Initialization of the segment data area prefix failed
INIT_DP2	No address for the segment data area prefix returned
INIT_SPR	No address for the SEGOUT segment name can be obtained
NO_DAP	Address of second data area in SEGOUT segment is not valid

## System action

PERFKIT operation continues, but formatted data collection terminates.

## User response

Check that the segment name is correctly defined and that it is large enough to contain your system's data. Use the FCXSEGSZ EXEC to check the size. If everything is correctly defined, contact your IBM representative.

---

<b>FCXxxx791E</b>	<b>SEGOUT collection stopped due to INTERIM or AVERAGE command</b>
-------------------	--

---

## Explanation

An INTERIM or AVERAGE command was issued in the PERFKIT virtual machine. The formatted data collectors must run in CURRENT mode in PERFKIT. The formatted data collection process is stopped.

## System action

PERFKIT operation continues, but formatted data collection terminates.



## User response

In the PERFKIT virtual machine, issue the CURRENT command and restart the formatted data collectors with the FC MONCOLL SEGOUT command.

---

<b>FCXxxx792E</b>	<b>SEGOUT is not enabled. Can not write to Segment <i>segname</i></b>
-------------------	---

---

## Explanation

Performance Toolkit cannot write to the *segname* segment, and therefore cannot begin collecting SEGOUT data.

## System action

Performance Toolkit operation continues, but formatted data collection is not enabled.

## User response

Make sure that you have a DCSS correctly defined to collect SEGOUT data. Then, restart the formatted data collectors using the FC MONCOLL SEGOUT command.

---

<b>FCXxxx798E</b>	<b>Error calling SEGOUT Storage Management, incorrect parameter passed</b>
-------------------	--

---

## Explanation

An internal processing error occurred in the formatted data collector's storage management routines.

## System action

PERFKIT operation continues, but formatted data collection terminates.

## User response

Restart the formatted data collectors with a new FC MONCOLL SEGOUT command. Contact your IBM representative if the problem persists.

---

<b>FCXxxx800A</b>	<b>Invalid <i>cccc</i> found for device <i>devno</i></b>
-------------------	--

---

## Explanation:

This message is generated by a trap set up to catch incorrect I/O device service time information which has caused Performance Toolkit abends due to bugs in the microcode. The string *cccc* indicates which service time component is invalid, and *devno* is the device number for which the incorrect service time was returned.

## System action:

Operation continues using the interval time. Service times and percent busy for this device are not valid.

## User response:

Contact your hardware support team and inform them of the problem.

---

<b>FCXxxx801A</b>	<b>RC <i>nn</i> for CMSSTOR OBTAIN from FCX<i>cccc+dddd</i> for <i>mmmmmm</i> DWs</b>
-------------------	---

---

## Explanation

The program FCX*ccc* at hex displacement *dddd* tried to obtain *mmmmmm* doublewords of free storage but did not obtain the space. Return code *nn* is the return code of the CMSSTOR OBTAIN macro and indicates the reason for the error. If the problem was not simply a case of insufficient storage, this return code helps the Performance Toolkit for z/VM developer determine what went wrong. If the problem was insufficient storage, this message can also be followed by message 327A.

This message is also sent to the maintenance user ID defined with the FC MAINTID command. Please note that only the first occurrence per day of this message will be sent to that user ID.

## System action

Operation continues, if possible, but data will be incomplete.

## User response

See the return codes of the CMSSTOR macro for the reason for the error. Increase the virtual storage size of the PERFSVM machine if the problem was insufficient storage.

---

<b>FCXxxx811T</b>	<b>'Proc/Reset <i>call_type</i> call timestamp_info [<i>reset_info</i>']</b>
-------------------	--

---

## Explanation

The message is produced during BATCH operation. It provides information about RESETs that are performed during the BATCH run.

## System action

Processing continues.

## User response

None.

---

<b>FCXxxx812I</b>	<b><i>report_name</i> not available; SMT is not enabled</b>
-------------------	---

---

**Explanation**

The requested report requires simultaneous multithreading (SMT) to be enabled:

- 1. For the monitored system when monitoring live data
- 2. On the system that generated a MONWRITE file that Performance Toolkit for z/VM is processing.

**System action:**

The selected report is not displayed or printed.

**User response:**

See the description of the report. It should state any special prerequisites.

---

**FCXxxx813I      There is no user activity on virtual IFLs during the current interval.**

**Explanation:**

The User Processor Time report displays data for users running on virtual IFL processors. There was

no activity during the current interval. It is possible that data is available when this report is selected in AVERAGE or INTERIM mode.

**System action:**

Operations continue.

**User response:**

None.

---

**FCXxxx999E      Message nr. *nnn* not found in FCXMSG MODULE**

**Explanation:**

Performance Toolkit programming error: The number *nnn* of the error message to be displayed could not be found in module FCXMSG.

**System action:**

Operation continues.

**User response:**

Inform your system programmer.

---

## Appendix A. Graphical Displays of Performance History Data

Many different ways of displaying performance history data graphically are supported. They differ in:

- The origin of the history data to be analyzed
- The variable to be used for the x-scale, and the scale factor
- The type of graphics which is to be generated

but they all use a very similar logic and command structure.

**Origin of history data:** You can use all of the many performance variables which can be displayed on the performance redisplay screen (see “[REDISP](#)” on page 269), the extended history data log display (see “[REDHIST](#)” on page 268), and also most of the variables that are contained in an object's benchmark log files for generating performance graphics. These variables can be found in

- *Performance redisplay buffer* (in storage). It contains data from the most recent measurements for many key performance variables, and it is automatically updated whenever new data are collected by the performance monitor. Its size can be defined by the 'FC MONCOLL REDISP' command (see “[FCONTROL MONCOLL](#)” on page 53) to accommodate data of between 60 and 14,400 of the most recent measurements.

The length of the period covered by these measurements depends on the length of the CP monitor sample interval to which all permanent data collection is linked.

- *Disk files with general system load performance data:*
  - Files *yyyymmdd HISTLOGn* (new) and/or *mmddyy PERFLOGn* (old) with detailed performance data for a single day (one record per measurement).
  - File 'ACUM HISTSUM' (new) and/or file 'ACUM PERFHIST' (old) which contain summary data (one record per hour), but for all days where data collection to history files has been active.

Both detailed and summary files will be available, however, only when data collection to disk had previously been activated by the command 'FC MONCOLL PERFLOG ON' (see “[FCONTROL MONCOLL](#)” on page 53).

- File '*systemid* FCXTREND' with extended summary history data: data for several weeks or months will usually be kept, and the data can include performance information on individual users, I/O devices, etc.

See the “Creating Extended Summary and Trend Files” section in the “Using Performance Toolkit in Performance Monitor Mode” chapter of the *[z/VM: Performance Toolkit Guide](#)* for a more detailed description of extended trend files and their creation.

*Collection of performance history data in disk files is an absolute pre-requisite for long-term trend analysis.*

See also the “Data Collection in History Files on Disk” section in the “Using Performance Toolkit in Performance Monitor Mode” chapter of the *[z/VM: Performance Toolkit Guide](#)* for a more detailed description of history data collection.

- *Benchmark log files:* These can optionally be created while benchmarking specific objects, and most of the log files on disk can be used as input for graphics creation just like the general system load files (user logs such as MTUSRLG, USERLOG, etc. and I/O device log files such as CACHELG, DEVLOG, etc.).

**Note:** The remainder of this section has been adapted to refer mainly to the new HISTLOG files as examples for detailed performance history data, but such references always apply to the old PERFLOG files as well, to *mmddyy* RMONLOGn files created by the central monitoring facility, and also to an object's benchmark log files.

Similarly, any references to the new ACUM HISTSUM files as input for condensed summary data should be interpreted to include also the old ACUM PERFHIST files and extended FCXTREND files.

The default source for all the graphical history displays is the redisplay buffer, other sources must explicitly be defined by the 'FILE *fn ft fm*' arguments of the selected PLOTxxx or GRAPHxxx command.

**Variable to be used for the x-scale:** The horizontal x-scale can be used for plotting

- *Elapsed time.* Depending on the desired amount of detail to be shown (and on the amount of detail available in the source data), different commands must be used for generating the graphical displays:
  - *Detailed graphics*, generated by the GRAPHDET or PLOTDET commands, can show data from up to 140 measurements on one display.
  - *Summary graphics*, generated by the GRAPHSUM or PLOTSUM commands, can show data for a period of up to 12 hours. They can be based on data from the current redisplay buffer (default), or from a 'yyyymmdd HISTLOGn' or benchmark log disk file with detailed performance data.
  - *Trend graphics*, also generated by the GRAPHSUM or PLOTSUM commands, can show data for a period of many months or even years. They are based on data from an 'ACUM HISTSUM' or 'systemid FCXTREND' history file.
  - *Daily performance 'profiles'* can also be generated by the GRAPHSUM or PLOTSUM commands. They will show how performance or system load varied during an average day of the selected period. They are based on data from an 'ACUM HISTSUM' or 'systemid FCXTREND' history file.
- *A selected performance variable.* The GRAPHVAR and PLOTVAR commands allow the creation of 'variable correlation' graphics, which show graphically how up to three Y-performance variables depend on another variable which is plotted on the X-scale.

**Type of display:** Data can be plotted using

- Simple plots based on ordinary display characters. Such plots can be shown on any kind of display terminal, and they do not require any additional software. The values of the different variables shown are distinguished by different plot characters. However, displaying more than two y-variables on the same screen with simple character plots is not recommended as it will usually lead to overloaded pictures which are difficult to interpret.

Use the 'PLOTxxx' commands to generate simple character plots.

- GDDM graphics displays. For host connected users these require both a display terminal with graphics capability and access to GDDM code (Graphical Data Display Manager V.3).

For users connected via the web interface and a web browser, the graphics capabilities and Java applet support of their workstation and web browser are used, and they do not need additional host software to create line graphics.

Different colors and line types are used for the different variables. This kind of graphics will be generated by the 'GRAPHxxx' commands.

The following table gives an overview of the commands used for generating a specific type of performance graphics, and of the type of source data needed.

Table 1. Available performance graphics			
Type of graphics	Source of history data		
	Redisplay buffer	Disk files *) yyyymmdd HISTLOGn or benchmark log file	Disk file *) ACUM HISTSUM or systemid FCXTREND
Detailed history graphics	GRAPHDET	GRAPHDET	-

Table 1. Available performance graphics (continued)

Type of graphics	Source of history data		
	Redisplay buffer	Disk files *) yyyyymmdd HISTLOGn or benchmark log file	Disk file *) ACUM HISTSUM or systemid FCXTREND
Summary history graphics	GRAPHSUM	GRAPHSUM	-
Performance trend graphics	-	-	GRAPHSUM
Daily performance profile	-	-	GRAPHSUM
Variable correlation graphics	GRAPHVAR	GRAPHVAR	GRAPHVAR

**Note:**

\*)

These history files will be available only for periods where data collection to disk had been activated.

The GRAPHxxx commands included in the above table will cause full GDDM graphics to be generated. Equivalent PLOTxxx commands are available which allow creation of simple plots for exactly the same data, these can be used if either GDDM or a graphics terminal are not available.

Examples for each of the different types of performance graphics will be given on the following pages, mostly in the form of GDDM graphics, but also including some simple plots.

**Some common features:**

- *Y-scale used:* Up to four (three for GRAPHVAR) of the many performance variables can be selected for concurrent display. Their values will be plotted on the vertical y-axis, and the scale automatically adapted so that
  - the highest value found for a variable will still fit into the plot
  - similar variables (same unit-type, e.g. 'pages/sec.') will be plotted using the same scale if they are shown on the same plot.

The scale used is indicated by the numbers shown on the left hand side of the y-axis. If more than one variable is being shown, the value at the top is the scale for the first variable, the next lower one is the scale for the second variable and so on.

Where only a few of the y-values are much higher than the remainder, such values may be ignored so as to allow displaying the bulk of the data in a better scale.

• *Saving / printing GDDM graphics:*

- press the PF-key corresponding to the SAVE function (PF6 if you work with the initial setting) to save the graphics part of the display as file 'xxxxxxx ADMGDF A', where xxxxxxx is the file name previously entered in the File Name field in the bottom line of the screen. The resulting file can later be retrieved by GDDM for re-displaying, or it can be converted and included in a script file for processing by DCF (Document Composition Facility, 5748-XX9).
- press the PF-key corresponding to the PRINT function (PF9 if you work with the initial setting) to save the graphics part of the display as file 'xxxxxxx ADMPRINT A'. It can then be printed using standard GDDM facilities. Select the desired number of copies by entering the corresponding number in the Copies-field in the bottom line.

# Graphics Selection Menu

The graphics selection menu will be shown after entering the GRAPHICS command. It is intended to help you in generating performance graphics based on history data.

FCX128

Graphics Selection MenuPerf. Monitor

General Specifications

Format : GDDM\_

Truncate X-var. at: \_\_\_\_\_

Data origin : STORAGE\_

Y-var. cumulative (Y/N): N

Graphics type : Detailed history\_

Scatter plot (Y/N): Y

Selected period: Last measurements\_\_\_\_\_

Selected hours : \_\_\_\_\_

Selected days : All days\_\_\_\_\_

Variables Selection

Trend Record Selection

Var	Name	Description	Display	Record Details
X	: ACT_____	- Active users	_____	_____
Y *	: CPU_____	- % total CPU	_____	_____
o	: IO/S_____	- SSCH rate/s	_____	_____
=	: _____	-	_____	_____
/	: _____	-	_____	_____

To select graphics

- either key in your choice directly in the fields above

- or enter '?' in the first position of any field for which you want additional information. A detailed selection menu will then be shown with further explanations.

pressing ENTER without any changes will start graphics creation.

Command ==> \_

F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

Figure 229. Graphics selection menu

The screen should essentially be self-explaining:

- Entering incorrect data will lead to a pop-up window being shown which displays the valid selections from which you can then select.
- Changing one field will automatically change other related fields to their default values if the 'old' value (from a previous graphics selection) is not valid for the new combination.
- Fields which are valid for input with the current graphics type selection are highlighted.

Graphics will be generated when the ENTER key is press without any changes to the input fields.

## Field Description:

General Specifications:

### Format

The kind of graphics to be generated. It can be

#### GDDM

for the creation of real GDDM-based graphics

#### PLOT

for the creation of simple plots, using normal display characters.

#### FILE

for creating a 'flat file' on disk, with one record per X-value, and each record containing the matching Y-values in EBCDIC format, and with a leading and trailing double quote sign (") per value. This format is suitable for importing the data into PC applications so that real graphics can be prepared even where GDDM is not available on a system, or where the resulting graphics are to be included in reports written with a PC text processing program.

The data will be written to your A-disk as file FCONDATA IMPORT. The following is an example for such a file:

```
"Time      " "% total CPU " "SSCH rate/s "
"7         " "          90" "          59"
"         " "          98" "          45"
"         " "         135" "          78"
"         " "         179" "          99"
"         " "         274" "         147"
"         " "         277" "         140"
"8         " "         388" "         166"
"         " "         397" "         180"
"         " "         391" "         197"
"         " "         417" "         218"
"         " "         432" "         223"
"         " "         399" "         196"
"9         " "         343" "         155"
"         " "         369" "         170"
...
...
```

### Data origin

The source data which are to be processed. Valid selections are

#### STORage

for graphics based on data from the current day's REDISP buffer

#### File fn ft fm

for graphics based on a performance history file.

### Graphics Type

The desired graphics format. Valid selections are

#### DETAILED history

for a detailed display of about two hours from the current redisplay buffer or from a detailed disk log file, with time being shown on the X-axis (see [“Detailed Performance History Display” on page 961](#)).

#### SUMmary history

for a more condensed display of about 12 hours from the current redisplay buffer, or from one of the detailed history log files, or for any length of time from the accumulated HISTSUM or PERFHIST file (see [“Summary Performance History Display” on page 963](#) and [“Performance Trend Display” on page 964](#)).

#### PROfile history

for displaying load variations during an average day in the selected interval (see [“Daily Performance Profile” on page 965](#)). This selection is valid only when the source data come from a HISTSUM or PERFHIST file.

#### VARIABLE correlation

for displaying up to three performance variables as a function of a another 'base'-variable (see [“Variable Correlation Display” on page 966](#)).

### Selected period

The period for which data are to be analyzed. Valid selection formats depend on the selected data origin:

#### All

to show all of the available data. Automatically translated to LAST when entered with a detailed history selection.

#### Last

to show the last part of the available detailed data. Valid with a detailed history selection only.

#### From hh:mm

indicates the time stamp of the first data record which is to be analyzed from the current redisplay buffer (STORage) or a detailed history file. Not valid with HISTSUM or PERFHIST files.

**From yyyy/mm/dd**

indicates the date of the first data record which is to be analyzed from the selected HISTSUM file. Not valid with HISTLOG or RMONLOG files, or when working with the redisplay buffer.

**To yyyy/mm/dd**

indicates the date of the last data record which is to be analyzed from the selected HISTSUM file. Not valid with HISTLOG or RMONLOG files, or when working with the redisplay buffer.

Both the 'From' and 'To' dates can also be entered in the formats mm/dd/yy or mmddyy.

**Selected hours**

Can be specified only when analyzing summary HISTSUM or PERFHIST data. Hours have to be specified in two different ways, depending on the selected type of graphics:

- *SUMmary history* graphics: Specifies the hours of each day for which data are to be analyzed (typically hours with peak load). Up to eight hours can be selected.
- *PROfile history* graphics: Specifies the first and last of a range of hours for which the average daily load is to be determined, i.e. only two values should be entered.

**Selected days**

Can be specified only when analyzing summary HISTSUM or PERFHIST data. Select the weekday(s) for which performance data are to be included, where your specification can be a combination of the following strings:

**ALL**

if all weekdays are to contribute to the analysis

**M-F**

if only the normal workdays Monday to Friday are to contribute data

**MON**

if all Mondays are to contribute data

**TUE**

if all Tuesdays are to contribute data

**WED**

if all Wednesdays are to contribute data

**THU**

if all Thursdays are to contribute data

**FRI**

if all Fridays are to contribute data

**SAT**

if all Saturdays are to contribute data

**SUN**

if all Sundays are to contribute data

All of the available data will be analyzed when no specific weekdays are specified.

When specified, only performance data collected for the selected weekdays are used for calculating the averages which will eventually be plotted.

Example: Specifying 'M-F SAT' will cause data to be selected for the weekdays Monday to Saturday, excluding any data for Sundays.

**Truncate X-var. at**

Specifies a maximum value for the X-variable above which measurements are to be ignored. The specification is valid for variable correlation graphics only; it will be ignored in all other cases.

This option has been introduced to give the user some means to influence the X-scale in cases where a few extreme X-values would otherwise make Performance Toolkit use a scale that is not suitable for the remainder of the data.



**Scatter plot (Y/N)**

Specifies whether dots are to be included in the graphics output for *all* measurements. The specification is meaningful only for 'variable correlation' graphics with a single Y-variable; it will be ignored in all other cases.

*Variables:*

**X**

Specifies the name of the performance variable which is to be plotted on the X-scale. Valid for variable correlation graphics only and ignored in all other cases.

**Y \***

Specifies

1. the display character to be used when plotting the variable in PLOT format, that is, non-GDDM graphics. The specification is ignored when creating GDDM graphics. The initial setting is '\*' or '=' for Y-variables 1 to 4, as indicated in the sample screen above.
2. the name of the Y-variable which is to be plotted.

What variables can be selected depends on the data origin specified. You can choose any of the many variables which are available on the general performance redisplay screen (see “REDIS” on page 269), and any 'user variables' which you have defined using the 'FC USERVAR ...' command (see “More about graphics” on page 968 for more information) as an X- or Y-variable if data from the redisplay buffer or from a HISTLOG/HISTSUM or PERFLOG/PERFHIST file are to be processed. When processing data from a benchmark log file, you will have to select from the variables available in that file.

Entering a question mark (or any invalid variable name) will let a prompting window be shown with the names and descriptions of the available variables. You can then select the variable directly from that window.

A short description of the selected variable will be appended when the variable name has been read and accepted.

**Y-var. cumulative (Y/N)**

Specifies whether the values of the selected Y-variables are to be added to each other for plotting.

This makes sense only if all of the selected Y-variables are similar (same units, e.g. '%') and if their total is also a meaningful value (see also “More about graphics” on page 968).

## Detailed Performance History Display

---

This type of graphics display will be generated by the GRAPHDET or PLOTDET commands. It allows creating detailed plots with the values of up to four variables. This often gives a much better overview over general system behaviour than the original figures as seen on the redisplay or benchmark screen.

## Example Using GDDM Graphics




Figure 230. Detailed performance history display with GDDM graphics

The example shown was created with the command 'GRAPHDET CPU C1ES', causing total CPU load and the class 1 elapsed time slice to be plotted in the form of a full GDDM graphics display. These are also the default variables which will be displayed if the GRAPHDET command is entered without additional arguments.

Data from up to 140 measurements can be shown in one display (the average values of 2 consecutive measurements are used to generate a total of up to 70 plot points if the default monitor interval of 60 seconds is used). If more data are available, you can select the period to be analyzed by specifying the time stamp of the first measurement to be shown by means of the

```
'FROM hh:mm:' argument.
```

The last measurements will be shown if no start time is specified.

Note that different colors and also different line types will be used for the graphical representation of the selected variables, and that the same colors and line types are used again for the variable descriptions and the underline-type in the title line.

See “GRAPHDET / PLOTDET” on page 182 for a detailed description of the GRAPHDET command.

## Example Using Simple Character Plot

The following example shows the kind of display which will be generated by the 'PLOTDET' command: a detailed display based on data from the current redisplay buffer, like the example on the preceding page, but using ordinary display characters for plotting instead of GDDM graphics.




Figure 231. Detailed performance history display using simple plot

This display was generated by entering the command 'PLOTDET CPU IO/S', i.e. it shows the total CPU load and the total I/O rate for the system.

Note that there are two big differences between the output created by the PLOTDET command and the previous GRAPHDET example, in addition to the obvious difference in the graphic aspect:

1. If sufficient data are available in the redisplay buffer, i.e. if it's size has been increased from the default 120 entries by means of the 'FC MONCOLL REDISP nnn' command, the remainder of the buffer contents can also be displayed graphically simply by shifting the screen window to the left or right, using either the 'LEFT' and 'RIGHT' sub-commands, or the corresponding PF-keys.
2. If the most recent data have been selected for display (default), the plot shown will automatically be updated with each new measurement. If you can spare a display terminal, you could so have a continuous graphic display of some key performance variables.

These features are available only with the detailed plots generated by the PLOTDET sub-command, but *not* with GDDM graphics generated by the GRAPHDET command.

## Summary Performance History Display

Summary history displays can be generated by the GRAPHSUM and PLOTSUM commands. They can be based either on data contained in the current redisplay buffer, or on data from a detailed 'yyyymmdd HISTLOGn', 'mmddyy PERFLOGn' or 'mmddyy RMONLOGn' file, or from one of the object's benchmark log files. Summary history displays allow plotting of the values of up to four redisplay variables for a period of up to 12 hours.




Figure 232. Summary performance history display using simple plot

The example above has been created using the command 'PLOTSUM TR/S CHAR T', i.e. the trivial transaction rate on the system has been plotted. A bar-chart type of plot has been created because only one variable was selected, and the character 'T' was used for plotting the variable as specified in the command. Since no 'FILE' argument has been included, data from the current redisplay buffer have been used as input.

Averages are calculated and plotted for each ten minute interval. Data from up to 12 hours can thus be shown in one plot.

See “[GRAPHSUM / PLOTSUM](#)” on page 185 for a detailed description of the GRAPHSUM command.

## Performance Trend Display

Trend analysis graphics are based on performance history files 'ACUM HISTSUM' or 'ACUM PERFHIST'. Depending on the selected period, the average load for one day or for several days will be represented by a single plot position, thus allowing periods of several months or even years to be included in a single display. The resulting graphics can so show performance trends over long periods, and they are a good basis for capacity planning.




Figure 233. Layout of performance trend display (GDDM)

The above example has been created by entering the command 'GRAPHS CPU IO/S FI ACUM PERFHIST FR 010103 TO 122303 H 8 9 10 14 15 16', i.e. the trends for total CPU load and the I/O request rate have been plotted for the year 2003. The 'HOURS' selection has been added to make sure that the resulting graphics show the trends for heavily loaded periods of the day only (averages which include many hours of low activity are of little help for capacity planning).

Be careful and make sure you are really plotting the data you want:

- All days found within the selected period are going to contribute to our trend graphics. Suppress weekends by specifying the DAYS argument of the GRAPHSUM command, and delete other records for days with non-representative load (such as the public holidays of January 1/2 and Easter in the example), if necessary, before generating trend graphics which are to serve as a base for capacity planning.
- All hours for which data are available in any given day will contribute to this day's average unless you specify the 'HOURS' argument of the GRAPHSUM or PLOTSUM command. If your history file includes data from many hours with low activity, these will cause correspondingly low averages to be generated if you do not explicitly select peak hours.

See "GRAPHSUM / PLOTSUM" on page 185 for a detailed description of the GRAPHSUM command.

## Daily Performance Profile

This type of display is based on history data from 'ACUM HISTSUM' or 'ACUM PERFHIST' summary files. Instead of using the data for analyzing trends in the average load per day as in the 'trend' type of display, you can also show the average load during each hour of the day within the selected period. The result is a performance 'profile' which illustrates how system load changes with each hour of an average day.




Figure 234. Layout of daily performance profile display (GDDM)

The example above has been created by the command 'GRAPHSUM TR/S NT/S FILE ACUM PERFHIST FROM 010103 TO 122303 PROFILE 6 19', i.e. the display shows the transaction rate for trivial and non-trivial transactions (UP and MP) on the system for the period of 06:00 to 19:59 of an average day of the year 2003 (these transaction rates are good indicators of user activity).

### Warning:

All days found within the selected period are going to contribute to our average day's values, and this may include weekends and holidays. While you can exclude weekends by means of the 'DAYS' specification, you may have to check the source file and delete records for holidays with non-representative load before generating performance profiles.

The example above is based on the same data as the performance trend display on the preceding page.

See "GRAPHSUM / PLOTSUM" on page 185 for a detailed description of the GRAPHSUM command.

## Variable Correlation Display

This display is similar to the history plots and graphics described on the preceding pages insofar as it also uses the data collected in the performance redisplay buffer, or in performance history files on disk, for creating a graphic image of the behavior of different performance indicators. However, while the (horizontal) x-coordinate is always a time scale for the history plots, any of the available performance variables can be chosen as the x-variable for a correlation plot. All other selected variables will then be inserted as y-variables on the vertical y-scale, based on the value of the x-variable in the same measurement interval.

The resulting plots show how the different variables depend on each other. They are, therefore, an excellent tool which can help you in analyzing performance bottlenecks on your current system, and they show trends which can tell you what to expect should the current average load increase.

## Variable Correlation Display with Simple Plot

This screen will be shown after entering the 'PLOTVAR' sub-command. It allows plotting of the values of up to three redisplay variables on the vertical y-scale, depending on the corresponding values of the

selected x-variable which will be plotted on the horizontal x-scale (see “GRAPHVAR / PLOTVAR” on page 189 for details).




Figure 235. Layout of screen for variable correlation plots

The above example has been created by entering the command 'PLOTVAR ACT CPU PG/S', thus generating a plot with the x-variable 'ACT' and the y-variables 'CPU' and 'PG/S', i.e. the total CPU load and the page rate have been plotted based on the number of users which were active during the corresponding measurement intervals.

Plots based on the number of active users are particularly useful when planning for the future of systems where the main load is caused by a high number of similar users (e.g. CMS machines). Make sure they are based on a sufficiently high number of samples, and that they also contain samples taken during peak load. Today's activity level during peak load may be an example of tomorrow's average system activity, and the trends shown for the different performance variables can tell you what to expect when system load is increased.

**Note:** The number of users which were found to have been active (used some virtual CPU time) during a certain interval obviously depends on the length of this interval. Correlation plots based on the number of active users will yield meaningful results only if the same measurement interval was used for collecting all the data which are now used as input. Be careful especially when doing correlation plots based on a summary HISTSUM or PERFHIST file, and make sure the same interval was used for the whole period which is to be analyzed.

**Scales:** All scales are automatically adapted, depending on the maximum values of the variables to be shown.

- The scale for the x-axis is selected depending on the maximum value of the x-variable. The corresponding figures, and the description of the variable, are inserted underneath the x-axis.
- The y-coordinate is used for plotting the values of the dependent y-variables. The scale is automatically adapted according to the same logic used for all other graphic displays.

If only a few x- or y-values are much higher than the remainder, such values may be ignored so as to allow displaying the bulk of the data in a better scale.

Where the resulting plots or graphics are not satisfactory because of an unsuitable x-scale, you can define the extent of the x-scale by means of the 'TRUNC' argument of the PLOTVAR or GRAPHVAR commands. See “GRAPHVAR / PLOTVAR” on page 189 for a detailed description of the two commands.

## Variable Correlation Display Using GDDM Graphics

---

GDDM graphics will be shown when a variable correlation display is requested with the 'GRAPHVAR' command (requires access to GDDM libraries and a display terminal with graphics support). As for the 'PLOTVAR' command, up to three of any of the performance variables shown on the performance redisplay screen can be selected as y-variables whose values are to be plotted for the corresponding values of a selected x-variable. See [“GRAPHVAR / PLOTVAR” on page 189](#) for a detailed description of the GRAPHVAR command.

The average values of the y-variables for each x-variable value are represented by the marker symbols which are also inserted in front of the variable description in the title line. In order to better show trends even where the y-values are widely scattered, averages are also calculated for each of the x-scale grid line values and, where possible, for one intermediate x-value. Where the resulting average carries sufficient weight, i.e. if it is based on a relatively high number of samples, these averages are interconnected using a polygon line.

Different colors and line types are used for the graphical representation of the selected variables, as indicated by the color and underline-type of the variable description in the title line.

When only one y-variable is selected you can also have 'scatter plot' type graphics generated by specifying the argument 'SCATTER'. ALL measurements will then be shown, in addition to the average values which are always displayed.

## More about graphics

---

While going through the preceding sections of this chapter, and thereby becoming an expert on the creation of performance graphics with Performance Toolkit, you may have wished for some additional features:

- Possibility to create graphics based on other than the original set of performance variables
- Possibility to generate 'cumulative' graphics, where the value of each y-variable is added to the sum of the preceding variables

The following paragraphs explain how to do this.

### Definition of 'User Variables'

The USERVAR argument of the FCONTROL command has been added to allow definition of 'User Variables', whose value is derived from the original performance variables in the REDISP screen or in the history files of Performance Toolkit according to some formula which you supply when defining the variable.

Once defined, such user variables can then be selected as input for performance graphics like the original performance variables which we referred to in the preceding sections. See [“FCONTROL USERVAR” on page 94](#) for details on the 'FC USERVAR ...' command.

### Generation of Cumulative Graphics

The CUMULAT argument of the GRAPHxxx and PLOTxxx commands allows creation of 'cumulative' graphics, where the position of each value on the Y-axis is determined by the *sum* of the values of the preceding Y-variables plus the value of the current Y-variable.

Just use the CUMULAT argument when you generate your graphics. The program will then

- translate the values to all use the same scale, and display only one scale value for all variables
- display the different y-values cumulatively, i.e. on top of each other
- insert the string 'Cumul.' over the Y-axis to indicate that cumulative values are being shown
- shade the areas between the different polygon lines which correspond to the variables' values in the same color as the the variables' lines.



Cumulative graphics can make sense only when *similar* variables are selected, i.e. variables whose values have the same units (e.g. percent, seconds, ...). However, Performance Toolkit will not prevent you from ignoring this basic rule, so you are free to create very pretty, but absolutely meaningless pictures ...

See “GRAPHDET / PLOTDET” on page 182, “GRAPHSUM / PLOTSUM” on page 185 and “GRAPHVAR / PLOTVAR” on page 189 for a detailed description of the GRAPHxxx and PLOTxxx commands.

## Example of Cumulative Graphics with User Variables

The basic performance variables with CPU load information are

- CPU, the total CPU load in percent
- %CP, the total CP load in percent (user + system)
- %EM, the total emulation mode load, in percent
- %WT, the total wait time, in percent
- %SY, the total SYSTEM CP time in percent.

We would like to show all this information in the form of a cumulative graphic display, as history graphics for the first half of 2003. Unfortunately we cannot just specify all the variables, because some of them are already contained within others: CPU is the sum of %CP and %EM, so it is of no use for our current purpose, and %SY is already contained in %CP. We want to display the components

- %EM, the total emulation mode load
- %SY, the total SYSTEM CP time
- the total USER CP time
- %WT, the total wait time

where the sum of the first three variables will be the total CPU load (i.e. equal to the original CPU value), and the sum of all four variables should be a constant: the totally available CPU power in percent. What is missing is the 'total user CP time', so we have to define the corresponding user variable for calculating it by the commands

```
FC USERVAR SET USERCP = %CP - %SY      (the formula)
FC USERVAR SET USERCP DESCR USER %CP   (the description)
FC USERVAR SET USERCP UNITS %CPU        (the units shown)
```

The desired graphics would then be created with the command

```
GRAPHSUM %EM USERCP %SY %WT FI ACUM PERFHIST FROM 010103 H 8 9 10 CUMUL
```

which results in the following display:




Figure 236. Layout of GDDM history data display in cumulative form

The sum of variables %EM (emulation time), USERCP (user CP time) and %SY (system CP time) is the total CPU load of the system. The wait time %WT has been added to show how much reserve power is left; in our example the result shows that a fifth processor has been added sometime in May 2003 (since averages are calculated for several days, we see two steps from a total of 400% to 500% instead of just one).

## Appendix B. Numerical Analysis of Performance History Data

The performance history files 'ACUM HISTSUM' and 'ACUM PERFHIST' contain one record per hour with the averages of many general system performance variables for that hour, and the detailed history log files 'yyyymmdd HISTLOGn' and 'mmdyy PERFLOGn' contain detailed data for the same variables, with one record per monitor interval.

You have seen, in the preceding chapter [Appendix A, “Graphical Displays of Performance History Data,”](#) on page 955, how the data in these history files can be analyzed graphically. Here you will find a description of the 'FCONXREP' procedure which is intended for numerical analysis of such history files.

All these files can also easily be read and analyzed by your own programs to produce special customized reports. The procedure 'FCONXREP', written in REXX, can then serve as an example on how to do that. See [Appendix D, “HISTLOG and HISTSUM File Records,”](#) on page 981 for the exact format of the new extended history files if you intend to write your own analysis procedures.

### Numerical Analysis Using FCONXREP Procedure

FCONXREP is a REXX procedure for analyzing summary performance history data files 'ACUM HISTSUM' and 'ACUM PERFHIST', and also the corresponding detailed history logs (HISTLOG and PERFLOG files). It can calculate averages for selected performance variables on an hourly, daily, monthly and yearly level for summary history files, or averages for smaller intervals for detailed history log files, and it will generate output files

- 'FCONXREP LISTING' with the output in table form, suitable for printing, and
- 'FCONXREP IMPORT' with an output format suitable for importing it into a PC application for further processing.


The procedure is not based on any of the components of the Performance Toolkit performance monitor, i.e. it can be executed from standard CMS, without first invoking Performance Toolkit. It can be called for either:

- interactive use, intended for cases where you want to create 'ad hoc' reports sporadically, or
- batch use, which allows you to let specific reports be created automatically (e.g. for the generation of monthly performance overview reports by a disconnected service machine).

Both methods are described here.

### Interactive Use of FCONXREP Procedure

This is the default operating mode. Just call the procedure by entering this command:



The following screen will be shown if a summary history file has been selected for input:

FCX201	Perf. Toolkit Numerical Performance Analysis	Perf. Report
<b>Period: Date</b>	<b>==&gt;</b> 2003/01/01 <b>to</b> 2003/02/28 (YYYY/MM/DD)	
<b>Time</b>	<b>==&gt;</b> 08 <b>to</b> 17 (all hours in period)	
<b>Sel. hours</b>	<b>==&gt;</b> -- -- -- -- (only selected hours)	
<b>Sel. days</b>	<b>==&gt;</b> M-F -- -- -- -- (only selected days)	
<b>Desired detail: HOURS</b>	<b>==&gt;</b> TOT (ALL or TOT)	
<b>DAYS</b>	<b>==&gt;</b> ALL (ALL or TOT)	
<b>MONTHS</b>	<b>==&gt;</b> ALL (ALL or TOT)	
<b>Variable selection</b>	<b>==&gt;</b> CPU Total CPU load (all processors)	
	<b>==&gt;</b> PG/S System page rate (reads & writes)	
	<b>==&gt;</b> C1ES 85% class 1 elapsed time slice	
	<b>==&gt;</b> TR/S Trivial transactions executed per sec.	
	<b>==&gt;</b> NT/S Non-triv. transact. executed per sec.	
	<b>==&gt;</b> -----	
	<b>==&gt;</b> -----	
	<b>==&gt;</b> -----	
Processing file 'ACUM HISTSUM A1' (new SUM format)		
<b>F1=Help F2=Browse F3=Quit F5=XEDIT F10=Execute F12=Return</b>		

Figure 237. Layout of FCONXREP Summary History Files Screen (FCX201)

The menu shown allows selection of a specific subset of the data contained in the history data file. You can specify:

- the time frame you are interested in,
- the amount of detail to be shown, and
- the performance variables for which you want averages to be calculated.

### Selecting Date and Time for the Period to be Analyzed

The 'ACUM PERFHIST' file contains one record per hour of data collection, with the averages for each of the selectable performance values during that hour. Both the creation date and hour are contained in each history record, and they can be specified as key variables to specify the period for which performance data are to be analyzed.

#### Date

the first and the last date of the period for which data are to be analyzed (format *yyyy/mm/dd*).

The period during each day for which performance data records are to be analyzed can be specified in *either* of the following two ways (not both):

#### Time

the first and the last hour of each day within the period for which data are to be analyzed (format *hh*).

Hours include all minutes within that hour, i.e. specify the hours '07' to '18' for a report of the period 07:00 to 18:59.

#### Sel. hours

all hours of each day for which data are to be analyzed (need not be consecutive, format *hh*). This selection would typically be used for specifying some hours with peak load (e.g. '09', '10', '14' and '15'). Enter only selected hours *or* start and end times for a range of hours, *not both*.

#### Sel. days

the weekdays for which data are to be analyzed (MON TUE WED THU FRI SAT SUN, ALL or M-F are accepted values). Specifying weekdays is important especially if the history file includes data for weekends that should be omitted from the analysis.

### Selecting the Amount of Detail to be Shown

Since data for each hour within the selected period are available, the output could show performance information split down to a certain hour's performance on a specific day. More general information will usually be desired, however, so one must be able to specify how much detail is to be shown. This is what the 'Hour', 'Day' and 'Month' specifications are for. Accepted values are 'ALL' and 'TOT', meaning

## HOURS

ALL: generates one output line for each hour of each day within the selected period

TOT: suppresses generation of hourly data

## DAYS

ALL: generates one output line for each day, with the average values of all selected hours of that day

TOT: suppresses generation of daily data

## MONTHS

ALL: generates one output line for each month, with the average values of all days of the selected period within that month

TOT: suppresses generation of monthly data

A selection of 'TOT' on a daily or monthly level is not accepted if a selection of 'ALL' is made on a lower (more detailed) level.

Summaries for each year, and a line with the overall average for the whole of the selected period, will always be included, regardless of the above specifications.

## Selecting Performance Variables

Up to eight performance variables can be chosen. What performance variables are selectable depends on the input file. See “REDHIST” on page 268 for a list of the performance variables available with the new extended history log files, or “REDISF” on page 269 for the variables contained in the old PERFLOG and PERFHIST files. When specifying one of the two-word 'user extremes' variables such as 'MAX. CPU', the separating blank must be replaced by an underscore character.

## Using FCONXREP with a Detailed History Log

When a detailed history log file is specified as input with the FCONXREP command, the following screen layout will be shown:

```
FCX201          Perf. Toolkit Numerical Performance Analysis          Perf. Report

Processing a detailed history log file for a single day. Please specify the
period to be analyzed, the interval to be used for calculating subtotals,
and the variables that should be included in the report.

Start and end time ==> 06:01 to 09:00 (hh Interval (minutes) ==> 20 (Range 1-120 minutes)

Variable selection ==> CPU          Total CPU load (all processors)
                   ==> PG/S        System page rate (reads & writes)
                   ==> C1ES        85% class 1 elapsed time slice
                   ==> TR/S        Trivial transactions executed per sec.
                   ==> NT/S        Non-triv. transact. executed per sec.
                   ==> -----
                   ==> -----
                   ==> -----

Processing file '20030214 HISTLOG2 B1' (new LOG format)

F1=Help  F2=Browse  F3=Quit  F5=XEDIT  F10=Execute  F12=Return
```

Figure 238. Layout of FCONXREP screen for detailed history files

The screen layout is different from the one shown for summary files, adapted to the needs for analyzing a detailed log file.

## Selecting Period and Interval

The detailed history logs contain one record per monitor interval, with the load values for that interval. For the analysis you have to specify


- the Start and end time of the period that should be analyzed, and

- the Interval length, in minutes, for which sub-averages are to be calculated. The interval should be chosen as a multiple of the monitor interval used in the creation of the input file.

No date or weekday selection are possible since one such detailed log file contains data for a single day only.

## Batch Use of FCONXREP Procedure

The FCONXREP procedure can be invoked in batch mode by entering the command



Since the idea is to let the procedure run without user interference, you have to provide the selection criteria (e.g. start date/time, end date/time, and performance variables to be reported) in machine readable form. The procedure expects all of this information to be contained in file FCONXREP SETTINGS, on any of the accessed disks. The following is an example for the FCONXREP SETTINGS file, with specifications corresponding exactly to what was shown in the previous example for the interactive selection menu of the procedure.

```

*** Start and End Date (YYYY/MM/DD)

SDATE 2003/01/01
EDATE 2003/02/28

*** Start and end hour (HH for summary file, HH *STIME 07:01 *ETIME 19:00

*** Specific weekdays (MON, TUE, WED, THU, FRI, SAT, SUN, or
*** M-F for all working days Monday-Friday, or
*** ALL for all days of the week

WEEKDAYS M-F

*** Selected hours (HH, up to 10 hours in a day can be defined)
*** If you define SELHOURS then start and end hours (STIME and
*** ETIME) will be ignored for processing

*SELHOURS 01 03 07 09 21

*** Hours Days Months (ALL or TOT), valid for summary files only

HOURS TOT
DAYS ALL
MONTHS ALL

*** Reporting interval in minutes, valid for detailed logs only

*INTERVAL 15

*** Performance Variables; up to 8 variables can be defined for selection

PERFVARS CPU PG/S C1ES TR/S NT/S

*** END of Input file
  
```

A sample file FCONXREP SETTINGS is supplied with the program. The above is not an exact copy, some comment lines have been removed.

### Explanation of keywords:

The required input is exactly the same as for the interactive menu; it is only explained in abbreviated form here. Refer to the descriptions for the interactive use of the program for more detail.

#### SDATE

The first date for which performance data are to be analyzed (format *yyyy/mm/dd*). The argument is valid for processing of summary history files only.

**EDATE**

The last date for which performance data are to be analyzed (format *yyyy/mm/dd*). The argument is valid for processing of summary history files only.

**WEEKDAYS**

The days of the week for which summary performance data are to be analyzed. Valid day specifications are the strings MON TUE WED THU FRI SAT SUN for single days, ALL for all days and M-F for all working days from Monday to Friday. The argument is valid for processing of summary history files only.

**STIME**

The first hour of each day for which performance data are to be analyzed (format *hh*) for summary files, or the start time in hours and minutes (format *hh:mm*) for detailed history logs.

**ETIME**

The last hour of each day for which performance data are to be analyzed (format *hh*) for summary files, or the end time in hours and minutes (format *hh:mm*) for detailed history logs.

**INTERVAL**

The time interval in minutes for which sub-averages are to be calculated and printed. The argument is valid for processing detailed history logs only.

**SELHOURS**

All hours of each day for which data are to be analyzed (need not be consecutive, entered as up to 10 hours in format *hh*). The argument is valid for processing of summary history files only.

You can only specify STIME/ETIME values **or** SELHOURS values, not both.

**HOURS**

Set to either ALL (values for each hour to be shown) or TOT (suppress generation of hourly data)

**DAYS**

Set to either ALL (values for each day to be shown) or TOT (suppress generation of daily data)

**MONTHS**

Set to either ALL (values for each month to be shown) or TOT (suppress generation of monthly data)

A selection of 'TOT' on a daily or monthly level is not valid if a selection of 'ALL' is made on a lower (more detailed) level. The HOURS, DAYS and MONTHS arguments are valid for processing of summary history files only.

**PERFVARS**

Specify up to eight performance variable names, out of the list given earlier in this section.

## Sample Output

The table with the averages for the selected variables will be written to disk in output file FCONXREP LISTING and, in different format, also in file FCONXREP IMPORT, regardless of whether the procedure is used interactively or in batch mode. The examples shown on the following page are the output created with the data selection made in the sample FCONXREP menu shown in [Figure 237 on page 972](#), or in the sample FCONXREP SETTINGS file, that is, for:

- the two months of January and February 2003
- the period from 8:00h to 17:59h of each day (prime shift)
- 'HOURS = TOT', 'DAY = ALL' and 'MONTH = ALL', i.e. hourly details are suppressed, but we can see each day's and month's average load, in addition to the yearly and the overall average which will always be shown.
- the performance variables

**CPU**

total CPU load

**PG/S**

total system page rate

**C1ES**

85% class 1 elapsed time slice (internal response time)

**TR/S**

the average trivial transaction rate, and

**NT/S**

the average non-trivial transaction rate.

PERFKIT Statistics for Period 2003/01/01 to 2003/02/28, Hours 8 to 17										
Year	Mon	Day	Hr	CPU	PG/S	C1ES	TR/S	NT/S		
2003	Jan	01		8	0	0.20	0.1	0.0		
		02		13	0	0.29	0.3	0.0		
		03		352	373	0.89	9.6	7.7		
		04		271	296	0.83	8.4	6.7		
		05		363	341	0.95	9.4	7.4		
		08		323	425	0.93	9.9	7.8		
		09		330	306	0.94	8.6	6.8		
		..		...	...	...	...	...		
		..		...	...	...	...	...		
		24		346	402	0.95	9.3	7.3		
		25		323	407	0.97	9.5	7.6		
		26		263	366	0.87	9.2	7.5		
		29		274	364	0.86	9.7	7.8		
		30		251	333	0.89	9.0	7.3		
		31		242	335	0.86	8.9	7.3		
		TOT		280	338	0.87	8.4	6.7		
							300054	239602		
		2003	Feb	01		265	359	0.88	8.7	6.9
				02		281	403	0.92	9.4	7.7
				05		267	425	0.93	9.6	7.8
				06		284	372	0.90	9.2	7.7
				07		266	339	0.88	9.0	7.4
				08		281	313	0.89	8.6	7.0
				..		...	...	...	...	...
				..		...	...	...	...	...
				21		277	339	0.86	9.6	7.5
				22		330	306	0.89	9.0	7.1
				23		245	349	0.86	9.6	7.8
26				327	442	0.96	10.0	8.4		
27				328	377	1.04	9.8	8.1		
28				267	305	0.95	8.7	7.2		
TOT				278	348	0.89	9.3	7.5		
							334088	268635		
2003 TOT				279	343	0.88	8.8	7.1		
Total				279	343	0.88	8.8	7.1		

Figure 239. Example of 'FCONXREP' output (file FCONXREP LISTING)

The following is the FCONXREP IMPORT file which corresponds to the FCONXREP LISTING file on the preceding page.

"	Year	Mon	Day	"	"	CPU	"	"	PG/S	"	"	C1ES	"	...
"	2003	Jan	01	"	"	8	"	"	0	"	"	0.20	"	...
"	2003	Jan	02	"	"	13	"	"	0	"	"	0.29	"	...
"	2003	Jan	03	"	"	352	"	"	373	"	"	0.89	"	...
"	2003	Jan	04	"	"	271	"	"	296	"	"	0.83	"	...
"	2003	Jan	05	"	"	363	"	"	341	"	"	0.95	"	...
"	2003	Jan	08	"	"	323	"	"	425	"	"	0.93	"	...
"	2003	Jan	09	"	"	330	"	"	306	"	"	0.94	"	...
"	...	...	..	"	"	...	"	"	...	"	"	...	"	...
"	...	...	..	"	"	...	"	"	...	"	"	...	"	...
"	2003	Jan	24	"	"	346	"	"	402	"	"	0.95	"	...
"	2003	Jan	25	"	"	323	"	"	407	"	"	0.97	"	...
"	2003	Jan	26	"	"	263	"	"	366	"	"	0.87	"	...
"	2003	Jan	29	"	"	274	"	"	364	"	"	0.86	"	...
"	2003	Jan	30	"	"	251	"	"	333	"	"	0.89	"	...
"	2003	Jan	31	"	"	242	"	"	335	"	"	0.86	"	...
"	2003	Feb	01	"	"	265	"	"	359	"	"	0.88	"	...
"	2003	Feb	02	"	"	281	"	"	403	"	"	0.92	"	...



"	2003	Feb	05	"	"	267	"	"	425	"	"	0.93	"	...
"	2003	Feb	06	"	"	284	"	"	372	"	"	0.90	"	...
"	2003	Feb	07	"	"	266	"	"	339	"	"	0.88	"	...
"	2003	Feb	08	"	"	281	"	"	313	"	"	0.89	"	...
....	...	...	...			...			...			....		...
....	...	...	...			...			...			....		...
"	2003	Feb	21	"	"	277	"	"	339	"	"	0.86	"	...
"	2003	Feb	22	"	"	330	"	"	306	"	"	0.89	"	...
"	2003	Feb	23	"	"	245	"	"	349	"	"	0.86	"	...
"	2003	Feb	26	"	"	327	"	"	442	"	"	0.96	"	...
"	2003	Feb	27	"	"	328	"	"	377	"	"	1.04	"	...
"	2003	Feb	28	"	"	267	"	"	305	"	"	0.95	"	...

Note that only the records with the most detailed data are shown (daily averages, in this case), while the higher averages are omitted.



# Appendix C. Printed Reports

While the format of the actual performance data on printed reports is identical to the format shown on display terminals, some special header lines and a table of contents will precede the performance information.

## Table of contents for printed reports

A table of contents is automatically created whenever multiple reports are printed with a PRINT ALL command. The table of contents will precede all of the actual performance data reports; its creation cannot be suppressed.

The following is an example for a table of contents:

FCXT0C Run 2003/04/03 16:30:57		Table of Contents		Page i	
From 2003/04/03 19:04:19		Performance Toolkit Reports in Order of Appearance			
To 2003/04/04 16:30:55				CHVM4	
For 77196 Secs 21:26:36		Performance reports for test system CHVM4		CPU 3090 SN 11619	
				z/VM V.n.n.n SLU 9401	
Rpt ID	Report Description	Page	From Date/Time	To Date/Time	
FCX100	General CPU Load and User Transactions	1	2003/04/04 13:29:55	- 2003/04/04 16:30:55	
FCX126	Logical Partition Activity	2	2003/04/04 13:29:55	- 2003/04/04 16:30:55	
FCX155	Count and Size of Processed Monitor Records	3	2003/04/04 13:30:55	- 2003/04/04 16:30:55	
FCX149	Monitor Settings: Initial and Changed	5	2003/04/03 19:04:19	- 2003/04/04 16:30:55	
FCX154	System Scheduler Settings, Initial and Changed	6	2003/04/03 19:04:19	- 2003/04/04 16:30:55	
FCX130	Overall Resource Usage per Transaction Statistic	7	2003/04/04 13:30:54	- 2003/04/04 16:30:55	
FCX111	Load and Performance for Cached Disks	9	2003/04/04 13:29:55	- 2003/04/04 16:30:55	
FCX177	Cache Extended Functions Performance	10	2003/04/04 13:29:55	- 2003/04/04 16:30:55	
FCX107	Channel Load and Channel Busy Distribution	11	2003/04/04 13:30:55	- 2003/04/04 16:30:55	
FCX176	Cached Control Units Overall Performance	12	2003/04/04 13:29:55	- 2003/04/04 16:30:55	
FCX108	General I/O Device Load and Performance	13	2003/04/04 13:29:55	- 2003/04/04 16:30:55	
FCX109	Load and Performance of CP Owned Disks	14	2003/04/04 13:29:55	- 2003/04/04 16:30:55	
....	....				
....	....				

Figure 240. Layout of print file table of contents

It lists the reports in the order in which they are appended, which is equal to the order in which they were specified in the FCONX REPORTS file.

### Field description:

General header lines: See “Layout of print file header:” on page 979 for their description.

### Table of contents:

#### Rpt ID

The internal report identification, equal to the report ID on the header of the report.

#### Report Description

The report description, equal to the description on the header part of the report, except that it will be truncated to a maximum length of 48 characters.

#### Page

The number of the first page of the report.

#### From Date/Time

The date and time of the first monitor sample which contributed to the data included in the report.

#### To Date/Time

The date and time of the last monitor sample which contributed to the data included in the report.

## Layout of print file header:

All print output is created with a common header which precedes the actual performance data. You can define the contents of two of the header fields; the remainder is generated automatically when building the print output. The following is an example for a print file header:

```

FCXnnn Run yyyy/mm/dd hh:mm:ss      reportname      Page      n
                                report description
From yyyy/mm/dd hh:mm:ss            systemid
To   yyyy/mm/dd hh:mm:ss            CPU nnnn      SN
nnnnn
For   nnnn Secs hh:mm:ss            system description      z/VM V.n.n.n SLU nnnn
-----
... .. Performance data ...
... ..

```

Figure 241. Layout of print file header

*Field description:*

**FCXnnn Run yyyy/mm/dd hh:mm:ss**

The first item is the internal report identification, and the Run information provides the date and time when the performance analysis run was made.

**From yyyy/mm/dd hh:mm:ss**

This is the date and time of the first monitor sample which contributed to the data included in the report.

**To yyyy/mm/dd hh:mm:ss**

This is the date and time of the last monitor sample which contributed to the data included in the report.

**For nnnn Secs hh:mm:ss**

This is the length of the report period, in seconds and in hours/minutes/seconds.

**reportname**

This is the short report name, equal to the command used for selecting it.

**report description**

This is a more detailed description of the data contained in the report.

**system description**

This is a field for user data, intended to describe the system where the analyzed monitor data come from. It can be set using the command 'FC SETTINGS SYSTEM ...'; see [“FCONTROL SETTINGS” on page 83](#) for details.

**systemid**

This is a short system identification. It is intended to contain the system identification or the RSCS nodeid of the system where the analyzed monitor data come from. The field can be set using the command 'FC SETTINGS SYSTEMID ...'; see [“FCONTROL SETTINGS” on page 83](#) for details. The system-ID of the VM system which provided the monitor data will be inserted if the user does not specify something else.

**CPU nnnn SN nnnnn**

These fields provide the CPU model and serial numbers of the VM system whose data were processed.

**z/VM V.n.n.n SLU nnnn**

These fields describe the CP software level of the VM system which provided the data.

---

## Appendix D. HISTLOG and HISTSUM File Records

The record layout of HISTLOG and HISTSUM files is described by HISTSECT DSECT. It is defined in COPY files in FCXGPI MACLIB, which is available on the Performance Toolkit production minidisk.

### Usage Notes:

1. Short floating-point format (E) is used storing the values of all performance variables. A value of X'FFFFFFFF' indicates that the corresponding value was not available, or could not be calculated because the calculation would have resulted in a division by zero. *Always test for a string of X'FFFFFFFF' before interpreting a value as valid data.*
2. When using these variable names in commands (FCONTROL LIMIT, for example), if you want to use a variable that begins with the pound sign (#), which would otherwise be interpreted as the logical line-end character, you must precede it with a double quotation mark ("), which is the logical escape character. For example, for the variable #CPU, you must specify "#CPU.

Note that these characters *cannot* be used when accessing Performance Toolkit remotely. For more information, see the "Logical Line End and Escape Characters" section of the "General System Operation" chapter in the [z/VM: Performance Toolkit Guide](#).



## Appendix E. Extended Trend File Records

The following trend records can be created by Performance Toolkit, subject to specifications in the FCONX TRENDREC file and the availability of the prerequisite monitor data.

<Record Number> Toolkit      Contents		Selection
X'FC00'	System configuration data	--
X'FC01'	General system load data	--
X'FC02'	Processor load data	--
X'FC03'	LPAR load data	--
X'FC04'	Minidisk / fulltrack cache data	--
X'FC05'	CP services activity data	--
X'FC06'	Channel busy (HF sampling)	CHANNEL
X'FC07'	Channel measurment facility data	CHANNEL
X'FC08'	Extended channel measurement facility data	CHANNEL
X'FC09'	System spin lock data	--
X'FC0A'	Privileged instructions counters	PRIVOPS
X'FC0B'	DIAGNOSE instructions counters	DIAGS
X'FC0C'	ISFC End Point configuration	--
X'FC0D'	ISFC Logical Link configuration	--
X'FC0E'	ISFC End Point Activity	--
X'FC0F'	ISFC Logical Link Activty	--
X'FC3A'	Overall user transaction & response time data	--
X'FC3C'	Shared segments data	NSS
X'FC3E'	Shared data spaces data	DSPACES
X'FC41'	User resource usage & wait states data	USER
X'FC42'	User class resource usage & wait states data	--
X'FC43'	Overall user resource usage & wait states data	--
X'FC44'	User transaction & response time data	USER
X'FC45'	User class transaction & response time data	--
X'FC46'	Overall user transaction & response time data	--
X'FC47'	CPU Pool data	--
X'FC51'	I/O processor activity data	--
X'FC54'	VSWITCH global networking object data	VSWITCH
X'FC55'	VSWITCH activity data	VSWITCH
X'FC56'	VNIC activity data	VNIC
X'FC57'	PCI Function configuration data	PCIFUNC
X'FC58'	PCI Function activity data	PCIFUNC
X'FC61'	General DASD data	DASD
X'FC65'	DASD cache data	DASD
X'FC68'	DASD CP owned data (system areas)	DASD
X'FC6D'	QDIO device data	QDIO
X'FC6F'	SCSI device data	SCSI
X'FC71'	DASD seeks data	SEEKS
X'FCA2'	SFS and BFS server data	SFS
X'FCA4'	Multitasking users data	MTUSER
X'FCA6'	TCP/IP server data	TCPIP
X'FCA7'	TCP/IP links data	TCPIP
X'FCA8'	Reusable server kernel summary data	RSK
X'FCA9'	Linux internal summary data	LINUX
X'FCB1'	SSI configuration	--
X'FCB2'	SSI State Change Synchronization Activity and State/Mode information	--

In the Selection field, dashes indicates that denotes those records dealing with general system performance data. They are always created (provided the prerequisite monitor domains were enabled) and need not be explicitly specified. The remaining records will be created only if the corresponding selection fields are specified on the RECORDS statement in the FCONX TRENDREC and FCONX SUMREC control files. See the "Creating Extended Summary and Trend Files" section in the "Operation in Performance Monitor Mode" chapter in the [z/VM: Performance Toolkit Guide](#) for more information.

The layout of these trend file records is described by a number of assembler DSECTs:

- TRNDRHDR DSECT for the common record header format
- HISxxSCT DSECT for the actual data area, where xx corresponds to the second half of the hexadecimal record number; the first byte of Performance Toolkit trend record numbers is always X'FC'.

These DSECTs are defined in COPY files in FCXGPI MACLIB, which is available on the Performance Toolkit production minidisk.

## Format of Physical Records

Trend records are accumulated and written as physical 4kB blocks, where the first record in the first block of each set of trend records for a new period is record X'FC00'. The remaining records are directly appended, i.e. you find the start of the following logical record by adding the length of the first (or current) record to its start address. Records are not spanned, i.e. where a record does not fit into the remaining space of a 4k block it will be written as the first record in a new block. The remainder of the block is cleared to X'00', i.e. you will have to read the next block and continue processing there if the offset to the start of the next record is 4096 or if you find a 'record' with a length of 0, and a record number of X'0000'.

First Block:

```

-----
| Rec.X'FC00'          | Rec. X'FCxx'          | ...          |
-----
| ...                  |                        |              |
| ...                  |                        |              |
|                        | Rec. X'FCxx'          |              |
-----

```

...  
Last Block:

```

-----
| Rec.X'FCxx'          | Rec. X'FCxx' ...      |
-----
| ...                  |                        |
| ...                  |                        |
|                        |                        |
-----

```

Figure 242. Trend file layout



---

## Appendix F. Extended Summary File Records

The summary file is intended for use as input to some other program, such as SLR. It contains only data for one single 'summary' period, and its records are written as separate variable length records to facilitate processing by other programs.

Similar to trend files, summary file creation is controlled by means of the 'FC MONCOLL RESET' command which allows you to define 'summary' periods that are independent of 'trend' and other periods. See the “Creating Extended Summary and Trend Files” and “Resetting Performance Data/Defining Periods” sections in the “Using Performance Toolkit in Performance Monitor Mode” chapter of the [\*z/VM: Performance Toolkit Guide\*](#) for the corresponding guidelines.

---

### Summary Record Headers

To allow easier processing, a different record header is used that contains the system identification and CPUID of the originating VM system.

SUMRHDR DSECT describes the layout of summary record headers. It is defined in COPY files in FCXGPI MACLIB, which is available on the Performance Toolkit production minidisk.

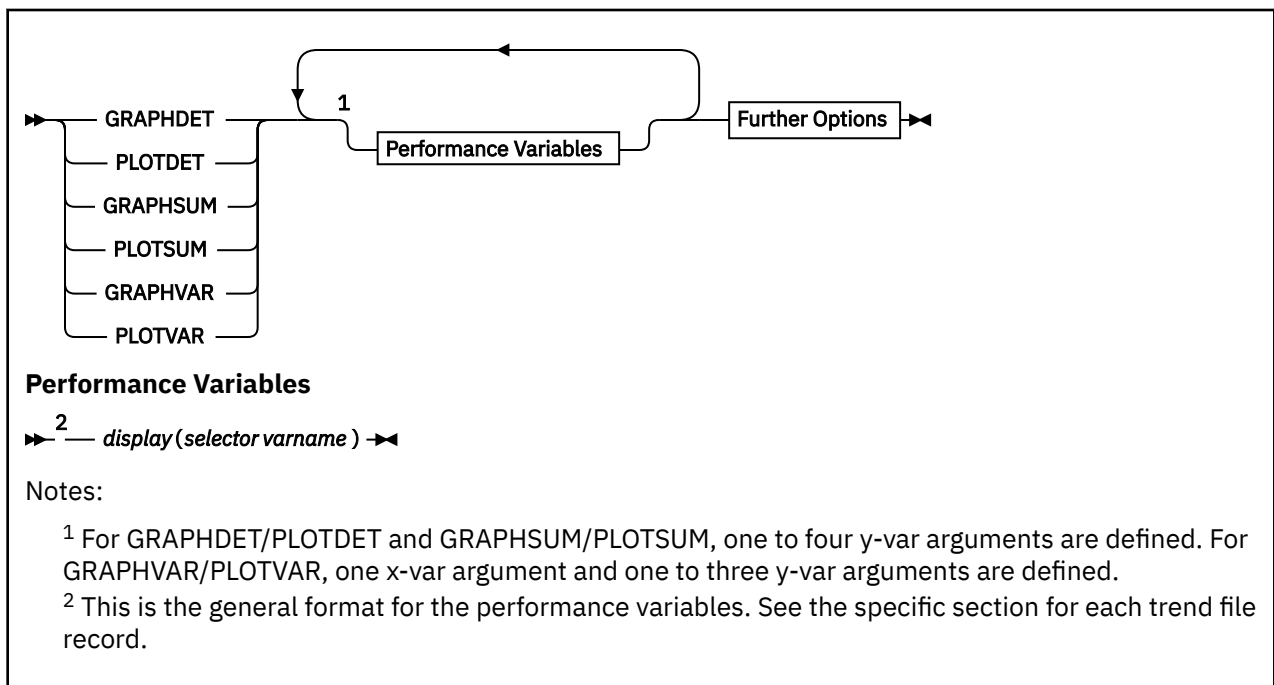
The actual data part after the header is identical in layout to one of the corresponding trend records. Please note, however, that due to the longer record header you will have to add 16 (X'10') to the field offsets in order to obtain the correct displacement of a field in a summary record.



## Appendix G. Trend File Variables Selection

This appendix lists the performance variables that can be selected for the creation of graphics from trend records. All of the variables listed are fields from a performance display. Since Performance Toolkit trend records only include data required for building the performance displays that could have been viewed based on the original monitor data, this covers most of the meaningful data that can be extracted from Performance Toolkit trend files.

The selectable variable names are listed in separate sections, under the heading of the trend record number on which they are based. Note that all of the syntax diagrams in this appendix are just fragments, to be used for specifying the composite trend variable selection arguments in lieu of the simple single strings documented for the original PLOTxxx and GRAPHxxx commands (see “GRAPHDET / PLOTDET” on page 182, “GRAPHSUM / PLOTSUM” on page 185, and “GRAPHVAR / PLOTVAR” on page 189). The complete graphics selection command would therefore look as follows:



where

### **display(**

is the name of the performance display that shows the variable. The following left parenthesis *must* be directly appended, without any intervening blanks.

### **selector**

is a further selector that is required for all record types where multiple similar records can be created, for instance a

- processor number
- channel path ID
- user identification
- device number

### **varname**

is the variable name. Wherever possible the variable name is chosen to be identical to the string shown in the lowest one of the header lines describing the variable, similar to the logic used for defining SORT arguments.

Where this logic could not be adhered to because the string is not unique on the display, new names have been defined and included with the affected variable description.

The selection format appropriate for specific trend file records, i.e. the specific 'selector' format to be used, is described at the beginning of the actual variable description section for each trend record.

## Performance Variables of Trend Records X'FC01'

Trend record X'FC01' contains overall system utilization and performance data.

The selection format for record 01 variables is:

► SYSTEM( *varname* ) ◄

where:

### ***varname***

is the name of the selected variable, as shown under the heading 'Variable Name' in the table below.

The variables shown in the following table have been pre-defined. Where possible, the same names have been chosen as for the equivalent variables defined for HISTLOG and HISTSUM files, i.e. they usually *do not* correspond directly to part of the header line of a performance display. Their description also includes the formula used to calculate the values from the original trend record data fields.

Not all of the data needed for calculating the old HISTLOG variables ended up in trend record X'FC01'. In order to simplify the migration to using the new trend file for capacity planning and trend analysis, the program has been set up to accept some of these old variables also when the prerequisite data comes from some other trend record. This is true at least in cases where the selection would have become more complex, otherwise, because some additional argument would have had to be entered (processor number for example).

Display	Variable Name	Variable Description and Derivation
<b>General Sample Collect Information</b>		
SYSTEM	EL_TIME	Total elapsed time covered by trend record ( TRNDRELT )
	SAMPLES	Number of CP monitor samples covered by trend record (to be used for dividing totals in order to obtain average values) ( TRNDRSMP )

### **CPU Load Fields**

All values are based on processor trend records X'FC02'. Values are calculated by processor, and the total of the values for all processors is returned. CPU load percentages can thus reach a value of 100 times the number of processors.

SYSTEM	#CPU	Number of processors ( TRNDRSMP / TRNDRSMP ) (adding one per HIS02SCT record)
	CPU	Total CPU load, in percent ( (H02USTIM+H02SYTIM) * 100 / TRNDRELT )
	%US	User CPU utilization, in percent ( H02USTIM * 100 / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	%CP	CP CPU utilization, in percent ( (H02USTIM+H02SYTIM-H02EMTIM) * 100 / TRNDRELT )
	%EM	Emulation CPU utilization, in percent ( H02EMTIM * 100 / TRNDRELT )
	%SY	System CPU utilization, in percent ( H02SYTIM * 100 / TRNDRELT )
	%SPIN	CPU spin utilization, in percent ( H02SPIN * 100 / TRNDRELT )
	%WT	CPU voluntary wait, in percent ( H02WAIT * 100 / TRNDRELT )
	%LOGLD	Logical CPU load, in percent (based on CPU active and voluntary wait time, but excluding any suspended time when running 2nd level or in an LPAR) ( (H02USTIM+H02SYTIM) * 100 / (H02USTIM+H02SYTIM+H02WAIT) )
	%VEC	Total vector utilization, in percent ( (H02VFTIM+H02VOTIM) * 100 / TRNDRELT )
	%VEM	Vector emulation load, in percent ( H02VFTIM * 100 / TRNDRELT )
	VECLD	Vector load rate, per second ( H02VFLOD / TRNDRELT )
	PRIV	Privileged instruction simulation rate per sec. ( H02PRVCT / TRNDRELT )
	DIAG	Diagnose instruction execution rate per sec. ( H02DIACT / TRNDRELT )
	SIGP	SIGP Ext Call receipt rate, per second ( H02SIGCT / TRNDRSMP )
	IO/S	SSCH and RSCH rate per second ( (H02SSCH+H02RSCH) / TRNDRSMP )
I	%EMPT	Percent of time DSVBK found empty ( H02EMPTY * 100 / H02HFSMP )

Display	Variable Name	Variable Description and Derivation
	#USRS	Mean VMDBKs in DSVBK while not empty ( H02VMDBK / (H02HFSMP-H02EMPTY) )
	#MAST	Mean VMDBKs in master processor DSVBK ( H02MASTR / H02HFSMP )
	VMSTL	VMDBKs stolen from processor, per second ( H02VMSTO / TRNDRELT )
	TOMAST	VMDBKs moved to master proc. per second ( H02VMMAS / TRNDRELT )
	AVEMP	Seconds between available list empty ( TRNDRELT / (sum of H02AVEMP) )

Display	Variable Name	Variable Description and Derivation
---------	---------------	-------------------------------------

### General System Activity

Where values are based on processor trend records X'FC02', these values are calculated by processor, and the total of the values for all processors is returned.

SYSTEM	DSPTCH	New users dispatched rate, per second ( H02LNGDS / TRNDRELT )
	SIE	SIE instruction rate, per second ( H02SIE / TRNDRELT )
	SIEIC	SIE intercept rate, per second ( H02SIEIC / TRNDRELT )
	PGFLT	Host page fault rate, per second ( H02PGFLT / TRNDRELT )
	FPSIM	Fastpath instruction simulation rate/s ( H02FPSIM / TRNDRELT )
	SPRDS	Spool I/O rate: reads per second ( H02SPRDS / TRNDRELT )
	SPWRT	Spool I/O rate: writes per second ( H02SPWRT / TRNDRELT )
	CSCH	Clear subchannel rate, per second ( H02CSCH / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	HSCH	Halt subchannel rate, per second ( H02HSCH / TRNDRELT )
	EXTINT	External interrupts rate, per second ( H02EXTIT / TRNDRELT )
	SOLINT	Solicited interrupts rate, per second ( H02SOLIT / TRNDRELT )
	UNSOLI	Unsolicited interrupts rate, per second ( H02USLIT / TRNDRELT )
	LOGN	Average users logged on ( H01UCLOG )
	ACT	Average users active in monitor interval ( H01UCACT )
	VIO/S	Virtual I/O rate, per second ( H01UCVIO )

#### Central Storage Variables

SYSTEM	PPAG	Number of pageable frames in the DPA ( (H01PGBL-H01NPAG+H01PGBL2) / TRNDRSMP )
	PPAG<2G	Number of pageable frames below the 2GB line ( (H01PGBL-H01NPAG) / TRNDRSMP )
	PPAG>2G	Number of pageable frames above the 2GB line ( H01PGBL2 / TRNDRSMP )
	NONPAG	Number of non-pageable DPA pages below 2GB line ( H01NPAG / TRNDRSMP )
	%ST	Percent of DPA pageable frames in use ( (H01SCWS0+H01SCWS1+H01SCWS2+H01SCWS3) * 100 * TRNDRSMP / (H01PGBL-H01NPAG+H01PGBL2) )
	SHAR	Number of resident shared pages ( H01RSHR / TRNDRSMP )
	FPGS	Number of FREE storage pages ( H01XFRE / TRNDRSMP )
	%FR	Percentage of FREE subpool storage in use ( H01FRUS / TRNDRSMP )



Display	Variable Name	Variable Description and Derivation
	V=R	V=R storage size in Bytes ( H01VRSIZ )
	VRFREE	V=R FREE storage in use (doublewords) ( H01VFRUS / TRNDRSMP )
	FEX/S	Frames extended rate, per second ( H02FRTAK / TRNDRELT ) (Sum of rates for all processors)
	FRMISS	Number of frames missing from free storage reserved frames list ( H01PGXTE / TRNDRSMP )
	FRSUB	Free subpool space available (pages) (converted from doublewords) ( H01SPAV / TRNDRSMP / 512 )
	FRSAV	Number of save area pages ( H01PGSAV / TRNDRSMP )
	RESRVD	Number of reserved pages ( H01PGRSV / TRNDRSMP )
	LOCKED	Number of locked pages ( (H01PGLOK+H01LOCKG) / TRNDRSMP )
	LOCK<2G	Number pages locked below the 2GB line ( H01PGLOK / TRNDRSMP )
	LOCK>2G	Number pages locked above the 2GB line ( H01LOCKG / TRNDRSMP )
	SXSTLOCK	Total SXS alias pages locked ( H01SXLPT / TRNDRSMP )
	SYSLLOCK	SXS alias pages locked by CP LOCK command ( H01SXLPL / TRNDRSMP )
	DI98LOCK	Pages locked for Diagnose X'98' in 31-bit mode ( H01DIA98 / TRNDRSMP )
	LOWTH<2G	Low threshold for <2GB available list ( H01ALLTH / TRNDRELT )
	LOWTH>2G	Low threshold for >2GB available list ( H01ALLTG / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	HITH<2G	High threshold for <2GB available list ( H01ALHTH / TRNDRELT )
	HITH>2G	High threshold for >2GB available list ( H01ALHTG / TRNDRELT )
	AVAIL	Number of pages on the available list ( H01AVAL / TRNDRSMP )
	AVAIL<2G	Number of pages on the <2GB available list ( H41AVALP / TRNDRSMP )
	AVAIL>2G	Number of pages on the >2GB available list ( H41AVALG / TRNDRSMP )
	MS-LIFE	Estimated central storage page life, based on Little's law, in seconds ( (H01PGBL-H01NPAG) / TRNDRSMP * TRNDRELT / (sum of H02PGWRT+H02PGOUT) ) The MS-LIFE value can be calculated only if the sum of all H02PGWRT and H02PGOUT values is nonzero.
	PRD/S	Page read rate, per second ( H02PGRDS / TRNDRELT ) (Sum of rates for all processors)
	PWR/S	Page write rate, per second ( H02PGWRT / TRNDRELT ) (Sum of rates for all processors)
	PG/S	Total page rate, per second ( (H02PGRDS+H02PGWRT) / TRNDRELT ) (Sum of rates for all processors)
	NSS/S	Page read rate/s for shared segments ( H02SHRRD / TRNDRELT ) (Sum of rates for all processors)
	GUEST/S	Single page read rate/s for guests ( H01SRDUS / TRNDRELT )
	SYST/S	Single page read rate/s for system ( H01SRDSY / TRNDRELT )
	SPAG/S	Total single page read rate, per second ( (H01SRDUS+H01SRDSY) / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	ALO/S	Pages allocation rate, per second ( H02PGALO / TRNDRELT ) (Sum of rates for all processors)
	#TW	Nr. of tasks waiting for a page, per second ( H02TSKWT / TRNDRELT ) (Sum of rates for all processors)
	FRWT	Number of tasks waiting for a frame ( H01DFRW / TRNDRSMP )
	SCAN1	Demand scan 1 successful rate, per second ( H02SCAN1 / TRNDRELT ) (Sum of rates for all processors)
	SCAN2	Demand scan 2 successful rate, per second ( H02SCAN2 / TRNDRELT ) (Sum of rates for all processors)
	E-SCAN	Emergency scan successful rate, per second ( H02SCANE / TRNDRELT ) (Sum of rates for all processors)
	SCFAIL	Not enough frames after scan, per second ( H02SCAFA / TRNDRELT ) (Sum of rates for all processors)
	ESCN	Percent of demand scans which did not complete with scan 1 ( (H02SCAN2+H02SCANE) * 100 / (H02SCAN1+H02SCAN2+H02SCANE) ) (Average value for all processors)
	BKRD	Average page read block size ( H01RDBSZ * H01RDBIO / H01RDBIO ) The H01RDBSZ field is valid only if the H01RDBIO field is nonzero. A dummy multiplication and division is used to force an 'invalid' value when H01RDBIO is zero.
	BKMG	Average page migrate block size ( H01MGBSZ * H01MGBIO / H01MGBIO ) The H01MGBSZ field is valid only if the H01MGBIO field is nonzero. A dummy multiplication and division is used to force an 'invalid' value when H01MGBIO is zero.

Display	Variable Name	Variable Description and Derivation
	BKST	Average page steal block size $(H01WRBSZ * H01WRBIO / H01WRBIO)$ The H01WRBSZ field is valid only if the H01WRBIO field is nonzero. A dummy multiplication and division is used to force an 'invalid' value when H01WRBIO is zero.

### Expanded Storage Variables

SYSTEM	XSTAV	Total number of expanded storage blocks in CP partition $(H01CPPAR / TRNDRSMP)$
	XSTPGBK	Total number of expanded storage paging blocks $((H01CPPAR - H01MDCNE) / TRNDRSMP)$
	%XS	Percent of expanded storage in use $((H01CPPAR - H01CPXAV) * 100 / H01CPPAR)$
	XAL/S	Expanded storage blocks allocation rate/s $(H01XSGET / TRNDRELT)$
	XPG/S	Expanded storage page rate (PGIN + PGOUT) $((H02FPGIN + H02PGINS + H02PGOUT) / TRNDRELT)$ (Sum of rates for all processors)
	XDEAL/S	Expanded storage blocks deallocation rate/s $(H01XSREL / TRNDRELT)$
	MIGVIS	Shared system migration visits rate/s $(H01XSHRV / TRNDRELT)$
	PGIN/S	PGIN rate, per second $((H02FPGIN + H02PGINS) / TRNDRELT)$ (Sum of rates for all processors)
	%FP_PGIN	Percent of PGINs done using fast-path logic $(H02FPGIN * 100 / (H02FPGIN + H02PGINS))$ (average for all processors)
	PGOUT/S	PGOUT rate, per second $(H02PGOUT / TRNDRELT)$ (Sum of rates for all processors)
	MIGR/S	Page migrate rate/s, exp. storage to DASD $((H01MBKDM + H01MBKND + H01MBKSH + H01MBKSY) / TRNDRELT)$

Display	Variable Name	Variable Description and Derivation
	XAGE	Average expanded storage page age (sec) ( H01AVAG / TRNDRSMP )
	MAGE	Average XSTORE page age at migration (sec) ( H01MIGAG * H01MBKDM / H01MBKDM ) The H01MIGAG field is valid only if the H01MBKDM field is nonzero. A dummy multiplication and division is used to force an 'invalid' value when H01MBKDM is zero.
	XSTLIFE	Estimated XSTORE page life, based on Little's law, in seconds ( (H01CPPAR-H01MDCNE) / TRNDRSMP * TRNDRELT / (sum of H02PGOUT) ) The XSTLIFE value can be calculated only if the sum of all H02PGOUT values is nonzero.

### Minidisk Cache Data

Please see the performance variables defined for trend records X'FC04' ([“Performance Variables of Trend Records X'FC04”](#) on page 1022, full-track cache data) for all variables related to the minidisk cache.

### Auxiliary Storage Variables (AUXLOG)

SYSTEM	PGSLOTS	Total page slots available on paging areas ( H01AUPSL )
	%PGSL	Paging area utilization, in percent ( H01AUPUS * 100 / H01AUPSL )
	SPSLOTS	Total spool slots available on spooling areas ( H01AUSSL )
	%SPSL	Spooling area utilization, in percent ( H01AUSUS * 100 / H01AUSSL )
	DUSLOTS	Total spool slots on dedicated dump areas ( H01AUDSL )
	%DUSL	Dump spooling area utilization, in percent ( H01AUDUS * 100 / H01AUDSL )
	SPCREA	Total spool files created ( H01AUFCE )
	SPCRE/S	Spool files created per second ( H01AUFCE / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	SPPURG	Total spool files purged ( H01AUFPU )
	SPPUR/S	Spool files purged per second ( H01AUFPU / TRNDRELT )
	PGMLOAD	Paging MLOAD (average page I/O response time, per page) in msec. ( H01AUAMP / 1000 )
	SPMLOAD	Spooling MLOAD (average spool I/O response time, per spool block) in msec. ( H01AUAMS / 1000 )

#### Scheduler Queues (SCHEDLOG)

SYSTEM	INQUE	Total users in dispatch and eligible lists ( H01SCDSP + H01SCEL1 )
	Q0	Class 0 users in dispatch list ( H01SCDSP - H01SCDS1 )
	Q1	Class 1 users in dispatch list ( H01SCDS1 - H01SCDS2 )
	QX	Class 0, 2 and 3 users in dispatch list ( H01SCDSP-H01SCDS1+H01SCDS2 )
	Q2	Class 2 users in dispatch list ( H01SCDS2 - H01SCDS3 )
	Q3	Class 3 users in dispatch list ( H01SCDS3 )
	Q0L	Loading class 0 users in dispatch list ( H01SCLDG - H01SCLD1 )
	Q1L	Loading class 1 users in dispatch list ( H01SCLD1 - H01SCLD2 )
	QXL	Loading class 0, 2 and 3 users in dispatch list ( H01SCLDG-H01SCLD1+H01SCLD2 )
	Q2L	Loading class 2 users in dispatch list ( H01SCLD2 - H01SCLD3 )
	Q3L	Loading class 3 users in dispatch list ( H01SCLD3 )

Display	Variable Name	Variable Description and Derivation
	LIMLST	Users in limit list ( H01SCLLS )
	E1	Class 1 users in eligible list ( H01SCEL1 - H01SCEL2 )
	EX	Class 2 and 3 users in eligible list ( H01SCEL2 )
	E2	Class 2 users in eligible list ( H01SCEL2 - H01SCEL3 )
	E3	Class 3 users in eligible list ( H01SCEL3 )
	E1L	Loading class 1 users in eligible list ( H01SCLE1 - H01SCLE2 )
	EXL	Loading class 2 and 3 users in elig. list ( H01SCLE2 )
	E2L	Loading class 2 users in eligible list ( H01SCLE2 - H01SCLE3 )
	E3L	Loading class 3 users in eligible list ( H01SCLE3 )
	C1ES	Class 1 elapsed time slice (sec.) ( H01SCETS / 1000 )
	ABSSHR	Sum of absolute shares for all in-queue users ( H01SCABS * 100 / 65534 )
	RELSHR	Sum of relative shares for all in-queue users ( H01SCREL )
	CONSID	Total storage considered (pages) ( H01SCPGS )
	Q0WSS	Total WSS of all Q0 users (pages) ( H01SCWS0 )
	Q1WSS	Total WSS of all Q1 users (pages) ( H01SCWS1 )
	Q2WSS	Total WSS of all Q2 users (pages) ( H01SCWS2 )

Display	Variable Name	Variable Description and Derivation
	Q3WSS	Total WSS of all Q3 users (pages) ( H01SCWS3 )
<b>Scheduler Activity (SCHDATA, Timer Pop Data)</b>		
SYSTEM	SC_PG/S	Page read rate per second ( H01SDPRT / H01SDSMP )
	SC_MS/P	Milliseconds per page read ( H01SDTRD / H01SDSMP / 10 )
	SC_PSV	Page steal value ( H01SDSTL / H01SDSMP )
	SC_SWGT	Storage resource weight ( H01SDSTW / H01SDSMP )
	SC_PWGT	Paging resource weight ( H01SDPGW / H01SDSMP )
	SC_C1P/S	Average class 1 user page rate/s ( H01SDEPR / H01SDSMP )
	SC_C1WSS	Average class 1 user WSS ( H01SDEWS / H01SDSMP )
	SC_ALP/S	Average page rate/s for all users ( H01SDAPR / H01SDSMP )
	SC_ALWSS	Average WSS for all users ( H01SDAWS / H01SDSMP )
	CL1EXP	Class 1 expansion factor ( H01SDE1X / H01SDSMP )
	CL2EXP	Class 2 expansion factor ( H01SDE2X / H01SDSMP )
	CL3EXP	Class 3 expansion factor ( H01SDE3X / H01SDSMP )
<b>Response Times / Transactions (RESPALL, RESPUP)</b>		
SYSTEM	TR_UP_T	Trivial UP users' response time (s) ( H01RSUTT / H01RSUTC )



Display	Variable Name	Variable Description and Derivation
	TR_MP_T	Trivial MP users' response time (s) ( H01RSMTT / H01RSMTC )
	TR-T	Average trivial response time (s) ( (H01RSUTT+H01RSMTT) / (H01RSUTC+H01RSMTC) )
	NT_UP_T	Non-trivial UP users' response time (s) ( H01RSUNT / H01RSUNC )
	NT_MP_T	Non-trivial MP users' response time (s) ( H01RSMNT / H01RSMNC )
	NT-T	Average non-trivial response time (s) ( (H01RSUNT+H01RSMNT) / (H01RSUNC+H01RSMNC) )
	UP_RESP	Average UP response time (s) ( (H01RSUTT+H01RSUNT) / (H01RSUTC+H01RSUNC) )
	QD_T	Quick-dispatch response time (s) ( H01RSQDT / H01RSQDC )
	MEANRESP	Mean response time (s) ( (H01RSUTT+H01RSUNT+H01RSMTT +H01RSMNT+H01RSQDT) / (H01RSUTC+H01RSUNC+H01RSMTC +H01RSMNC+H01RSQDC) )
	TR_UP/S	Trivial UP transaction rate per second ( H01RSUTC / TRNDRELT )
	TR_MP/S	Trivial MP transaction rate per second ( H01RSMTC / TRNDRELT )
	TR/S	Average trivial transaction rate/s ( (H01RSUTC+H01RSMTC) / TRNDRELT )
	NT_UP/S	Non-trivial UP transaction rate per second ( H01RSUNC / TRNDRELT )
	NT_MP/S	Non-trivial MP transaction rate per second ( H01RSMNC / TRNDRELT )
	NT/S	Average non-trivial transaction rate/s ( (H01RSUNC+H01RSMNC) / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	UP_TR/S	Average UP transaction rate per second ( (H01RSUTC+H01RSUNC) / TRNDRELT )
	QD_TR/S	Quick-dispatch transaction rate per second ( H01RSQDC / TRNDRELT )
	TRANS/S	Total transaction rate per second ( (H01RSUTC+H01RSUNC+H01RSMTC +H01RSMNC+H01RSQDC) / TRNDRELT )
	TR-Q	Average trivial transactions active ( (H01RSUTT+H01RSMTT) / TRNDRELT )
	NT-Q	Average non-trivial transactions active ( (H01RSUNT+H01RSMNT) / TRNDRELT )
	%TR_UP	Percent trivial UP transactions, of all trans. ( H01RSUTC * 100 / (H01RSUTC+H01RSMTC+H01RSUNC +H01RSMNC+H01RSQDC) )
	%TR_MP	Percent trivial MP transactions, of all trans. ( H01RSMTC * 100 / (H01RSUTC+H01RSMTC+H01RSUNC +H01RSMNC+H01RSQDC) )
	%NT_UP	Percent non-trivial UP transact., of all trans. ( H01RSUNC * 100 / (H01RSUTC+H01RSMTC+H01RSUNC +H01RSMNC+H01RSQDC) )
	%NT_MP	Percent non-trivial MT transact., of all trans. ( H01RSMNC * 100 / (H01RSUTC+H01RSMTC+H01RSUNC +H01RSMNC+H01RSQDC) )
	%Q_DSP	Percent quick-dispatch transact., of all trans. ( H01RSQDC * 100 / (H01RSUTC+H01RSMTC+H01RSUNC +H01RSMNC+H01RSQDC) )
	#TR_UP	Number of trivial UP transactions ( H01RSUTC )
	#NT_UP	Number of non-trivial UP transactions ( H01RSUNC )
	#TR_MP	Number of trivial MP transactions ( H01RSMTC )

Display	Variable Name	Variable Description and Derivation
	#NT_MP	Number of non-trivial MP transactions ( H01RSMNC )
	#TRANS	Total number of transactions ( H01RSUTC+H01RSMTC+H01RSUNC +H01RSMNC+H01RSQDC )
	#ALLUP	Total number of UP transactions ( H01RSUTC + H01RSUNC )
	%TRUP/UP	% of UP transactions which were trivial ( H01RSUTC * 100 / (H01RSUTC+H01RSUNC) )
	%NTUP/UP	% of UP transactions which were non-trivial ( H01RSUNC * 100 / (H01RSUTC+H01RSUNC) )

#### User Wait State Averages (USTLOG)

SYSTEM	%ACT	Percent of time an average user was found to be active (not dormant) ( (H01STHFC+H01STDRS-H01STDRM) * 100 / H01STHFC )
	%RUN	Percent of active time users were found running ( H01STRUN * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%CPU	Percent of active time users were in CPU wait ( H01STCPW * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%LDG	Percent of active time users were loading ( H01STLOD * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%PGW	Percent of active time users were in page wait ( H01STPGW * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%IOW	Percent of active time users were in I/O wait ( H01STIOW * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%SIM	% of active time users were in inst. simul. wait ( H01STSIM * 100 / (H01STHFC+H01STDRS-H01STDRM) )

Display	Variable Name	Variable Description and Derivation
	%TIW	Percent of active time users were in test-idle wait ( H01STTID * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%CFW	Percent of active time users were in CF wait ( H01STCFW * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%TI	Percent of active time users were in SVM & TI wait ( H01STTIS * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%EL	Percent of active time users were in SVM wait and in the eligible list ( H01STELS * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%DM	Percent of active time users were in SVM wait and dormant ( H01STDRS * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%IOA	% of active time users were in 'I/O active' wait ( H01STIOA * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%PGA	% of active time users were in 'page active' wait ( H01STPGA * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%LIM	Percent of active time users were in limit list wait ( H01STLMT * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%OTH	% of active time users were in another wait state ( H01STOTH * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%Q0	Percent of active time users were in queue 0 ( H01STIQ0 * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%Q1	Percent of active time users were in queue 1 ( H01STIQ1 * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%Q2	Percent of active time users were in queue 2 ( H01STIQ2 * 100 / (H01STHFC+H01STDRS-H01STDRM) )

Display	Variable Name	Variable Description and Derivation
	%Q3	Percent of active time users were in queue 3 ( H01STIQ3 * 100 / (H01STHFC+H01STDRS-H01STDRM) )
	%E0-3	Percent of active time users were in elig. list ( H01STELI * 100 / (H01STHFC+H01STDRS-H01STDRM) )

#### **DASD Load and Performance (DASDLOG)**

SYSTEM	LINKS	Average minidisk links per DASD ( H01DSLNK / H01DSCNT )
	PATHS	Average number of channel paths per DASD ( H01DSPTH / H01DSCNT )
	DASD/S	Total SSCH & RSCH rate per second to DASD ( H01DSSSC / TRNDRELT )
	AVOID/S	Total I/Os avoided rate per second ( H01DSNIO / TRNDRELT )
	PEND	Average 'function pending' time, in msec. ( H01DSPND * 128 / 1000 / H01DSSSC )
	DISC	Average 'disconnected' time, in msec. ( H01DSDSC * 128 / 1000 / H01DSSSC )
	CONN	Average 'connected' time, in msec. ( H01DSCON * 128 / 1000 / H01DSSSC )
	SERV	Average I/O service time, in msec. ( H01DSSRV / 10 )
	RESP	Average I/O response time, in msec. ( H01DSRSP / 10 / H01DSSSC )
	CUWT	Average control unit queueing time, in msec. ( H01DSCUW * 128 / 1000 / H01DSSSC )
	QUED	Average IO request block queue length ( H01DSHFQ / H01DSHFC )
	BUSY	Average DASD busy percentage ( H01DSBSY / H01DSSMP )

Display	Variable Name	Variable Description and Derivation
	%READ	Average percentage of I/O requests that were for READ operations only ( H01DSRDS * 100 / (H01DSRDS+H01DSWRT) )
	CYLS	Average nr. of cylinders skipped per SEEK order ( H01DSMOV / H01DSSKS )
	RECOV	Average number of SSCHs executed for error recovery purposes ( H01DSERR / H01DSCT2 )
	DLY/S	Average throttle I/O delay rate, per second ( H01DSTHD / H01DSCT2 / TRNDRELT )

### CP Services Activity

Although CP system services activity variables are also described for trend record X'FC05' (see [“Performance Variables of Trend Records X'FC05”](#) on page 1025) they have been redefined for being selected as 'SYSTEM' variables, too, in which case the value returned will always be the total activity for all processors. (The formulas in the following descriptions show only the calculations made per processor.)

SYSTEM	IUVM_T	Total IUCV and VMCF activity per second ( (H05ISEND+H05IRECV+H05IFAIL +H05VMCFS+H05VMCFR+H05VMCFF) / TRNDRELT )
	IUCV_T	Total IUCV activity per second ( (H05ISEND+H05IRECV+H05IFAIL) / TRNDRELT )
	VMCF_T	Total VMCF activity per second ( (H05VMCFS+H05VMCFR+H05VMCFF) / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	*TOT_T	Total system services activity per second ( (H05IMSAS+H05IMSGS+H05IRPIS +H05IBIOS+H05IMONS+H05ISIGS +H05ICCSS+H05ISPLS+H05ISYMS +H05IACCS+H05ILOGS+H05ICRMS +H05IIDES+H05ICNFS +H05IMSAR+H05IMSGR+H05IRPIR +H05IBIOR+H05IMONR+H05ISIGR +H05ICCSR+H05ISPLR+H05ISYMR +H05IACCR+H05ILOGR+H05ICRMR +H05IIDER+H05ICNFR +H05IMSAF+H05IMSGF+H05IRPIF +H05IBIOF+H05IMONF+H05ISIGF +H05ICCSF+H05ISPLF+H05ISYMF +H05IACCF+H05ILOGF+H05ICRMF +H05IIDEF+H05ICNFF) / TRNDRELT )
	*MSA_T	Total *MSGALL IUCV activity per second ( (H05IMSAS+H05IMSAR+H05IMSAF) / TRNDRELT )
	*MSG_T	Total *MSG IUCV activity per second ( (H05IMSGS+H05IMSGR+H05IMSGF) / TRNDRELT )
	*RPI_T	Total *RPI IUCV activity per second ( (H05IRPIS+H05IRPIR+H05IRPIF) / TRNDRELT )
	*BLK_T	Total *BLOCKIO IUCV activity per second ( (H05IBIOS+H05IBIOR+H05IBIOF) / TRNDRELT )
	*MON_T	Total *MONITOR IUCV activity per second ( (H05IMONS+H05IMONR+H05IMONF) / TRNDRELT )
	*SIG_T	Total *SIGNAL IUCV activity per second ( (H05ISIGS+H05ISIGR+H05ISIGF) / TRNDRELT )
	*CCS_T	Total *CCS IUCV activity per second ( (H05ICCSS+H05ICCSR+H05ICCSF) / TRNDRELT )
	*SPL_T	Total *SPL IUCV activity per second ( (H05ISPLS+H05ISPLR+H05ISPLF) / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	*SYM_T	Total *SYMPTOM IUCV activity per second ( (H05ISYMS+H05ISYMR+H05ISYMF) / TRNDRELT )
	*ACC_T	Total *ACCOUNT IUCV activity per second ( (H05IACCS+H05IACCR+H05IACCF) / TRNDRELT )
	*LOG_T	Total *LOGREC IUCV activity per second ( (H05ILOGS+H05ILOGR+H05ILOGF) / TRNDRELT )
	*CRM_T	Total *CRM IUCV activity per second ( (H05ICRMS+H05ICRMR+H05ICRMF) / TRNDRELT )
	*IDE_T	Total *IDENT IUCV activity per second ( (H05IIDES+H05IIDER+H05IIDEF) / TRNDRELT )
	*CNF_T	Total *CONFIG IUCV activity per second ( (H05ICNFS+H05ICNFR+H05ICNFF) / TRNDRELT )

#### CP Services from CP to Virtual Machines, Rates per Second

SYSTEM	IUVM>V	Total IUCV and VMCF rate/s, CP to VM ( (H05ISEND+H05VMCFS) / TRNDRELT )
	IUCV>V	Total IUCV activity/s, CP to VM ( H05ISEND / TRNDRELT )
	VMCF>V	Total VMCF activity/s, CP to VM ( H05VMCFS / TRNDRELT )
	*TOT>V	Total system services rate/s, CP to VM ( (H05IMSAS+H05IMSGS+H05IRPIS +H05IBIOS+H05IMONS+H05ISIGS +H05ICCSS+H05ISPLS+H05ISYMS +H05IACCS+H05ILOGS+H05ICRMS +H05IIDES+H05ICNFS) / TRNDRELT )
	*MSA>V	Total *MSGALL IUCV rate/s, CP to VM ( H05IMSAS / TRNDRELT )
	*MSG>V	Total *MSG IUCV rate/s, CP to VM ( H05IMSGS / TRNDRELT )



Display	Variable Name	Variable Description and Derivation
	*RPI>V	Total *RPI IUCV rate/s, CP to VM ( H05IRPIS / TRNDRELT )
	*BLK>V	Total *BLOCKIO IUCV rate/s, CP to VM ( H05IBIOS / TRNDRELT )
	*MON>V	Total *MONITOR IUCV rate/s, CP to VM ( H05IMONS / TRNDRELT )
	*SIG>V	Total *SIGNAL IUCV rate/s, CP to VM ( H05ISIGS / TRNDRELT )
	*CCS>V	Total *CCS IUCV rate/s, CP to VM ( H05ICCSS / TRNDRELT )
	*SPL>V	Total *SPL IUCV rate/s, CP to VM ( H05ISPLS / TRNDRELT )
	*SYM>V	Total *SYMPTOM IUCV rate/s, CP to VM ( H05ISYMS / TRNDRELT )
	*ACC>V	Total *ACCOUNT IUCV rate/s, CP to VM ( H05IACCS / TRNDRELT )
	*LOG>V	Total *LOGREC IUCV rate/s, CP to VM ( H05ILOGS / TRNDRELT )
	*CRM>V	Total *CRM IUCV rate/s, CP to VM ( H05ICRMS / TRNDRELT )
	*IDE>V	Total *IDENT IUCV rate/s, CP to VM ( H05IIDES / TRNDRELT )
	*CNF>V	Total *CONFIG IUCV rate/s, CP to VM ( H05ICNFS / TRNDRELT )

#### CP Services from Virtual Machines to CP, Rates per Second

SYSTEM	IUVM>C	Total IUCV and VMCF activity/s, VM to CP ( H05IRECV+H05VMCFR ) / TRNDRELT )
	IUCV>C	Total IUCV activity/s, VM to CP ( H05IRECV / TRNDRELT )
	VMCF>C	Total VMCF activity/s, VM to CP ( H05VMCFR / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	*TOT>C	Total system services rate/s, VM to CP ( (H05IMSAR+H05IMSGR+H05IRPIR +H05IBIOR+H05IMONR+H05ISIGR +H05ICCSR+H05ISPLR+H05ISYMR +H05IACCR+H05ILOGR+H05ICRMR +H05IIDER+H05ICNFR) / TRNDRELT )
	*MSA>C	Total *MSGALL IUCV rate/s, VM to CP ( H05IMSAR / TRNDRELT )
	*MSG>C	Total *MSG IUCV rate/s, VM to CP ( H05IMSGR / TRNDRELT )
	*RPI>C	Total *RPI IUCV rate/s, VM to CP ( H05IRPIR / TRNDRELT )
	*BLK>C	Total *BLOCKIO IUCV rate/s, VM to CP ( H05IBIOR / TRNDRELT )
	*MON>C	Total *MONITOR IUCV rate/s, VM to CP ( H05IMONR / TRNDRELT )
	*SIG>C	Total *SIGNAL IUCV rate/s, VM to CP ( H05ISIGR / TRNDRELT )
	*CCS>C	Total *CCS IUCV rate/s, VM to CP ( H05ICCSR / TRNDRELT )
	*SPL>C	Total *SPL IUCV rate/s, VM to CP ( H05ISPLR / TRNDRELT )
	*SYM>C	Total *SYMPTOM IUCV rate/s, VM to CP ( H05ISYMR / TRNDRELT )
	*ACC>C	Total *ACCOUNT IUCV rate/s, VM to CP ( H05IACCR / TRNDRELT )
	*LOG>C	Total *LOGREC IUCV rate/s, VM to CP ( H05ILOGR / TRNDRELT )
	*CRM>C	Total *CRM IUCV rate/s, VM to CP ( H05ICRMR / TRNDRELT )
	*IDE>C	Total *IDENT IUCV rate/s, VM to CP ( H05IIDER / TRNDRELT )
	*CNF>C	Total *CONFIG IUCV rate/s, VM to CP ( H05ICNFR / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
<b>CP Services, Failure Rates per Second</b>		
SYSTEM	IUVM_F	Total IUCV and VMCF failures per second ( (H05IFAIL+H05VMCFF) / TRNDRELT )
	IUCV_F	Total IUCV failures per second ( H05IFAIL / TRNDRELT )
	VMCF_F	Total VMCF failures per second ( H05VMCFF / TRNDRELT )
	*TOT_F	Total system services failures per second ( (H05IMSAF+H05IMSGF+H05IRPIF +H05IBIOF+H05IMONF+H05ISIGF +H05ICCSF+H05ISPLF+H05ISYMF +H05IACCF+H05ILOGF+H05ICRMF +H05IIDEF+H05ICNFF) / TRNDRELT )
	*MSA_F	Total *MSGALL IUCV failures per second ( H05IMSAF / TRNDRELT )
	*MSG_F	Total *MSG IUCV failures per second ( H05IMSGF / TRNDRELT )
	*RPI_F	Total *RPI IUCV failures per second ( H05IRPIF / TRNDRELT )
	*BLK_F	Total *BLOCKIO IUCV failures per second ( H05IBIOF / TRNDRELT )
	*MON_F	Total *MONITOR IUCV failures per second ( H05IMONF / TRNDRELT )
	*SIG_F	Total *SIGNAL IUCV failures per second ( H05ISIGF / TRNDRELT )
	*CCS_F	Total *CCS IUCV failures per second ( H05ICCSF / TRNDRELT )
	*SPL_F	Total *SPL IUCV failures per second ( H05ISPLF / TRNDRELT )
	*SYM_F	Total *SYMPTOM IUCV failures per second ( H05ISYMF / TRNDRELT )
	*ACC_F	Total *ACCOUNT IUCV failures per second ( H05IACCF / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	*LOG_F	Total *LOGREC IUCV failures per second ( H05ILOGF / TRNDRELT )
	*CRM_F	Total *CRM IUCV failures per second ( H05ICRMF / TRNDRELT )
	*IDE_F	Total *IDENT IUCV failures per second ( H05IIDEF / TRNDRELT )
	*CNF_F	Total *CONFIG IUCV failures per second ( H05ICNFF / TRNDRELT )
<b>System Log Data (SYSLOG)</b>		
SYSTEM	DONE	Fast CCW translation: successful rate/s ( H01SLFDO / TRNDRELT )
	ABORT	Fast CCW translation: Aborted rate/s ( H01SLFAB / TRNDRELT )
	NOTELIG	Fast CCW translation: Not eligible rate/s ( H01SLFNE / TRNDRELT )
	TOTCCW	Fast CCW translation: Total rate/s ( (H01SLFDO+H01SLFAB+H01SLFNE) / TRNDRELT )
	EXITTIME	Microseconds spent in user exits ( H01SLEXT * 1000000 )
<b>User Communication Totals (UCOMMLOG)</b>		
SYSTEM	MSG/S	Total message rate per second, for IUCV and VMCF communication ( H01UCTOT / TRNDRELT )
	ISEND	Total IUCV messages sent per second ( H01UCUSN / TRNDRELT )
	IRECV	Total IUCV messages received per second ( H01UCURC / TRNDRELT )
	IFAIL	Total IUCV messages failed per second ( H01UCUFL / TRNDRELT )
	VSEND	Total VMCF messages sent per second ( H01UCVSN / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	VRECV	Total VMCF messages received per second ( H01UCVRC / TRNDRELT )
	VFAIL	Total VMCF messages failed per second ( H01UCVFL / TRNDRELT )
	SENDQ	Average IUCV msgs. in send queue ( H01UCUSQ / H01UCUS1 )
	RECVQ	Average IUCV msgs. in receive queue ( H01UCUCQ / H01UCUS1 )
	REPLQ	Average IUCV msgs. in reply queue ( H01UCURQ / H01UCUS1 )
	QUEUE	Average VMCF messages queued ( H01UCVQU / H01UCUS1 )

# Performance Variables of Trend Records X'FC02'

Trend record X'FC02' contains data describing processor utilization and performance.

The selection format for record 02 variables is:

➤ *display*(           \*                                  ) ➤  
                  *cpuno*                  *varname*

***cpuno***

is the number of the processor for which data are to be extracted.

Totals for the values from all processors will be calculated if the processor number is specified as an asterisk '\*'.

***varname***

is the name of the selected variable, as shown under the heading 'Variable Name' in the table below.

The variables shown in the following table have been pre-defined. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
<b>Processor Activity Data</b>		
PRCLOG/ DSVSLOG/ PROCLOG	TOTAL	Total CPU utilization, in percent ( (H02USTIM+H02SYTIM) * 100 / TRNDRELT )
	USER	User CPU utilization, in percent ( H02USTIM * 100 / TRNDRELT )
	SYST	System CPU utilization, in percent ( H02SYTIM * 100 / TRNDRELT )
	EMUL	Emulation CPU utilization, in percent ( H02EMTIM * 100 / TRNDRELT )
	%CP	CP CPU utilization, in percent ( (H02USTIM+H02SYTIM-H02EMTIM) * 100 / TRNDRELT )
	%SPIN	CPU spin utilization, in percent ( H02SPIN * 100 / TRNDRELT )
	%LOGLD	Logical CPU load, in percent (based on CPU active and voluntary wait time, but excluding any suspended time when running 2nd level or in an LPAR) ( (H02USTIM+H02SYTIM) * 100 / (H02USTIM+H02SYTIM+H02WAIT))

Display	Variable Name	Variable Description and Derivation
	VECT	Total vector utilization, in percent ( (H02VFTIM+H02VOTIM) * 100 / TRNDRELT )
	%VEM	Vector emulation load, in percent ( H02VFTIM * 100 / TRNDRELT )
	VECLD	Vector load rate, per second ( H02VFLOD / TRNDRELT )
	SIML	Privileged instruction simulation rate per sec. ( H02PRVCT / TRNDRELT )
	DIAG	Diagnose instruction execution rate per sec. ( H02DIACT / TRNDRELT )
	SIGP	SIGP Ext Call receipt rate, per second ( H02SIGCT / TRNDRSMP )
	SSCH	SSCH and RSCH rate per second ( (H02SSCH+H02RSCH) / TRNDRSMP )
I	EMPTY	Percent of time DSVBK found empty ( H02EMPTY * 100 / H02HFSMP )
I	#USRS	Mean VMDBKs in DSVBK while not empty ( H02VMDBK / (H02HFSMP-H02EMPTY))
I	#MAST	Mean VMDBKs in master processor DSVBK ( H02MASTR / H02HFSMP )
	VMSTL	VMDBKs stolen from processor, per second ( H02VMSTO / TRNDRELT )
	TOMAST	VMDBKs moved to master proc. per second ( H02VMMAS / TRNDRELT )
	AVEMP	Seconds between available list empty ( TRNDRELT / H02AVEMP )
	BEL2GB	Pages moved/s below 2GB for translation ( H02PGB2G / TRNDRELT )
	PGIN	PGIN rate per second from expanded storage ( (H02FPGIN+H02PGINS) / TRNDRELT )
	%FAST	Percent of fastpath PGINS ( H02FPGIN * 100 / (H02FPGIN+H02PGINS) )

Display	Variable Name	Variable Description and Derivation
	PGRD/S	Page read rate per second ( H02PGRDS / TRNDRELT )
	MSG/S	Total IUCV and VMCF message rate per second ( (H02ISEVM+H02ISTVM+H02ISUVM +H02VSEVM+H02VSTVM+H02VSUVM) / TRNDRELT )

### Available List Management Data

Includes activity data shown on both the AVAILLOG and DEMNDLOG reports.

AVAILLOG	OBTAINBL	Page frames obtained/s from <2GB available list ( H02PGALO / TRNDRELT )
	OBTAINAB	Page frames obtained/s from >2GB available list ( H02AVOBT / TRNDRELT )
	RETURNBL	Page frames returned/s to <2GB available list ( H02AVRET / TRNDRELT )
	RETURNAB	Page frames returned/s to >2GB available list ( H02AVREG / TRNDRELT )
	EMPTYBL	Number of times <2GB available list was empty ( H02AVEMP )
	EMPTYAB	Number of times >2GB available list was empty ( H02AVEMG )
	SCAN1	Number of times demand scan ended with pass 1 ( H02SCAN1 )
	FRAMES1	Number of frames obtained with scan 1 ( H02FR1LD+H02FR1DM+H02FR1NS + H02FR1EL+H02FR1DL )
	SCAN2	Number of times demand scan ended with pass 2 ( H02SCAN2 )
	FRAMES2	Number of frames obtained with scan 2 ( H02FR2LD+H02FR2DM+H02FR2NS + H02FR2EL+H02FR2DL )
	ESCAN	Number of times demand scan ended with emergency pass ( H02SCANE )



Display	Variable Name	Variable Description and Derivation
	FRAMESE	Number of frames obtained with scan 1 ( H02FR1DM+H02FR1NS+H02FR1EL +H02FR1DL )
	SCANFAIL	Number of times demand scan failed obtaining enough page frames ( H02SCNFA )
	SCN1LDRM	Number of times scan ended with pass 1 of users in the long term dormant list ( H02SE1LD )
	SCN1DORM	Number of times scan ended with pass 1 of users in the dormant list ( H02SE1DM )
	SCN1NSS	Number of times scan ended with pass 1 of the system VMDBK, NSSs, DCSSs, ... ( H02SE1NS )
	SCN1ELIG	Number of times scan ended with pass 1 of users in the eligible list ( H02SE1EL )
	SCN1DISP	Number of times scan ended with pass 1 of users in the dispatch list ( H02SE1DL )
	SCN2LDRM	Number of times scan ended with pass 2 of users in the long term dormant list ( H02SE2LD )
	SCN2DORM	Number of times scan ended with pass 2 of users in the dormant list ( H02SE2DM )
	SCN2ELIG	Number of times scan ended with pass 2 of users in the eligible list ( H02SE2EL )
	SCN2DISP	Number of times scan ended with pass 2 of users in the dispatch list ( H02SE2DL )
	SCNEDORM	Number of times scan ended with the emergency pass of users in the dormant list ( H02SEEDM )

Display	Variable Name	Variable Description and Derivation
	SCNENSS	Number of times scan ended with the emergency pass of the system VMDBK, NSSs, DCSSs, ... ( H02SEENS )
	SCNEELIG	Number of times scan ended with the emergency pass of users in the eligible list ( H02SEEL )
	SCNEDISP	Number of times scan ended with the emergency pass of users in the dispatch list ( H02SEEDL )
	FRM1LDRM	Number of frames obtained with pass 1 from users in the long term dormant list ( H02FR1LD )
	FRM1DORM	Number of frames obtained with pass 1 from users in the dormant list ( H02FR1DM )
	FRM1NSS	Number of frames obtained with pass 1 from system VMDBK, NSSs, DCSSs, ... ( H02FR1NS )
	FRM1ELIG	Number of frames obtained with pass 1 from users in the eligible list ( H02FR1EL )
	FRM1DISP	Number of frames obtained with pass 1 from users in the dispatch list ( H02FR1DL )
	FRM2LDRM	Number of frames obtained with pass 2 from users in the long term dormant list ( H02FR2LD )
	FRM2DORM	Number of frames obtained with pass 2 from users in the dormant list ( H02FR2DM )
	FRM2NSS	Number of frames obtained with pass 2 from system VMDBK, NSSs, DCSSs, ... ( H02FR2NS )
	FRM2ELIG	Number of frames obtained with pass 2 from users in the eligible list ( H02FR2EL )

Display	Variable Name	Variable Description and Derivation
	FRM2DISP	Number of frames obtained with pass 2 from users in the dispatch list ( H02FR2DL )
	FRMEDORM	Number of frames obtained with the emergency pass from users in the dormant list ( H02FREDM )
	FRMENSS	Number of frames obtained with the emergency pass from system VMDBK, NSSs, DCSSs, ... ( H02FRENS )
	FRMEELIG	Number of frames obtained with the emergency pass from users in the eligible list ( H02FREEL )
	FRMEDISP	Number of frames obtained with the emergency pass from users in the dispatch list ( H02FREDL )

# Performance Variables of Trend Records X'FC03'

Trend record X'FC03' contains data describing logical partition load and performance.

The selection format for record 03 variables is:

➤ LPAR(            \*            ) ➤  
          lpaname      varname

**lpaname**

is the name of the logical partition for which data are to be extracted. Overall load values for the whole system will be calculated if the LPAR name is specified as an asterisk (\*).

**varname**

is the name of the selected variable, as shown under the **Variable Name** heading in the following table.

The variables shown in the following table have been pre-defined. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
For the physical load values we add the load values obtained for each of the involved logical processors, and finally divide the total by the number of physical processors in the complex.		
LPAR	%LOAD	Overall logical partition load, in percent of total system capacity. The formula for calculating the load for a <i>single</i> LPAR is: ( (sum of H03ACTM) * 100 / TRNDRELT / H03NRPRC) For overall totals ('*' selection) the general LPAR management overhead is also added, as follows: + (H03MGMTM * 100 / H03MGMEI / H03NRPRC)
For all of the remaining variables we calculate the average load for a logical processor, i.e. where values from multiple logical logical processors are to be considered we first build the sum of the values obtained for each logical processor and finally divide it by the number of logical processors we obtained data for. The formula shown indicates only how the values per logical processor are obtained.		
LPAR	%BUSY	Logical processor busy percentage ( H03ACTM * 100 / TRNDRELT )
	%OVHD	Logical processor LPAR management overhead, in percent ( (H03ACTM-H03NETT)* 100 / TRNDRELT )
	%SUSP	Logical processor suspended percentage ( (TRNDRELT-H03UTIM-H03STIM -H03WTIM) * 100 / TRNDRELT)
	%VMLD	Logical processor load as seen by guest ( (H03UTIM+H03STIM) * 100 / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	%LOGLD	'Logical' load of logical processor, based on non-suspended elapsed time ( (H03UTIM+H03STIM) * 100 / (H03UTIM+H03STIM+H03WTIM) )

## Performance Variables of Trend Records X'FC04'

Trend record X'FC04' contains data describing full-track cache (or minidisk cache) performance.

The selection format for record 04 variables is

➤ MDC(*varname*) ➤

### ***varname***

is the name of the selected variable, as shown under the heading 'Variable Name' in the table below.

The variables shown in the following table have been pre-defined. The same names have been chosen as for the equivalent variables defined for HISTLOG and HISTSUM files, i.e. they usually *do not* correspond directly to part of the header line of the MDCACHE or MDCSTOR performance display. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
MDC	MNIDEAL	FTC main storage frames count, ideal ( H04IDL MN )
	MNACT	FTC main storage frames count, actual ( H04ACT MN )
	MN>2GB	FTC main storage frames count, actual frames used above 2GB line ( H04G2GMN )
	MNMIN	FTC main storage frames, minimum set ( H04MIN MN )
	MNMAX	FTC main storage frames, maximum set ( H04MAX MN )
	MNDEL/S	FTC main storage page delete rate/s ( H04MNDEL / TRNDRELT )
	MNSTL/S	FTC main storage page steal rate/s ( H04MNSTL / TRNDRELT )
	MNBIAS	FTC main storage bias ( H04BIAMN / 100 )
	XSTAV	Nr. of XSTORE blocks, CP partition size ( H04CPPAR )
	XSIDEAL	FTC expanded storage blocks, ideal ( H04IDLXS )
	XSACT	FTC expanded storage blocks, actual ( H04MDCNE )

Display	Variable Name	Variable Description and Derivation
	MDNE	FTC total nr. of blocks (main & XSTORE) ( H04ACTMN + H04MDCNE )
	XSMIN	FTC exp. storage blocks, minimum set ( H04MINXS )
	XSMAX	FTC exp. storage blocks, maximum set ( H04MDCAC )
	FAIRSHR	Fair share limit (max. FTC blocks per user) ( H04FSHAR )
	XSDel/S	FTC XSTORE page delete rate per sec. ( H04XSREL / TRNDRELT )
	XSSTL/S	FTC XSTORE page steal rate per second ( H04XSSTL / TRNDRELT )
	XSBIAS	FTC expanded storage bias ( H04BIAXS / 100 )
	MDR/S	Minidisk cache read request rate/s ( H04RDREQ / TRNDRELT )
	MDH/S	Minidisk cache read request hits/s ( H04AVOID / TRNDRELT )
	%MDH	Percent minidisk cache read hits ( H04AVOID * 100 / H04RDREQ )
	MDPIN/S	FTC XSTORE block PGIN rate per second ( H04PGIN / TRNDRELT )
	MDPOUT/S	FTC XSTORE block PGOUT rate per second ( H04PGOUT / TRNDRELT )
	MDAVAGE	Average age of FTC XSTORE blocks, in sec. ( H04AGEXS )
	FSEX/S	FTC fair share exceeded rate per second ( H04FSEXC / TRNDRELT )
	BKINS/S	'In transit' inserts into waiting queue ( H04WQINS / TRNDRELT )
	INVR/S	Invalidation request rate per second ( H04INVRQ / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	INVB/S	Blocks invalidated per second ( H04INVBK / TRNDRELT )
	MDUS	Users with MDC inserts ( H04USRIN )



## Performance Variables of Trend Records X'FC05'

Trend record X'FC05' contains data describing CP system services' activity.

The selection format for record 05 variables is:

➤ CPSERV(            \*            ) ➤

*cpuno*                      *varname*

**српско**

is the number of the processor for which data are to be extracted.

Totals for the values from all processors will be calculated if the processor number is specified as an asterisk (\*).

***varname***

is the name of the selected variable, as shown under the **Variable Name** heading in the following table.

The variables shown in the following table have been pre-defined. The same names have been chosen as for the equivalent variables defined for HISTLOG and HISTSUM files, that is, they generally *do not* correspond directly to part of the header line of a performance display. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
<b>CP Services (ALL), Rates per Second</b>		
CPSERV	IUVM_T	Total IUCV and VMCF activity per second ( (H05ISEND+H05IRECV+H05IFAIL+H05VMCFS +H05VMCFR+H05VMCFF) / TRNDRELT )
	IUCV_T	Total IUCV activity per second ( (H05ISEND+H05IRECV+H05IFAIL) / TRNDRELT )
	VMCF_T	Total VMCF activity per second ( (H05VMCFS+H05VMCFR+H05VMCFF) / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	*TOT_T	Total system services activity per second ( (H05IMSAS+H05IMSGS+H05IRPIS +H05IBIOS+H05IMONS+H05ISIGS +H05ICCSS+H05ISPLS+H05ISYMS +H05IACCS+H05ILOGS+H05ICRMS +H05IIDES+H05ICNFS+H05ISEVS +H05ISEAS+H05ISESC+H05ISEVE +H05IMSAR+H05IMSGR+H05IRPIR +H05IBIOR+H05IMONR+H05ISIGR +H05ICCSR+H05ISPLR+H05ISYMR +H05IACCR+H05ILOGR+H05ICRMR +H05IIDER+H05ICNFR+H05ISTVS +H05ISTAS+H05ISTSC+H05ISTVEVE +H05IMSAF+H05IMSGF+H05IRPIF +H05IBIOF+H05IMONF+H05ISIGF +H05ICCSF+H05ISPLF+H05ISYMF +H05IACCF+H05ILOGF+H05ICRMF +H05IIDEF+H05ICNFF+H05ISUVS +H05ISUAS+H05ISUSC+H05ISUVE) / TRNDRELT )
	*MSA_T	Total *MSGALL IUCV activity per second ( (H05IMSAS+H05IMSAR+H05IMSAF) / TRNDRELT )
	*MSG_T	Total *MSG IUCV activity per second ( (H05IMSGS+H05IMSGR+H05IMSGF) / TRNDRELT )
	*RPI_T	Total *RPI IUCV activity per second ( (H05IRPIS+H05IRPIR+H05IRPIF) / TRNDRELT )
	*BLK_T	Total *BLOCKIO IUCV activity per second ( (H05IBIOS+H05IBIOR+H05IBIOF) / TRNDRELT )
	*MON_T	Total *MONITOR IUCV activity per second ( (H05IMONS+H05IMONR+H05IMONF) / TRNDRELT )
	*SIG_T	Total *SIGNAL IUCV activity per second ( (H05ISIGS+H05ISIGR+H05ISIGF) / TRNDRELT )
	*CCS_T	Total *CCS IUCV activity per second ( (H05ICCSS+H05ICCSR+H05ICCSF) / TRNDRELT )
	*SPL_T	Total *SPL IUCV activity per second ( (H05ISPLS+H05ISPLR+H05ISPLF) / TRNDRELT )
	*SYM_T	Total *SYMPTOM IUCV activity per second ( (H05ISYMS+H05ISYMR+H05ISYMF) / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	*ACC_T	Total *ACCOUNT IUCV activity per second ( (H05IACCS+H05IACCR+H05IACCF) / TRNDRELT )
	*LOG_T	Total *LOGREC IUCV activity per second ( (H05ILOGS+H05ILOGR+H05ILOGF) / TRNDRELT )
	*CRM_T	Total *CRM IUCV activity per second ( (H05ICRMS+H05ICRMR+H05ICRMF) / TRNDRELT )
	*IDE_T	Total *IDENT IUCV activity per second ( (H05IIDES+H05IIDER+H05IIDEF) / TRNDRELT )
	*CNF_T	Total *CONFIG IUCV activity per second ( (H05ICNFS+H05ICNFR+H05ICNFF) / TRNDRELT )

#### CP Services from CP to Virtual Machines, Rates per Second

CPSERV	IUVM>V	Total IUCV and VMCF rate/s, CP to VM ( (H05ISEND+H05VMCFS) / TRNDRELT )
	IUCV>V	Total IUCV activity/s, CP to VM ( H05ISEND / TRNDRELT )
	VMCF>V	Total VMCF activity/s, CP to VM ( H05VMCFS / TRNDRELT )
	*TOT>V	Total system services rate/s, CP to VM ( (H05IMSAS+H05IMSGS+H05IRPIS +H05IBIOS+H05IMONS+H05ISIGS +H05ICCSS+H05ISPLS+H05ISYMS +H05IACCS+H05ILOGS+H05ICRMS +H05IIDES+H05ICNFS+H05ISEVS +H05ISEAS+H05ISESC+H05ISEVE) / TRNDRELT )
	*MSA>V	Total *MSGALL IUCV rate/s, CP to VM ( H05IMSAS / TRNDRELT )
	*MSG>V	Total *MSG IUCV rate/s, CP to VM ( H05IMSGS / TRNDRELT )
	*RPI>V	Total *RPI IUCV rate/s, CP to VM ( H05IRPIS / TRNDRELT )
	*BLK>V	Total *BLOCKIO IUCV rate/s, CP to VM ( H05IBIOS / TRNDRELT )
	*MON>V	Total *MONITOR IUCV rate/s, CP to VM ( H05IMONS / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	*SIG>V	Total *SIGNAL IUCV rate/s, CP to VM ( H05ISIGS / TRNDRELT )
	*CCS>V	Total *CCS IUCV rate/s, CP to VM ( H05ICCSS / TRNDRELT )
	*SPL>V	Total *SPL IUCV rate/s, CP to VM ( H05ISPLS / TRNDRELT )
	*SYM>V	Total *SYMPTOM IUCV rate/s, CP to VM ( H05ISYMS / TRNDRELT )
	*ACC>V	Total *ACCOUNT IUCV rate/s, CP to VM ( H05IACCS / TRNDRELT )
	*LOG>V	Total *LOGREC IUCV rate/s, CP to VM ( H05ILOGS / TRNDRELT )
	*CRM>V	Total *CRM IUCV rate/s, CP to VM ( H05ICRMS / TRNDRELT )
	*IDE>V	Total *IDENT IUCV rate/s, CP to VM ( H05IIDES / TRNDRELT )
	*CNF>V	Total *CONFIG IUCV rate/s, CP to VM ( H05ICNFS / TRNDRELT )

#### CP Services from Virtual Machines to CP, Rates per Second

CPSERV	IUVM>C	Total IUCV and VMCF activity/s, VM to CP ( (H05IRECV+H05VMCFR) / TRNDRELT )
	IUCV>C	Total IUCV activity/s, VM to CP ( H05IRECV / TRNDRELT )
	VMCF>C	Total VMCF activity/s, VM to CP ( H05VMCFR / TRNDRELT )
	*TOT>C	Total system services rate/s, VM to CP ( (H05IMSAR+H05IMSGR+H05IRPIR +H05IBIOR+H05IMONR+H05ISIGR +H05ICCSR+H05ISPLR+H05ISYMR +H05IACCR+H05ILOGR+H05ICRMR +H05IIDER+H05ICNFR+H05ISTVS +H05ISTAS+H05ISTSC+H05ISTVE) / TRNDRELT )
	*MSA>C	Total *MSGALL IUCV rate/s, VM to CP ( H05IMSAR / TRNDRELT )
	*MSG>C	Total *MSG IUCV rate/s, VM to CP ( H05IMSGR / TRNDRELT )
	*RPI>C	Total *RPI IUCV rate/s, VM to CP ( H05IRPIR / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	*BLK>C	Total *BLOCKIO IUCV rate/s, VM to CP ( H05IBIOR / TRNDRELT )
	*MON>C	Total *MONITOR IUCV rate/s, VM to CP ( H05IMONR / TRNDRELT )
	*SIG>C	Total *SIGNAL IUCV rate/s, VM to CP ( H05ISIGR / TRNDRELT )
	*CCS>C	Total *CCS IUCV rate/s, VM to CP ( H05ICCSR / TRNDRELT )
	*SPL>C	Total *SPL IUCV rate/s, VM to CP ( H05ISPLR / TRNDRELT )
	*SYM>C	Total *SYMPTOM IUCV rate/s, VM to CP ( H05ISYMR / TRNDRELT )
	*ACC>C	Total *ACCOUNT IUCV rate/s, VM to CP ( H05IACCR / TRNDRELT )
	*LOG>C	Total *LOGREC IUCV rate/s, VM to CP ( H05ILOGR / TRNDRELT )
	*CRM>C	Total *CRM IUCV rate/s, VM to CP ( H05ICMR / TRNDRELT )
	*IDE>C	Total *IDENT IUCV rate/s, VM to CP ( H05IIDER / TRNDRELT )
	*CNF>C	Total *CONFIG IUCV rate/s, VM to CP ( H05ICNFR / TRNDRELT )

#### CP Services, Failure Rates per Second

CPSERV	IUVM_F	Total IUCV and VMCF failures per second ( (H05IFAIL+H05VMCFF) / TRNDRELT )
	IUCV_F	Total IUCV failures per second ( H05IFAIL / TRNDRELT )
	VMCF_F	Total VMCF failures per second ( H05VMCFF / TRNDRELT )
	*TOT_F	Total system services failures per second ( (H05IMSAF+H05IMSGF+H05IRPIF +H05IBIOF+H05IMONF+H05ISIGF +H05ICCSF+H05ISPLF+H05ISYMF +H05IACCF+H05ILOGF+H05ICRMF +H05IIDEF+H05ICNFF+H05ISUVS +H05ISUAS+H05ISUSC+H05ISUVE) / TRNDRELT )
	*MSA_F	Total *MSGALL IUCV failures per second ( H05IMSAF / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	*MSG_F	Total *MSG IUCV failures per second ( H05IMSGF / TRNDRELT )
	*RPI_F	Total *RPI IUCV failures per second ( H05IRPIF / TRNDRELT )
	*BLK_F	Total *BLOCKIO IUCV failures per second ( H05IBIOF / TRNDRELT )
	*MON_F	Total *MONITOR IUCV failures per second ( H05IMONF / TRNDRELT )
	*SIG_F	Total *SIGNAL IUCV failures per second ( H05ISIGF / TRNDRELT )
	*CCS_F	Total *CCS IUCV failures per second ( H05ICCSF / TRNDRELT )
	*SPL_F	Total *SPL IUCV failures per second ( H05ISPLF / TRNDRELT )
	*SYM_F	Total *SYMPTOM IUCV failures per second ( H05ISYMF / TRNDRELT )
	*ACC_F	Total *ACCOUNT IUCV failures per second ( H05IACCF / TRNDRELT )
	*LOG_F	Total *LOGREC IUCV failures per second ( H05ILOGF / TRNDRELT )
	*CRM_F	Total *CRM IUCV failures per second ( H05ICRMF / TRNDRELT )
	*IDE_F	Total *IDENT IUCV failures per second ( H05IIDEF / TRNDRELT )
	*CNF_F	Total *CONFIG IUCV failures per second ( H05ICNFF / TRNDRELT )

#### CP Services (ALL), Rates per Second

CPSERV	*VSW_T	Total *VSWITCH IUCV activity per second ( (H05ISEVS+H05ISTVS+H05ISUVS) / TRNDRELT )
	*ASY_T	Total *ASYNCMD IUCV activity per second ( (H05ISEAS+H05ISTAS+H05ISUAS) / TRNDRELT )
	*SCL_T	Total *SCLP IUCV activity per second ( (H05ISESC+H05ISTSC+H05ISUSC) / TRNDRELT )
	*VME_T	Total *VMEVENT IUCV activity per second ( (H05ISEVE+H05ISTVE+H05ISUVE) / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
<b>CP Services from CP to Virtual Machines, Rates per Second</b>		
	*VSW>V	Total *VSWITCH IUCV rate/s, CP to VM ( H05ISEVS / TRNDRELT )
	*ASY>V	Total *ASYNCMD IUCV rate/s, CP to VM ( H05ISEAS / TRNDRELT )
	*SCL>V	Total *SCLP IUCV rate/s, CP to VM ( H05ISESC / TRNDRELT )
	*VME>V	Total *VMEVENT IUCV rate/s, CP to VM ( H05ISEVE / TRNDRELT )
<b>CP Services from Virtual Machines to CP, Rates per Second</b>		
	*VSW>C	Total *VSWITCH IUCV rate/s, VM to CP ( H05ISTVS / TRNDRELT )
	*ASY>C	Total *ASYNCMD IUCV rate/s, VM to CP ( H05ISTAS / TRNDRELT )
	*SCL>C	Total *SCLP IUCV rate/s, VM to CP ( H05ISTSC / TRNDRELT )
	*VME>C	Total *VMEVENT IUCV rate/s, VM to CP ( H05ISTVE / TRNDRELT )
<b>CP Services, Failure Rates per Second</b>		
	*VSW_F	Total *VSWITCH IUCV failures per second ( H05ISUVS / TRNDRELT )
	*ASY_F	Total *ASYNCMD IUCV failures per second ( H05ISUAS / TRNDRELT )
	*SCL_F	Total *SCLP IUCV failures per second ( H05ISUSC / TRNDRELT )
	*VME_F	Total *VMEVENT IUCV failures per second ( H05ISUVE / TRNDRELT )

# Performance Variables of Trend Records X'FC06'

Trend record X'FC06' contains data describing channel utilization, based on HF sampling data. The values refer to the overall (physical) utilization of the channel, regardless of whether data collection was from a system running in an LPAR or not.

The selection format for record 06 variables is

➤ CHANNEL( *chpidvarname* ) ➤

**chpid**

is the hexadecimal channel path identifier (two hexadecimal digits) of the channel for which data are to be extracted.

**varname**

is the name of the selected variable, as shown under the heading 'Variable Name' in the table below.

The variable shown in the following table has been pre-defined. The description also includes the formula used to calculate the value from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
CHANNEL	AVE	Average channel busy percentage ( H06HFBSY * 100 / H06HFCNT )



## Performance Variables of Trend Records X'FC07'

---

Trend record X'FC07' contains data describing channel utilization, based on information collected by the Channel Measurement Facility. When collected by a system running in an LPAR, the values describe 'logical' channel utilization, i.e. the utilization by the system that also collected the data.

The selection format for record 07 variables is

►► LCHANNEL( *chpidvarname* ) ►◄

### ***chpid***

is the hexadecimal channel path identifier (two hexadecimal digits) of the channel for which data are to be extracted.

### ***varname***

is the name of the selected variable, as shown under the heading 'Variable Name' in the table below.

The variable shown in the following table has been pre-defined. The description also includes the formula used to calculate the value from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
LCHANNEL	AVE	Average logical channel busy percentage ( H07TOTBS / H07SAMPS )

# Performance Variables of Trend Records X'FC08'

Trend record X'FC08' contains data describing channel utilization, based on information collected by the Extended Channel Measurement Facility.

The selection format for record 08 variables is

➤ *display(chpidvarname)* ➤

**display(**  
is the name of a performance display that contains the variable, shown under the **Display** heading in the following table of selectable variables, and immediately followed by a left parenthesis.

**chpid**  
is the hexadecimal channel path identifier (two hexadecimal digits) of the channel for which data are to be extracted.

**varname**  
is the name of the selected variable, as shown under the **Variable Name** heading in the following table.

The variables shown in the following table have been pre-defined. The descriptions also include the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
<b>FICON Channel Utilization Data</b>		
FCHANNEL	%BUSY	Average total channel busy percentage ( H08TOTBS / H08SAMPS )
	%LBUSY	Average logical channel busy percentage, with channel load by the own partition ( H08LPABS / H08SAMPS )
The following fields are defined for FICON channels only; they are not valid for other channels.		
	T_BCY	Average total internal bus cycle utilization, in percent ( H08TBCYC * 100 * 1000000 / H08MXIBC / H08TSTMP / 128 )
	T_WUN	Average total channel work units utilization, in percent ( H08TWUNI * 100 * 1000000 / H08MXWUN / H08TSTMP / 128 )
	T_DUW	Average total data unit writes utilization, in percent ( H08TDUWR * 100 * 1000000 / H08MXDUW / H08TSTMP / 128 )

Display	Variable Name	Variable Description and Derivation
	T_DUR	Average total data unit reads utilization, in percent ( H08TDURD * 100 * 1000000 / H08MXDUR / H08TSTMP / 128 )
	L_WUN	Average logical channel work units utilization by the own partition, in percent ( H08LWUNI * 100 * 1000000 / H08MXWUN / H08TSTMP / 128 )
	L_DUW	Average logical data unit writes utilization by the own partition, in percent ( H08LDUWR * 100 * 1000000 / H08MXDUW / H08TSTMP / 128 )
	L_DUR	Average logical data unit reads utilization by the own partition, in percent ( H08LDURD * 100 * 1000000 / H08MXDUR / H08TSTMP / 128 )
	WRITE/S	Average total bytes written per second ( H08TDUWR * H08BYTES * 1000000 / H08TSTMP / 128 )
	READ/S	Average total bytes read per second ( H08TDURD * H08BYTES * 1000000 / H08TSTMP / 128 )

#### Hipersocket Channel Activity Data

HIPSOCK	T_MSGS	Average total messages sent rate per second. ( H08TMSNT * 1000000 / H08TSTMP / 128 )
	T_DUNITS	Average total data units sent rate per second. ( H08TDUSN * 1000000 / H08TSTMP / 128 )
	T_NOBUFF	Average total unsuccessful send rate due to insufficient receive buffers, per second. ( H08TUNBF * 1000000 / H08TSTMP / 128 )
	L_MSGS	Average messages sent rate per second for the own partition. ( H08LMSNT * 1000000 / H08TSTMP / 128 )

Display	Variable Name	Variable Description and Derivation
	L_DUNITS	Average data units sent rate per second for the own partition. ( H08LDUSN * 1000000 / H08TSTMP / 128 )
	L_NOBUFF	Average unsuccessful send rate per second due to insufficient receive buffers, for the own partition. ( H08LUNBF * 1000000 / H08TSTMP / 128 )
	L_OTHER	Average unsuccessful send rate per second due to other reasons, for the own partition. ( H08LUNSS * 1000000 / H08TSTMP / 128 )

# Performance Variables of Trend Records X'FC3A'

Trend record X'FC3A' contains overall transaction and response time data statistics for all users, as shown on the SYSTRANS display.

The selection format for record 3A variables is:

➤ SYSTRANS( *varname* ) ➤

### *varname*

is the name of the selected variable, as shown under the **Variable Name** heading in the following table.

Transaction counts are provided for the number of transactions falling below one of twelve thresholds for a specific variable, or above the highest one, that is, 13 values are available for each group of trivial and non-trivial transaction counts. The percentage of transactions for which the value of the selected variable fell below a particular threshold can thus be calculated as

$$\frac{(\text{sum of transaction counts below threshold}) * 100}{(\text{sum of transaction counts for all 13 fields})}$$

The variables shown in the following table have been pre-defined. Since the full formulas for all variables would take up a lot of space, they have not been included here.

Display	Variable Name	Variable Description and Derivation
<b>Transaction Response Times:</b>		
SYSTRANS	%TTM <i>n.nn</i>	Percentage of trivial transactions with a response time lower than <i>n.nn</i>
	%NTM <i>n.nn</i>	Percentage of non-trivial transactions with a response time lower than <i>n.nn</i> The string <i>n.nn</i> must be specified as one of the following threshold values: 0.01, 0.02, 0.04, 0.06, 0.10, 0.20, 0.40, 0.60, 1.00, 2.00, 4.00, or 6.00. The values must be entered exactly as shown, i.e. as three digits with a decimal point after the first digit. Example: Specify SYSTRANS(%TTM0.20) to select the percentage of trivial transactions that had a response time of less than 0.20 seconds.
<b>CPU Usage per Transaction:</b>		
SYSTRANS	%TCP <i>nnnn</i>	Percentage of trivial transactions with a CPU usage lower than <i>nnnn</i> milliseconds

Display	Variable Name	Variable Description and Derivation
	%NCP $nnnn$	Percentage of non-trivial transactions with a CPU usage lower than $nnnn$ milliseconds The string $nnnn$ must be specified as one of the following threshold values, exactly as shown: 1, 2, 5, 10, 20, 50, 100, 200, 500, 1000, 2000 or 5000.
<b>Disk I/Os per Transaction:</b>		
SYSTRANS	%TIO $nnn$	Percentage of trivial transactions with less than $nnn$ disk I/O operations
	%NIO $nnn$	Percentage of non-trivial transactions with less than $nnn$ disk I/O operations The string $nnn$ must be specified as one of the following threshold values, exactly as shown: 1, 2, 4, 6, 10, 20, 40, 60, 100, 200, 400 or 600.
<b>Page Faults per Transaction:</b>		
SYSTRANS	%TPF $nnn$	Percentage of trivial transactions with less than $nnn$ page faults
	%NPF $nnn$	Percentage of non-trivial transactions with less than $nnn$ page faults The string $nnn$ must be specified as one of the following threshold values, exactly as shown: 1, 2, 4, 6, 10, 20, 40, 60, 100, 200, 400 or 600.
<b>XSTORE Page Reads per Transaction:</b>		
SYSTRANS	%TXR $nnn$	Percentage of trivial transactions with less than $nnn$ XSTORE page reads
	%NXR $nnn$	Percentage of non-trivial transactions with less than $nnn$ XSTORE page reads The string $nnn$ must be specified as one of the following threshold values, exactly as shown: 1, 2, 4, 6, 10, 20, 40, 60, 100, 200, 400 or 600.
<b>IUCV Messages per Transaction:</b>		
SYSTRANS	%TIU $nnn$	Percentage of trivial transactions with less than $nnn$ IUCV interactions

Display	Variable Name	Variable Description and Derivation
	%NIU <i>nnn</i>	Percentage of non-trivial transactions with less than <i>nnn</i> IUCV interactions The string <i>nnn</i> must be specified as one of the following threshold values, exactly as shown: 1, 2, 4, 6, 10, 20, 40, 60, 100, 200, 400 or 600.

#### VMCF Messages per Transaction:

SYSTRANS	%TVM <i>nnn</i>	Percentage of trivial transactions with less than <i>nnn</i> VMCF interactions
	%NVM <i>nnn</i>	Percentage of non-trivial transactions with less than <i>nnn</i> VMCF interactions The string <i>nnn</i> must be specified as one of the following threshold values, exactly as shown: 1, 2, 4, 6, 10, 20, 40, 60, 100, 200, 400 or 600.

#### Spool Reads per Transaction:

SYSTRANS	%TSR <i>nnn</i>	Percentage of trivial transactions with less than <i>nnn</i> spool reads
	%NSR <i>nnn</i>	Percentage of non-trivial transactions with less than <i>nnn</i> spool reads The string <i>nnn</i> must be specified as one of the following threshold values, exactly as shown: 1, 2, 4, 6, 10, 20, 40, 60, 100, 200, 400 or 600.

#### Spool Writes per Transaction:

SYSTRANS	%TSW <i>nnn</i>	Percentage of trivial transactions with less than <i>nnn</i> spool writes
	%NSW <i>nnn</i>	Percentage of non-trivial transactions with less than <i>nnn</i> spool writes The string <i>nnn</i> must be specified as one of the following threshold values, exactly as shown: 1, 2, 4, 6, 10, 20, 40, 60, 100, 200, 400 or 600.

## Performance Variables of Trend Records X'FC3C'

Trend record X'FC3C' contains data describing NSS performance for individual shared segments.

The selection format for record 3C variables is:

➤ NSS(*segnovname*) ➤

### ***segno***

is the spool file number of a shared segment for which data are to be extracted, in up to four decimal digits. Note that shared segment data cannot be extracted by segment name because multiple segments with the same name may be active concurrently.

### ***varname***

is the name of the selected variable, as shown under the heading 'Variable Name' in the table below.

The variables shown in the following table have been pre-defined. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
NSS	SHARE	Number of users accessing the segment in shared mode ( H3CUSHRD / TRNDRSMP )
	N-SHR	Number of users accessing the segment in non-shared mode ( H3CUSNSH / TRNDRSMP )
	SAVD	Number of pages saved in this segment ( H3CPGSAV )
	NDAT	Number of pages defined with a 'no data saved' page descriptor ( H3CPGNSV )
	RES.	Number of private resident pages in the segment ( H3CPGRES / TRNDRSMP )
	XSTOR	Number of expanded storage blocks occupied ( H3CXSPGS / TRNDRSMP )
	TRANS	Page translation rate, per second ( H3CPGTRN / TRNDRELT )
	STEAL	Page frame steal rate, per second ( H3CPGSTL / TRNDRELT )
	PGRDS	Page read rate, per second ( H3CPGRDS / TRNDRELT )
	PGWRT	Page write rate, per second ( H3CPGWRT / TRNDRELT )




Display	Variable Name	Variable Description and Derivation
	X-RDS	Page-in rate per second from expanded storage ( H3CXPGRD / TRNDRELT )
	X-WRT	Page-out rate per second to expanded storage ( H3CXPGWR / TRNDRELT )
	X-MIG	Page migration rate per second from expanded storage to disk ( H3CXMIGR / TRNDRELT )

# Performance Variables of Trend Records X'FC41', X'FC42' & X'FC43'

Trend record X'FC41' contains data describing user resource utilization and wait states for individual users. The same information is available for user classes in record X'FC42' and as totals for all users on the system in record X'FC43'.

The selection format for record 41/42/43 variables is:



### *display*(

Is the name of a performance display that contains the variable, shown under the heading 'Display' in the following table of selectable variables, and immediately followed by a left parenthesis.

### *userid*

Is the user identification of a virtual machine, i.e. data will be extracted from trend records X'FC41'.

### *'user class'*

Is the user class name, i.e. data will be extracted from user class trend records X'FC42'. The string may contain up to 8 characters including blanks, and it *must* be enclosed by single quotes.

### *\*ALL\**

Indicates data extraction from trend record X'FC43', for system totals of all users.

### *varname*

Is the name of the selected variable, as shown under the heading 'Variable Name' in the table below.

The variables shown in the following table have been pre-defined. Their description also includes the formula used to calculate the values from the original trend record data fields. When processing user class records X'FC42' or user totals records X'FC43' the values will also be divided by the number of users (H41USRCT) to obtain averages per user, where appropriate.

Display	Variable Name	Variable Description and Derivation
<b>General User Load Data</b>		
USER	%CPU	Total CPU load, in percent ( H41TCPU * 100 / TRNDRELT )
	TCPU	Total CPU time used, in seconds ( H41TCPU )
	VCPU	Virtual CPU time used, in seconds ( H41VCPU )
	RATIO	Total/virtual CPU time ratio ( H41TCPU / H41VCPU )
	%VEC	Vector facility load, in percent ( (H41VVEC+H41OVEC) * 100 / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	TOTAL	Total virtual I/O activity, in IOs per sec. ( (H41IOCON+H41IOCTC+H41IODSD +H41IOUR+H41IOOTH+H41IOD98) / TRNERELT )
	DASD	I/O activity to DASD, in IOs per second ( H41IODSD / TRNDRELT )
	AVOID	DASD I/O rate avoided due to minidisk cache ( (H41IOAVO+H41IOVDI) / TRNDRELT )
	DIAG98	Diagnose X'98' I/O rate, diagnoses per sec. ( H41IOD98 / TRNDRELT )
	UR	I/O activity to UR devices, in IOs per sec. ( H41IOUR / TRNDRELT )
	PG/S	Pages read and written per second ( (H41PGRD+H41PGWRT) / TRNDRELT )
	LOGGED	Minutes logged on ( H41TIMLG / 60 )
	ACTIVE	Minutes active ( H41TIMAC / 60 )
	PAGES	Total spool pages processed ( H41SPPGS )
	SPG/S	Spool pages processed per second ( H41SPPGS / TRNDRELT )
	MDC/S	Minidisk cache inserts per second ( H41MDCIN / TRNDRELT )

#### User Paging and Storage Utilization Data

UPAGE	OWNED	Non-base address spaces owned ( H41DSOWN / TRNDRSMP )
	READS	Page read rate per second ( H41PGRD / TRNDRELT )
	WRITE	Page write rate per second ( H41PGWRT / TRNDRELT )
	STEALS	Page steal rate per second ( H41PGSTL / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	>2GB>	Pages moved/s below 2GB for translation ( H41PGB2G / TRNDRELT )
	X>MS	XSTORE to main PGIN rate per second ( H41XTOMN / TRNDRELT )
	MS>X	Main to XSTORE PGOUT rate per second ( H41MNTOX / TRNDRELT )
	X>DS	XSTORE to DASD migrations per second ( H41XTODS / TRNDRELT )
	WSS	Average nr. of pages in working set ( H41PGWSS / TRNDRSMP )
	L<2GB	Average pages locked below 2GB ( H41PGLCK / TRNDRSMP )
	L>2GB	Average pages locked above 2GB ( H41PGLCG / TRNDRSMP )
	RESRVD	Average number of reserved pages ( H41PGRSV / TRNDRSMP )
	R<2GB	Average pages resident below 2GB ( H41PGRES / TRNDRSMP )
	R>2GB	Average pages resident above 2GB ( H41PGRS2 / TRNDRSMP )
	XSTOR	Average nr. of expanded storage pages ( H41PGXST / TRNDRSMP )
	DASD	Average nr. of pages on DASD ( H41PGDSD / TRNDRSMP )
	SIZE	Virtual storage size, in bytes ( H41STSIZ )
<b>User Communication Data</b>		
UCOMM	MSG/S	Total message rate per second, for IUCV and VMCF communication ( (H41IUSND+H41IURCV+H41IUFAI +H41VMSND+H41VMRCV+H41VMFAI) / TRNDRELT )
	ISEND	IUCV messages sent per second ( H41IUSND / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	Irecv	IUCV messages received per second ( H41IURCV / TRNDRELT )
	IFail	IUCV messages failed per second ( H41IUFAI / TRNDRELT )
	VSEND	VMCF messages sent per second ( H41VMSND / TRNDRELT )
	VRECV	VMCF messages received per second ( H41VMRCV / TRNDRELT )
	VFail	VMCF messages failed per second ( H41VMFAI / TRNDRELT )
	SENDQ	Average IUCV msgs. in send queue ( H41IUSNQ / TRNDRSMP )
	RECVQ	Average IUCV msgs. in receive queue ( H41IURCQ / TRNDRSMP )
	REPLQ	Average IUCV msgs. in reply queue ( H41IURPQ / TRNDRSMP )
	QUEUE	Average VMCF messages queued ( H41VMCFQ / TRNDRSMP )
	TOTAL	Average number of IUCV connections ( H41IUCON / TRNDRSMP )
	MAX.	Maximum number of IUCV connections allowed ( H41MXCON / TRNDRSMP )

#### User Wait States

USTAT	%ACT	Percent of time user was found to be active (not dormant) ( (H41HFSMP+H41HFSVD-H41HFDRM) * 100 / H41HFSMP )
	%RUN	Percent of active time user was found running ( H41HFRUN * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	%CPU	Percent of active time user was in CPU wait ( H41HFCPU * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )

Display	Variable Name	Variable Description and Derivation
	%LDG	Percent of active time user was loading ( H41HFLDG * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	%PGW	Percent of active time user was in page wait ( H41HFPGW * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	%IOW	Percent of active time user was in I/O wait ( H41HFIOW * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	%SIM	% of active time user was in inst. simul. wait ( H41HFSIM * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	%TIW	Percent of active time user was in test-idle wait ( H41HFTIW * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	%CFW	Percent of active time user was in CF wait ( H41HFCFW * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	%TI	Percent of active time user was in SVM & TI wait ( H41HFSVT * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	%EL	Percent of active time user was in SVM wait and in the eligible list ( H41HFSVE * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	%DM	Percent of active time user was in SVM wait and dormant ( H41HFSVD * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	%IOA	% of active time user was in 'I/O active' wait ( H41HFIOA * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	%PGA	% of active time user was in 'page active' wait ( H41HFPGA * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	%LIM	Percent of active time user was in limit list wait ( H41HFLIM * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )

Display	Variable Name	Variable Description and Derivation
	%OTH	% of active time user was in another wait state ( H41HFOTH * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	Q0	Percent of active time user was in queue 0 ( H41HFQU0 * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	Q1	Percent of active time user was in queue 1 ( H41HFQU1 * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	Q2	Percent of active time user was in queue 2 ( H41HFQU2 * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	Q3	Percent of active time user was in queue 3 ( H41HFQU3 * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )
	E0-3	Percent of active time user was in elig. list ( (H41HFEL0+H41HFEL1+H41HFEL2+H41HFEL3) * 100 / (H41HFSMP+H41HFSVD-H41HFDRM) )


#### User QDIO Related Activity Data

UQDIO	TVS	TVS instruction intercept rate per second ( H41QDTV / TRNDRSMP )
	SVS	SVS instruction intercept rate per second ( H41QDSVS / TRNDRELT )
	/SEC	Alert interrupt rate per second ( H41QDALE / TRNDRELT )
	VIATPI	Adapter interrupts delivered/s to guest via TPI ( H41QDTPI / TRNDRELT )
	INSIE	Adapter interrupts delivered/s to guest in SIE ( H41QDHDW / TRNDRELT )
	BYCP	Adapter interrupts stacked per secondn by CP ( H41QDVIR / TRNDRELT )

# Performance Variables of Trend Records X'FC44', X'FC45' & X'FC46'

Trend record X'FC44' contains transaction and response time data for individual users. The same information is available for user classes in record X'FC45' and as totals for all users on the system in record X'FC46'.

The selection format for record 44/45/46 variables is:



**display(**  
Is the name of a performance display that contains the variable, shown under the heading 'Display' in the following table of selectable variables, and immediately followed by a left parenthesis.

**userid**  
Is the user identification of a virtual machine, i.e. data will be extracted from trend records X'FC44'.

**'user class'**  
Is the user class name, i.e. data will be extracted from user class trend records X'FC45'. The string may contain up to 8 characters including blanks, and it *must* be enclosed by single quotes.

**\*ALL\***  
Indicates data extraction from trend record X'FC46', for system totals of all users.

**varname**  
Is the name of the selected variable, as shown under the heading 'Variable Name' in the table below.

The variables shown in the following table have been pre-defined. The variable description also includes the formula used to calculate the values from the original trend record data fields. When processing user class records X'FC45' or user totals records X'FC46' the values will also be divided by the number of users (H44USRCT) to obtain averages per user, where appropriate.

Display	Variable Name	Variable Description and Derivation
<b>User Transaction Data</b>		
UTRANS	COUNT	Total number of transactions ( H44TRNCT )
	/HOUR	Number of transactions per hour ( H44TRNCT * 3600 / TRNDRELT )
	MSEC	Total CPU time per transaction, in msec ( H44CPU * 1000 / H44TRNCT )
	I/O	Number of disk I/Os per transaction ( H44DSKIO / H44TRNCT )
	FLTS	Number of page faults per transaction ( H44PGFLT / H44TRNCT )



Display	Variable Name	Variable Description and Derivation
	READS	Number of PGINs per transaction ( H44XTOMN / H44TRNCT )
	SP_RD	Number of spool pages read per transaction ( H44SPRDS / H44TRNCT )
	SP_WR	Number of spool pages written per transaction ( H44SPWRT / H44TRNCT )
	IU_BY	Number of IUCV & VMCF sends per transaction ( (H44IUCBY+H44VMCBY) / H44TRNCT )
	IU_TO	Number of IUCV & VMCF receives per transaction ( (H44IUCTO+H44VMCTO) / H44TRNCT )
<b>User Response Time Data</b>		
URESP	TIME	Think time between transactions, in seconds ( H44THNK / H44THCNT )
	TRIV.	Trivial response time, in seconds ( H44TRVTM / H44TRVCT / 1000 )
	NTRIV	Non-trivial response time, in seconds ( H44NTRTM / H44NTRCT / 1000 )
	MEAN	Average response time for all transactions ( (H44TRVTM+H44NTRTM) / (H44TRVCT+H44NTRCT) / 1000 )
	TRV/H	Trivial transaction rate per hour ( H44TRVCT * 3600 / TRNDRELT )
	NTRV/H	Non-trivial transaction rate per hour ( H44NTRCT * 3600 / TRNDRELT )
	TOTL/H	Total transaction rate per hour ( (H44TRVCT+H44NTRCT) * 3600 / TRNDRELT )
	TRV-T	Trivial transaction count ( H44TRVCT )
	NTRV-T	Non-trivial transaction count ( H44NTRCT )
	TOTL-T	Total transaction count ( H44TRVCT + H44NTRCT )

# Performance Variables of Trend Records X'FC51'

Trend record X'FC51' contains data describing I/O processor utilization and performance.

The selection format for record 51 variables is:



**cpuno**  
is the number of the I/O processor for which data are to be extracted. Totals for the values from all I/O processors will be calculated if the processor number is specified as an asterisk '\*'.

**varname**  
is the name of the selected variable, as shown under the heading 'Variable Name' in the table below.

The variables shown in the following table have been pre-defined. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
IOPROCLG	BEG_SSCH	Begin SSCH rate per second ( H51BSSCH / TRNDRELT )
	I/O_INT	I/O interrupts generated per second ( H51IOINT / TRNDRELT )
	%BUSY	I/O processor busy percentage ( H51PRBSY * 100 / (H51PRBSY+H51PRIDL) )
	CHANNEL	Percentage of 'begin SSCH' which found a channel to be already busy ( H51CHBSY * 100 / H51BSSCH )
	SWITCH	Percentage of 'begin SSCH' which found a switch to be already busy ( H51SWBSY * 100 / H51BSSCH )
	CU	Percentage of 'begin SSCH' which found a ctl. unit to be already busy ( H51CUBSY * 100 / H51BSSCH )
	DEVICE	Percentage of 'begin SSCH' which found the target device to be already busy ( H51DVBSY * 100 / H51BSSCH )

## Performance Variables of Trend Records X'FC61'

Trend record X'FC61' contains data describing DASD performance for individual disks.

The selection format for record 61 variables is:

➤➤ **DEVICE**( *devnovarname* ) ➤➤

### **devno**

is the device number of the disk for which data are to be extracted, in normal hexadecimal notation.

### **varname**

is the name of the selected variable, as shown under the heading 'Variable Name' in the table below.

The variables shown in the following table have been pre-defined. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
DEVICE	LINKS	Number of minidisk links to the disk ( H61LINKS )
	I/O	Total SSCH & RSCH rate per second ( H61SSCH / TRNDRELT )
	AVOID	I/Os avoided rate per second ( H61AVOID / TRNDRELT )
	PEND	Average 'function pending' time, in msec. ( H61FPNT * 128 / 1000 / H61SSCH )
	DISC	Average 'disconnected' time, in msec. ( (H61DSCT+H61DAOT) * 128 / 1000 / H61SSCH )
	CONN	Average 'connected' time, in msec. ( H61CONT * 128 / 1000 / H61SSCH )
	SERV	Average I/O service time, in msec. ( PEND + DISC + CONN )
	RESP	Average I/O response time, in msec. ( (H61HFQUD * TRNDRELT / H61HFCNT / H61SSCH) + SERV )
	CUWT	Average control unit queueing time, in msec. ( H61CUBSY * 128 / 1000 / H61SSCH )
	QUED	Average IO request block queue length ( H61HFQUD / H61HFCNT )

Display	Variable Name	Variable Description and Derivation
	BUSY	Percentage of elapsed time that DASD was busy ( (H61FPNT+H61DSCT+H61CAOT+H61CONT) * 128 / 10000 / TRNDRELT )
	READ	Percentage of I/O requests that were for READ operations only ( H61READS * 100 / (H61READS+H61WRITE) )
	CYLS	Average nr. of cylinders skipped per SEEK order ( H61SKCYL / H61SKCNT )
	SSCH	Number of SSCHs executed for error recovery purposes ( H61RECOV )
	SET/S	Throttle I/O rate set for the disk (I/Os per sec.) ( H61THRRRT )
	DLY/S	Throttle I/O delay rate, per second ( H61THRDL / TRNDRELT )

## Performance Variables of Trend Records X'FC65'

Trend record X'FC65' contains data describing DASD cache activity and performance for individual disks.  
The selection format for record 65 variables is:

►► *display(devnovarname)* ◄◄

### **display(**

is the name of a performance display that contains the variable, shown under the heading 'Display' in the following table of selectable variables, and immediately followed by a left parenthesis.

### **devno**

is the device number of the disk for which data are to be extracted, in normal hex. notation.

### **varname**

is the name of the selected variable, as shown under the **Variable Name** heading in the following table.

The variables shown in the following table have been pre-defined. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
<b>General CU-Cached Disks Data</b>		
CACHDBSE	IO/S	Total I/O rate per second seen by cache CU ( (H65RNORR+H65WNORR+H65RSEQR+H65WSEQR +H65RCFWR+H65WCFWR+H65ICREQ+H65BYREQ) / TRNDRELT )
	%READ	Percent cachable read operations, out of total cache I/O activity ( (H65RNORR+H65RSEQR+H65RCFWR) * 100 / (H65RNORR+H65RSEQR+H65RCFWR +H65WNORR+H65WSEQR+H65WCFWR) ) +H65ICREQ+H65BYREQ) )
	%RDHIT	Overall read hit percentage ( (H65RNORH+H65RSEQH+H65RCFWH) * 100 / (H65RNORR+H65RSEQR+H65RCFWR) )
	%WRHIT	Overall write hit percentage ( (H65WNORH+H65WSFWH+H65WCFWH) * 100 / (H65WNORR+H65WSEQR+H65WCFWR) )
	ICL/S	Inhibit cache load rate, per second ( H65ICREQ / TRNDRELT )
	BYP/S	Bypass cache rate, per second ( H65BYREQ / TRNDRELT )

### **Cache Extended Function Data**

Display	Variable Name	Variable Description and Derivation
CACHEXT	CACHE	Total I/O rate per second seen by cache CU ( (H65RNORR+H65WNORR+H65RSEQR+H65WSEQR +H65RCFWR+H65WCFWR+H65ICREQ+H65BYREQ) / TRNDRELT )
	SCMBK	Total I/O rate per second by our system ( H61SSCH / TRNDRELT )
	N-SEQ	Non-sequential read rate per second ( (H65RNORR+H65RCFWR) / TRNDRELT )
	SEQ	Sequential read rate per second ( H65RSEQR / TRNDRELT )
	FW	Total 'fast write' rate (cache FW & DASD FW) ( (H65WCFWR+H65WNFWR+H65WSFWR) / TRNDRELT )
	READ	Percent cachable read operations ( (H65RNORR+H65RSEQR+H65RCFWR) * 100 / (H65RNORR+H65RSEQR+H65RCFWR +H65WNORR+H65WSEQR+H65WCFWR) )
	TOT	Overall read and write hit percentage ( (H65RNORH+H65RSEQH+H65RCFWH +H65WNORH+H65WSFWH+H65WCFWH) * 100 / (H65RNORR+H65RSEQR+H65RCFWR +H65WNFWR+H65WSFWR+H65WCFWR) )
	RDHT	Overall read hit percentage ( (H65RNORH+H65RSEQH+H65RCFWH) * 100 / (H65RNORR+H65RSEQR+H65RCFWR) )
	WRT	Overall write hit percentage ( (H65WNORH+H65WSFWH+H65WCFWH) * 100 / (H65WNFWR+H65WSFWR+H65WCFWR) )
	DFW	DASD 'fast write' hit percentage ( (H65WNORH+H65WSFWH) * 100 / (H65WNFWR+H65WSFWR) )
	CFW	Cache 'fast write' write hit percentage ( H65WCFWH * 100 / H65WCFWR )
	%DFW	Percent of DASD fast writes, out of total writes ( (H65WNFWR+H65WSFWR) * 100 / (H65RNORR+H65WSEQR+H65WCFWR) )

Display	Variable Name	Variable Description and Derivation
	%CFW	Percent of write cache fast writes, out of total writes ( H65WCFWR * 100 / (H65WNORR+H65WSEQR+H65WCFWR) )
	NORM	Percent of normal (non-sequential) stage transfers, out of total cache transfer operations ( H65NDACA * 100 / (H65RNORR+H65WNORR+H65RSEQR+H65WSEQR +H65RCFWR+H65WCFWR+H65ICREQ+H65BYREQ)
	%SEQ	Percent of sequential stage transfers, out of total cache transfer operations ( H65SDACA * 100 / (H65RNORR+H65WNORR+H65RSEQR+H65WSEQR +H65RCFWR+H65WCFWR+H65ICREQ+H65BYREQ)
	DEST	Percent of destage (cache to DASD) transfers, out of total cache transfer operations ( H65DESTA * 100 / (H65RNORR+H65WNORR+H65RSEQR+H65WSEQR +H65RCFWR+H65WCFWR+H65ICREQ+H65BYREQ)
	BYPS	Percentage of DASD FW requests bypassed (forced destages to DASD because of nonvolatile space constraints) ( H65NCADA * 100 / (H65RNORR+H65WNORR+H65RSEQR+H65WSEQR +H65RCFWR+H65WCFWR+H65ICREQ+H65BYREQ)
	+ICL	Percentage of 'bypass cache' and 'inhibit cache load' requests, out of total cache transfer operations ( (H65BYREQ+H65ICREQ) * 100 / (H65RNORR+H65WNORR+H65RSEQR+H65WSEQR +H65RCFWR+H65WCFWR+H65ICREQ+H65BYREQ)

## Performance Variables of Trend Records X'FC68'

Trend record X'FC68' contains data describing utilization and performance for system areas on CP owned disks.

The selection format for record 68 variables is:

►► CPOWNED( *devnovarname* ) ►◄

### ***devno***

is the device number of the disk for which data are to be extracted, in normal hexadecimal notation.

### ***varname***

is the name of the selected variable, as shown under the heading 'Variable Name' in the table below.

The variables shown in the following table have been pre-defined. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
CPOWNED	P-RDS	Page read rate from the disk, per second ( H68PAGRD / TRNDRELT )
	P-WRT	Page write rate to the disk, per second ( H68PAGWR / TRNDRELT )
	S-RDS	Spool read rate from the disk, per second ( H68SPLRD / TRNDRELT )
	S-WRT	Spool write rate to the disk, per second ( H68SPLWR / TRNDRELT )
	TOTAL	Total spool and page read and write rate ( (H68PAGRD+H68PAGWR+H68SPLRD +H68SPLWR) / TRNDRELT )
	+RSCH	I/O rate per second (SSCH plus RSCH oper.) ( H68SSCH / TRNDRELT )
	FERES	Total user interferes (user I/Os interfering with CP I/O activity) per mon. sample interval ( H68USINT / TRNDRSMP )
	LNGTH	Average I/O queue length for paging and spooling I/Os ( H68RWQUD / TRNDRSMP )
	/PAGE	I/O service time per page, in msec. ( H68SERV * 16 / 1000 / TRNDRSMP )
	TIME	MLOAD time (smoothed I/O response time per page) in msec. ( H68MLOAD * 16 / 1000 / TRNDRSMP )



Display	Variable Name	Variable Description and Derivation
	SIZE	Average paging block size ( $H68BKSIZ * H68BKIOS / H68BKIOS$ ) Note that a multiplication and division by H68BKIOS is done to force a 'zero divide' condition, and an invalid SIZE value, when no page block I/Os were executed.
	ALLOC	Percentage device was used for allocation after being looked at ( $H68CHOSN * 100 / H68LOKAT$ )

## Performance Variables of Trend Records X'FC6D'

Trend record X'FC6D' contains QDIO device activity data.

The selection format for record 6D variables is:

➤ QDIO( *devnovarname* ) ➤

### **devno**

is the device number of the QDIO device for which data are to be extracted.

### **varname**

is the name of the selected variable, as shown under the **Variable Name** heading in the following table.

The variables shown in the following table have been pre-defined. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
<b>QDIO Device Activity Data</b>		
QDIO	IN	Number of input queues ( H6DINPQU )
	OUT	Number of output queues ( H6DOUTQU )
	GST-R	Guest SIGA-r instruction rate ( H6DGSIGR / TRNDRELT )
	GST-W	Guest SIGA-w instruction rate ( H6DGSIGW / TRNDRELT )
	GST-S	Guest SIGA-s instruction rate ( H6DGSIGS / TRNDRELT )
	CP-R	CP SIGA-r instruction rate ( H6DCSIGR / TRNDRELT )
	CP-W	CP SIGA-w instruction rate ( H6DCSIGW / TRNDRELT )
	CP-S	CP SIGA-s instruction rate ( H6DCSIGS / TRNDRELT )
	WBUFFER	QDIO buffers written to adapter rate ( H6DBUFWR / TRNDRELT )
	WBYTES	Bytes written to adapter per second ( H6DBYTWR / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	RBUFFER	QDIO buffers read from adapter rate ( H6DBUFRD / TRNDRELT )
	RBYTES	Bytes read from adapter per second ( H6DBYTRD / TRNDRELT )
	TOTAL	Total PCI interrupt rate per second ( H6DPCITO / TRNDRELT )
	REQD	PCI interrupt rate requiring a response ( H6DPCIRR / TRNDRELT )
	CARDED	PCI interrupts discarded per second ( H6DPCIDS / TRNDRELT )

#### QEBSM Specific Device Activity Data

QDIO	SCOMPL	SQBS instr/s completely processed by CP ( H6DSQVM / TRNDRELT )
	SPARTL	SQBS instr/s partially processed by CP ( H6DSQVMP / TRNDRELT )
	ECOMPL	EQBS instr/s completely processed by CP ( H6DEQVM / TRNDRELT )
	EPARTL	EQBS instr/s partially processed by CP ( H6DEQVMP / TRNDRELT )
	SIGA-WT	SIGA-wt instructions executed/s by CP ( H6DSIGWT / TRNDRELT )
	SIGA-RT	SIGA-rt instructions executed/s by CP ( H6DSIGRT / TRNDRELT )
	SIGA-LT	SIGA-lt instructions executed/s by CP ( H6DSIGLT / TRNDRELT )
	GENERAL	Unsuccessful attempts per second to obtain a general QIOBD/QSD lock for serializing input or output ( H6DQLKF / TRNDRELT )
	OUTPUT	The number of unsuccessful attempts per second to obtain a QIOBD/QSD output lock. ( H6DQOLKF / TRNDRELT )
	PIN	'Pin' requests per second processed by CP ( H6DPPIN / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	UNPIN	'Unpin' requests per second processed by CP ( H6DPUPIN / TRNDRELT )
	EXH/S	FOBLK resources exhausted rate per second ( H6DRESEX / TRNDRELT )

## Performance Variables of Trend Records X'FC6F'

Trend record X'FC6F' contains SCSI device activity data.

The selection format for record 6F variables is:

➤ SCSI( *devnovarname* ) ➤

### **devno**

is the device number of the SCSI device for which data are to be extracted.

### **varname**

is the name of the selected variable, as shown under the **Variable Name** heading in the following table.

The variables shown in the following table have been pre-defined. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
<b>SCSI Device Activity Data</b>		
SCSI	BLOCK	Number of bytes in a block ( H6ZDNB )
	KBYTE	Kilobytes transferred per second ( H6ZDXR )
	FERS	Transfers per second ( H6ZDNX / TRNDRELT )
	SEEKS	Number of seek operations per second ( H6ZDNS / TRNDRELT )
	READ	Number of blocks read per second ( H6ZDNR / TRNDRELT )
	WRITTEN	Number of blocks written per second ( H6ZDNW / TRNDRELT )
	CONNECT	Percent of time connected (device busy) ( H6ZDTM * 128 / 10000 / TRNDRELT )

## Performance Variables of Trend Records X'FC71'

Trend record X'FC71' contains data describing DASD SEEK activity for individual disks.

The selection format for record 71 variables is:

►► *display(devnovarname)* ►►

### ***display***(

is the name of a performance display that contains the variable, shown under the heading 'Display' in the following table of selectable variables, and immediately followed by a left parenthesis.

### ***devno***

is the device number of the disk for which data are to be extracted, in normal hexadecimal notation.

### ***varname***

is the name of the selected variable, as shown under the heading 'Variable Name' in the table below.

The variables shown in the following table have been pre-defined. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
<b>SEEK Distances Data</b>		
SEEKDIST	MSEC	Device disconnected time, in msec. ( H71DSCTM+H71DAOTM ) * 128 / 1000 / H71SSCH )
	READ	Percentage of SEEKS with only read operations ( H71RDS * 100 / H71SKCT )
	NON0	Percentage of SEEKS with a nonzero cylinder displacement ( (H71SKCT-H71SKC0) * 100 / H71SKCT )
	DIST	Mean nonzero SEEKS distance ( H71SKDI / (H71SKCT-H71SKC0) )
	SEEKS	Total number of SEEK operations ( H71SKCT )
	0	Percentage of SEEKS with a displacement of zero ( H71DSP0 * 100 / H71TSMPS )
	5	Percentage of SEEKS with a displacement in the range of 1 to 5 cylinders ( H71DSP1 * 100 / H71TSMPS )
	10	Percentage of SEEKS with a displacement in the range of 6 to 10 cylinders ( H71DSP2 * 100 / H71TSMPS )

Display	Variable Name	Variable Description and Derivation
	20	Percentage of SEEKs with a displacement in the range of 11 to 20 cylinders ( H71DSP3 * 100 / H71TSMPS )
	50	Percentage of SEEKs with a displacement in the range of 21 to 50 cylinders ( H71DSP4 * 100 / H71TSMPS )
	100	Percentage of SEEKs with a displacement in the range of 51 to 100 cylinders ( H71DSP5 * 100 / H71TSMPS )
	200	Percentage of SEEKs with a displacement in the range of 101 to 200 cylinders ( H71DSP6 * 100 / H71TSMPS )
	400	Percentage of SEEKs with a displacement in the range of 201 to 400 cylinders ( H71DSP7 * 100 / H71TSMPS )
	800	Percentage of SEEKs with a displacement in the range of 401 to 800 cylinders ( H71DSP8 * 100 / H71TSMPS )
	>800	Percentage of SEEKs with a displacement exceeding 800 cylinders ( H71DSP9 * 100 / H71TSMPS )
<b>SEEK Locations Data</b>		
SEEKLOC	MSEC	Device disconnected time, in msec. ( (H71DSCTM+H71DAOTM) * 128 / 1000 / H71SSCH )
	READ	Percentage of SEEKs with only read operations ( H71RDS * 100 / H71SKCT )
	NON0	Percentage of SEEKs with a nonzero cylinder displacement ( (H71SKCT-H71SKC0) * 100 / H71SKCT )
	DIST	Mean nonzero SEEKs distance ( H71SKDI / (H71SKCT-H71SKC0) )
	SEEKS	Total number of SEEK operations ( H71SKCT )

Display	Variable Name	Variable Description and Derivation
	<=10	Percentage of SEEKs with a target cylinder in the range of 0-10% ( H71LOC0 * 100 / H71TSMPS )
	<=20	Percentage of SEEKs with a target cylinder in the range of 11-20% ( H71LOC1 * 100 / H71TSMPS )
	<=30	Percentage of SEEKs with a target cylinder in the range of 21-30% ( H71LOC2 * 100 / H71TSMPS )
	<=40	Percentage of SEEKs with a target cylinder in the range of 31-40% ( H71LOC3 * 100 / H71TSMPS )
	<=50	Percentage of SEEKs with a target cylinder in the range of 41-50% ( H71LOC4 * 100 / H71TSMPS )
	<=60	Percentage of SEEKs with a target cylinder in the range of 51-60% ( H71LOC5 * 100 / H71TSMPS )
	<=70	Percentage of SEEKs with a target cylinder in the range of 61-70% ( H71LOC6 * 100 / H71TSMPS )
	<=80	Percentage of SEEKs with a target cylinder in the range of 71-80% ( H71LOC7 * 100 / H71TSMPS )
	<=90	Percentage of SEEKs with a target cylinder in the range of 81-90% ( H71LOC8 * 100 / H71TSMPS )
	90<	Percentage of SEEKs with a target cylinder in the range of 91-100% ( H71LOC9 * 100 / H71TSMPS )



## Performance Variables of Trend Records X'FCA2'

Trend record X'FCA2' contains shared file system data for individual servers.

The selection format for record A2 variables is:

►► *display(useridvarname)* ►◄

### **display(**

Is the name of a performance display that contains the variable, shown under the heading 'Display' in the following table of selectable variables, and immediately followed by a left parenthesis.

### **userid**

Is the user identification of the SFS server for which data are to be extracted.

### **varname**

is the name of the selected variable, as shown under the **Variable Name** heading in the following table.

The variables shown in the following table have been pre-defined. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
<b>Overall SFS performance:</b>		
SFSLOG	FPRCOUNT	Number of file pool requests ( sum of HA2CF010 to HA2CF054 + sum of HA2CF097 to HA2CF101 + sum of HA2CF103 to HA2CF109 + sum of HA2CF117 to HA2CF121 + sum of HA2CF124 to HA2CF130 + sum of HA2CF139 to HA2CF142 + sum of HA2CF145 to HA2CF146 ) This value is referenced as 'total file pool requests' in many of the formulas for calculating the remaining variables.
	FPRRATE	File pool requests rate per second ( (total file pool requests) / TRNDRELT )
	FPRSERV	Service time per file pool request ( (HA2CF055+HA2CF110) / 10000 / (total file pool requests) )
	FPRCPU	CPU time per file pool request ( HA2TTIME / (total file pool requests) )
	FPRLOCK	Lock wait time per file pool request ( HA2CF072 / 10000 / (total file pool requests) )

Display	Variable Name	Variable Description and Derivation
	FPRBLKIO	Block I/O time per file pool request ( (HA2CF090+HA2CF116) / 10000 / (total file pool requests) )
	FPRESM	Ext. security manager wait time per FPR ( HA2CF009 / 10000 / (total file pool requests) )
	FPROTH	Other time needed per file pool request, calculated as ( FPRSERV - FPRCPU - FPRLOCK - FPRBLKIO - FPRESM )
	%BUSY	Total server busy percentage ( %CPU + %PGWT + %CKPT + %QSAM )
	%CPU	Percentage of time using CPU ( HA2TTIME * 100 / TRNDRELT )
	%PGWT	Percentage of time spent in page wait ( HA2HFPGW * 100 / HA2HFSMP )
	%CKPT	Percentage of time spent in checkpoint wait ( HA2CF005 / 100 / TRNDRELT )
	%QSAM	Percentage of time waiting for QSAM I/O ( HA2CF075 / 100 / TRNDRELT )
	AGNTACT	Average number of agents that were active ( HA2CF055 / 100000 / TRNDRELT )
	AGNTHELD	Average number of agents that were held ( HA2CF060 / 100000 / TRNDRELT )
	ROLLBACK	Number of deadlocks with roll-backs ( HA2CF123 )
Display	Variable Name	Variable Description and Derivation

#### SFS I/O Activity per File Pool Request:

Display	Variable Name	Variable Description and Derivation
SFSIOLOG	FPRCOUNT	Number of file pool requests ( sum of HA2CF010 to HA2CF054 + sum of HA2CF097 to HA2CF101 + sum of HA2CF103 to HA2CF109 + sum of HA2CF117 to HA2CF121 + sum of HA2CF124 to HA2CF130 + sum of HA2CF139 to HA2CF142 + sum of HA2CF145 to HA2CF146 )
	IOTOTAL	Total I/O requests per file pool request ( (HA2CF082+HA2CF115+HA2CF144 +(sum of HA2CF091 to HA2CF096) ) / (total file pool requests) )
	IOREAD	File read requests per file pool request ( HA2CF091 / (tot. file pool req.) )
	IOWRITE	File write requests per file pool request ( HA2CF092 / (tot. file pool req.) )
	IOCATRD	Catalog read requests per file pool request ( HA2CF093 / (tot. file pool req.) )
	IOCATWR	Catalog write requests per file pool request ( HA2CF094 / (tot. file pool req.) )
	IOCMDRD	Control minidisk read requests per file pool req. ( HA2CF095 / (tot. file pool req.) )
	IOCMDWR	Control minidisk write requests per file pool request ( HA2CF096 / (tot. file pool req.) )
	IOLOGRD	Log read requests per file pool request ( HA2CF082 / (tot. file pool req.) )
	IOLOGWR	Log write requests per file pool request ( (HA2CF115+HA2CF144) / (total file pool requests) )
	BLK/BIO	Average number of blocks transferred per block I/O request ( (sum of HA2CF076 to HA2CF083) / ((sum of HA2CF084 to HA2CF089) +HA2CF082+HA2CF143)

Display	Variable Name	Variable Description and Derivation
	BLK/IO	Average number of blocks transferred per I/O request ( (sum of HA2CF076 to HA2CF083) / ((sum of HA2CF091 to HA2CF096) +HA2CF082+HA2CF144)
	SACCALL	Average Storage Access Component calls per file pool request ( HA2CF062 / (tot. file pool req.) )
	BIOTIME	Average time per block I/O request ( (HA2CF090+HA2CF116) / ((sum of HA2CF084 to HA2CF089) +HA2CF082+HA2CF143) / 10000 )
	LOCKTIME	Average time per lock wait ( HA2CF072 / ((sum of HA2CF063 to HA2CF071) +HA2CF122) / 10000 )
	CKPTTIME	Average time spent processing a checkpoint ( HA2CF005 / HA2CF004 / 10000 )
	ESMTIME	Average time spent per ESM exit call ( HA2CF009 / HA2CF008 / 10000 )

#### SFS Requests Percentages:

Variables indicating the relative frequency of selected request types, compared to the total of all file pool requests.

SFSREQ	FPRCOUNT	Number of file pool requests ( sum of HA2CF010 to HA2CF054 + sum of HA2CF097 to HA2CF101 + sum of HA2CF103 to HA2CF109 + sum of HA2CF117 to HA2CF121 + sum of HA2CF124 to HA2CF130 + sum of HA2CF139 to HA2CF142 + sum of HA2CF145 to HA2CF146 )
	OPENRD	Percentage of open file read requests ( HA2CF031 * 100 / (total file pool requests) )
	OPENUPD	Percentage of open update requests ( (HA2CF030+HA2CF032+HA2CF033) * 100 / (total file pool requests) )

Display	Variable Name	Variable Description and Derivation
	READ	Percentage of read file requests ( HA2CF040 * 100 / (total file pool requests) )
	WRITE	Percentage of write file requests ( HA2CF054 * 100 / (total file pool requests) )
	CLOSE	Percentage of close file requests ( HA2CF014 * 100 / (total file pool requests) )
	DELETE	Percentage of delete file requests ( HA2CF020 * 100 / (total file pool requests) )
	GETDIR	Percentage of get directory entry requests ( HA2CF024 * 100 / (total file pool requests) )
	RENAME	Percentage of rename requests ( HA2CF047 * 100 / (total file pool requests) )
	CRALIAS	Percentage of create alias requests ( HA2CF017 * 100 / (total file pool requests) )
	CRDELDIR	Percentage of create and delete directory req. ( (HA2CF018+HA2CF019) * 100 / (total file pool requests) )
	OGCDIR	Percentage of open/get/close directory requests ( (HA2CF013+HA2CF023+HA2CF029) * 100 / (total file pool requests) )
	AUTH	Percentage of grant and revoke authorization req. ( (HA2CF026+HA2CF049) * 100 / (total file pool requests) )
	LOCK	Percentage of lock and unlock requests ( (HA2CF028+HA2CF052) * 100 / (total file pool requests) )
	QUERY	Percentage of query requests ( ((sum of HA2CF034 to HA2CF039) +HA2CF119+HACF0120) * 100 / (total file pool requests) )

Display	Variable Name	Variable Description and Derivation
	REFRESH	Percentage of refresh directory requests ( HA2CF045 * 100 / (total file pool requests) )
	OTHER	Percentage of all other requests, i.e. ( 100 -OPENRD -OPENUPD -READ - WRITE -CLOSE -DELETE -GETDIR - RENAME-CRALIAS-CRDELDIR-OGCDIR - AUTH -LOCK -QUERY -REFRESH )

## Performance Variables of Trend Records X'FCA4'

Trend record X'FCA4' contains multitasking user data for individual users.

The selection format for record A4 variables is:

➤ MTUSERS( *useridvarname* ) ➤

### **userid**

Is the user identification of the virtual machine for which data are to be extracted.

### **varname**

Is the name of the selected variable, as shown under the heading 'Variable Name' in the table below.

The variables shown in the following table have been pre-defined. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
MTUSERS	CRE/S	Thread create rate per second ( HA4CRCNT / TRNDRELT )
	MS/CRE	Average elapsed time per thread creation, in milliseconds ( HA4CRTIM * 1000 / HA4CRCNT )
	DEL/S	Thread delete rate per second ( HA4DLCNT / TRNDRELT )
	MS/DEL	Average elapsed time per thread deletion, in milliseconds ( HA4DLTIM * 1000 / HA4DLCNT )
	SLOW	Regular path switch rate per second ( HA4SLWSW / TRNDRELT )
	FAST	Fast path switch rate per second ( HA4FSTSW / TRNDRELT )
	BLOCKED	Average number of threads blocked ( HA4BLOKD / TRNDRSMP )
	PROC.	Highest number of processes that were concurrently defined ( HA4MXPRC )
	THRDS	Highest number of threads that were concurrently defined ( HA4MXTHR )
	FAILS	Number of POSIX create failures during the period covered by the trend record ( HA4CRFAI )

## Performance Variables of Trend Records X'FCA6'

Trend record X'FCA6' contains data describing overall TCP/IP server activity and performance for individual servers.

The selection format for record A6 variables is:

►► *display(useridvarname)* ►►

### **display(**

Is the name of a performance display that contains the variable, shown under the heading 'Display' in the following table of selectable variables, and immediately followed by a left parenthesis.

### **userid**

Is the user identification of the TCP/IP server for which data are to be extracted.

### **varname**

is the name of the selected variable, as shown under the **Variable Name** heading in the following table.

The variables shown in the following table have been pre-defined. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
<b>TCP/IP Activity Log Data</b>		
TCPACTLG	CONINI	TCP connection opens initiated per second ( HA6COPI / TRNDRELT )
	CONACC	TCP connection opens accepted per second ( HA6COPA / TRNDRELT )
	CONFAIL	TCP connection open failures per second ( HA6COPF / TRNDRELT )
	CONRESET	TCP connections resets per second ( HA6CRST / TRNDRELT )
	SEGRECV	TCP segments received per second ( HA6SRCV / TRNDRELT )
	SEGXMIT	TCP segments transmitted per second ( HA6STRN / TRNDRELT )
	SEGRXMIT	TCP segments re-transmitted per second ( HA6SRTR / TRNDRELT )
	SEGRCERR	TCP segments received with errors per second ( HA6SRCE / TRNDRELT )
	SEGXRST	TCP segments transmitted with reset per second ( HA6STRR / TRNDRELT )



Display	Variable Name	Variable Description and Derivation
	ARPREQRC	ARP requests received per second ( HA6ARCV / TRNDRELT )
	ARPREPXM	ARP replies transmitted per second ( HA6ARPL / TRNDRELT )
	ARPREQXM	ARP requests transmitted per second ( HA6AREQ / TRNDRELT )

#### **TCP/IP Buffer Pool Data**

*Average buffer pool levels:*

TCPBPLOG	AVEACB	Average number of free activity control blocks ( HA6RACBL )
	AVECCB	Average number of free client control blocks ( HA6RCCBL )
	AVEENV	Average number of free envelope control blocks ( HA6RENVL )
	AVELENV	Average number of free large envelope ctl. blks ( HA6RLEVL )
	AVERCB	Average number of raw IP control blocks ( HA6RRCBL )
	AVESCB	Average number of socket control blocks ( HA6RSCBL )
	AVEBSD	Average number of BSD style socket ctl. blocks ( HA6RBSDL )
	AVETCB	Average number of TCP control blocks ( HA6RTCBL )
	AVEUCB	Average number of UDP control blocks ( HA6RUCBL )
	AVEDBF	Average number of regular data buffers ( HA6RDBFL )
	AVESDB	Average number of small data buffers ( HA6RSDBL )

Display	Variable Name	Variable Description and Derivation
	AVETDB	Average number of tiny data buffers ( HA6RTDBL )
	AVESGA	Average number of segment acknowledgment control blocks ( HA6RSAPL )
	AVEFPSP	Average number of fixed page storage pool pages ( HA6RFPSL )
<i>Allocated buffer pool numbers:</i>		
TCPBPLOG	AVEFPAVL	Average locked fixed page storage pool pages that are available (not in use) ( HA6RFPLA )
	AVEFPUSE	Average locked fixed page storage pool pages that are allocated to users ( HA6RFPLU )
	ALLACB	Allocated activity control blocks ( HA6IACB )
	ALLCCB	Allocated client control blocks ( HA6ICCB )
	ALLENV	Allocated envelope control blocks ( HA6IENV )
	ALLLENV	Allocated large envelope control blocks ( HA6ILEV )
	ALLRCB	Allocated raw IP control blocks ( HA6IRCB )
	ALLSCB	Allocated socket control blocks ( HA6ISCB )
	ALLBSD	Allocated BSD-style socket control blocks ( HA6IBSD )
	ALLTCB	Allocated TCP control blocks ( HA6ITCB )
	ALLUCB	Allocated UDP control blocks ( HA6IUCB )

Display	Variable Name	Variable Description and Derivation
	ALLDBF	Allocated regular data buffers ( HA6IDBF )
	ALLSDB	Allocated small data buffers ( HA6ISDB )
	ALLTDB	Allocated tiny data buffers ( HA6ITDB )
	ALLSGA	Allocated segment acknowledgment control blocks ( HA6ISGA )
	ALLHOST	Allocated host pool blocks ( HA6IHOS )
	ALLATR	Allocated address translation control blocks ( HA6IATR )
	ALLIPR	Allocated IP routing control blocks ( HA6IIPR )
	ALLFPSP	Allocated fixed page storage pool pages ( HA6IFPSP )

*Minimum buffer pool levels found:*

For most buffer pools this is the actual minimum level, updated with each monitor sample interval, but for some pools the information is only available from the initial configuration data: these are marked with (I).

TCPBPLOG	MINACB	Minimum activity control blocks level ( HA6NACBM )
	MINCCB	Minimum client control blocks level ( HA6NCCBM )
	MINENV	Minimum regular envelope control blocks level ( HA6NENVM )
	MINLENV	Minimum large envelope control blocks level ( HA6NLEVM )
	MINRCB	Minimum raw IP control blocks level ( HA6NRCBM )
	MINSCB	Minimum socket control blocks level ( HA6NSCBM )

Display	Variable Name	Variable Description and Derivation
	MINBSD	Minimum BSD-style socket control blocks level ( HA6NBSDM )
	MINTCB	Minimum TCP control blocks level ( HA6NTCBM )
	MINUCB	Minimum UDP control blocks level ( HA6NUCBM )
	MINDBF	Minimum regular data buffers level ( HA6NDBFM )
	MINSDB	Minimum small data buffers level ( HA6NSDBM )
	MINTDB	Minimum tiny data buffers level ( HA6NTDBM )
	MINSGA	Minimum segment acknowledgment control blocks level ( HA6NSAPM )
	MINHOST	Minimum host pool blocks level (I) ( HA6IHOS+16 )
	MINATR	Minimum address translation control blocks level (I) ( HA6IATR+16 )
	MINIPR	Minimum IP routing control blocks level (I) ( HA6IIPR+16 )
	MINFPSP	Minimum fixed page storage pool level ( HA6NFPSM )
<b>TCP/IP Data Transfer</b>		
TCPDATLG	IPRECV	IP packets received per second ( HA6PRRC / TRNDRELT )
	IPRCHDER	IP packets received with header errors per second ( HA6PRHE / TRNDRELT )
	IPRCADER	IP packets received with address errors per second ( HA6PRAE / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	DGFORW	Incoming IP datagrams forwarded per second ( HA6DFOR / TRNDRELT )
	DGUKPROT	Incoming IP datagrams with unknown protocol received per second ( HA6DUPR / TRNDRELT )
	DGINDISC	Incoming IP datagrams discarded per second ( HA6DDIS / TRNDRELT )
	DGTOUSER	Incoming IP datagrams delivered to IP user protocols, per second ( HA6DDUP / TRNDRELT )
	DGBYUSER	Outgoing IP datagrams supplied by IP user protocols, per second ( HA6DSUP / TRNDRELT )
	DGOUTDSC	Outgoing IP datagrams discarded before delivery, per second ( HA6DDBD / TRNDRELT )
	DGNOROUT	Outgoing IP datagrams that had no route to their destination, per second ( HA6DNRT / TRNDRELT )
	DGFRAGRC	IP fragments received per second that required reassembly ( HA6DFRR / TRNDRELT )
	DGREASS	IP datagrams reassembled per second ( HA6DRAS / TRNDRELT )
	DGREAERR	IP datagram reassembly errors per second ( HA6DRAE / TRNDRELT )
	DGFRAGM	IP datagrams fragmented per second ( HA6DFRG / TRNDRELT )
	DGFRFAIL	IP datagram fragmentation failures per second ( HA6DFRF / TRNDRELT )
	DGFRAGCR	IP datagram fragments created per second ( HA6DFRC / TRNDRELT )
	UDPRECV	UDP datagrams received per second ( HA6URCV / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	UDPNOLST	UDP datagrams received per second for ports that had no listener ( HA6URNL / TRNDRELT )
	UDPRCERR	UDP datagrams received per second that had errors ( HA6URER / TRNDRELT )
	UDPXMIT	UDP datagrams transmitted per second ( HA6UTRN / TRNDRELT )

#### TCP/IP I/O Denial of Service Data

TCPDOSLG	DOSSMURF	Smurf denial of service packet discard/s ( HA6SMDOS / TRNDRELT )
	DOSFRAGL	Fraggle denial of service packet discard/s ( HA6FRDOS / TRNDRELT )
	DOSPODEA	Ping-o-Death den. of serv. packet discard/s ( HA6FDDOS / TRNDRELT )
	DOSBLAT	BLAT denial of service packet discard/s ( HA6BLDOS / TRNDRELT )
	DOSSTRM	Stream denial of service packet discard/s ( HA6STDOS / TRNDRELT )
	DOSR4P3D	R4P3D denial of service packet discard/s ( HA6R4DOS / TRNDRELT )
	DOSKOD	KOD denial of service packet discard/s ( HA6KDDOS / TRNDRELT )
	DOSKOX	KOX denial of service packet discard/s ( HA6KXDOS / TRNDRELT )
	DOSFLOOD	SynFlood denial of service packet discard/s ( HA6SFDOS / TRNDRELT )

#### TCP/IP I/O Activity Data

TCPIOLOG	READREQ	Read I/O requests executed per second ( HA6READ / TRNDRELT )
	WRITEREQ	Write I/O requests executed per second ( HA6WRIT / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	BYTERCV	Number of bytes received per second ( HA6BRCV / TRNDRELT )
	BYTEXMIT	Number of bytes transmitted per second ( HA6BTRN / TRNDRELT )
	QDIOIN	QDIO inbound data transfers per second ( HA6QDIN / TRNDRELT )
	QDIOOUT	QDIO outbound data transfers per second ( HA6QDOUT / TRNDRELT )
	QDIOPCI	QDIO PCI interrupts received per second ( HA6QDPCI / TRNDRELT )
	QDIOTPOL	QDIO total poll rate per second ( HA6QDPOL / TRNDRELT )
	QDIOIPOL	QDIO idle poll rate per second ( HA6QDIPO / TRNDRELT )
	IUCVRECV	IUCV receive operations per second ( HA6IURC / TRNDRELT )
	IUCVRJCT	IUCV reject operations per second ( HA6IURJ / TRNDRELT )
	IUCVREP	IUCV reply operations per second ( HA6IURP / TRNDRELT )
	IUCVSEND	IUCV send operations per second ( HA6IUSN / TRNDRELT )
	VMCFSEND	VMCF send operations per second ( HA6VMCS / TRNDRELT )
	VMCFASND	VMCF abnormal sends and send failures per second ( (HA6VMCA+HA6VMCF) / TRNDRELT )

## Performance Variables of Trend Records X'FCA8'

Trend record X'FCA8' contains data describing Reusable Server Kernel activity and performance for individual servers.

The selection format for record A8 variables is:

►► *display(useridvarname)* ►◄

### **display(**

Is the name of a performance display that contains the variable, shown under the heading 'Display' in the following table of selectable variables, and immediately followed by a left parenthesis.

### **userid**

Is the user identification of the RSK server for which data are to be extracted.

### **varname**

is the name of the selected variable, as shown under the **Variable Name** heading in the following table.

The variables shown in the following table have been pre-defined. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
<b>RSK Disk Performance Log Data</b>		
RSKDISK	AUTPERM	Authorize function permit rate per second ( HA8AUPER / TRNDRELT )
	AUTINQ	Authorize function inquire rate per second ( HA8AUINQ / TRNDRELT )
	AUTROWS	Authorize function rows retrieval rate per second ( HA8AURET / TRNDRELT )
	AUTRHIT	Authorize function row retrieval hit percentage ( HA8AUHIT * 100 / HA8AURET )
	ENRACT	Average enrollment sets active ( HA8ERROW / HA8SAMPS )
	ENRREC	Total enrollment sets records ( HA8ERREC / HA8SAMPS )
	ENRBYTE	Total enrollment sets bytes ( HA8ERBYT / HA8SAMPS )
	ENRINS	Total enrollment sets insert rate per second ( HA8ERINS / TRNDRELT )
	ENRREM	Total enrollment sets remove rate per second ( HA8ERREM / TRNDRELT )



Display	Variable Name	Variable Description and Derivation
	ENRRET	Total enrollment sets retrieve rate per second ( HA8ERRET / TRNDRELT )
	CACACT	Average caches active ( HA8CAROW / HA8SAMPS )
	CACBYTE	Total cache bytes ( HA8CABYT / HA8SAMPS )
	CACUSED	Total cache bytes used percentage ( HA8CAUSE * 100 / HA8CABYT )
	CACFILE	Total cached files ( HA8CAFIL / HA8SAMPS )
	CACOPEN	Total files opened rate per second ( HA8CAOPN / TRNDRELT )
	CACHIT	Percentage of hits out of all opens ( HA8CAHIT * 100 / HA8CAOPN )
	CACDSCRD	Total files discarded rate per second ( HA8CADSC / TRNDRELT )
<b>RSK Performance Log Data</b>		
RSKLOG	SRVACT	Average number of active services ( HA8SRROW / HA8SAMPS )
	SRVTRANS	Total services transactions per second ( HA8SRTRC / TRNDRELT )
	SRVBYREC	Total services bytes received per second ( HA8SRBRC / TRNDRELT )
	SRVBYSNT	Total services bytes sent per second ( HA8SRBSN / TRNDRELT )
	DRVACT	Average number of active line drivers ( HA8LDROW / HA8SAMPS )
	DRVTRANS	Total line drivers transactions per second ( HA8LDTRC / TRNDRELT )
	DRVBYREC	Total line drivers bytes received per second ( HA8LDBRC / TRNDRELT )

Display	Variable Name	Variable Description and Derivation
	DRVBYSNT	Total line drivers bytes sent per second ( HA8LDBSN / TRNDRELT )
	TRIACT	Average number of active tries ( HA8TRROW / HA8SAMPS )
	TRIMBYT	Total data space storage defined by tries (MB) ( HA8TRSIZ / HA8SAMPS / 1048576 )
	TRIUSED	Data space storage utilization percentage ( HA8TRUSE * 100 / HA8TRSIZ )
	TRIREC	Total records indexed ( HA8TRREC / HA8SAMPS )
	TRINODE	Total nodes ( HA8TRNOD / HA8SAMPS )
	TRILUP	Total trie lookups per second ( HA8TRLUP / TRNDRELT )
	TRIRRET	Total records returned per second ( HA8TRRRRT / TRNDRELT )
<b>RSK Storage Log Data</b>		
RSKSTOR	SGPACT	Average number of active storage groups ( HA8SGROW / HA8SAMPS )
	SGPRDIO	Total storage group page read I/Os per second ( HA8SGRDS / TRNDRELT )
	SGPPGRD	Total storage group page reads per second ( (HA8SGPRD+HA8SGPR2) / TRNDRELT )
	SGPPGRTM	Average time per page read, in msec ( HA8SGPRT / HA8SGPR2 / 4096000 )
	SGPWRIIO	Total storage group page write I/Os per second ( HA8SGWRT / TRNDRELT )
	SGPPGWR	Total storage group page writes per second ( (HA8SGPWR+HA8SGPW2) / TRNDRELT )
	SGPPGWTM	Average time per page write, in msec ( HA8SGPWT / HA8SGPW2 / 4096000 )

Display	Variable Name	Variable Description and Derivation
	MEMACT	Average number of memory rows active ( HA8MEROW / HA8SAMPS )
	MEMBYTE	Total storage size managed by memory ( (HA8MEFRE+HA8MEUSE) / HA8SAMPS )
	MEMUSED	Memory managed storage utilization ( HA8MEUSE * 100 / (HA8MEFRE+HA8MEUSE) )
	MEMALO	Total ssMemoryAllocate calls per second ( HA8MEALO / TRNDRELT )
	MEMALBYT	Total bytes taken through ssMemoryAllocate ( HA8MEALB )
	MEMREL	Total ssMemoryRelease calls per second ( HA8MEREL / TRNDRELT )
	MEMRLBYT	Total bytes returned through ssMemoryRelease ( HA8MERLB )
	MEMEXT	Total times storage extended through CMSSTOR ( HA8MEEXT )
	MEMEXBYT	Total bytes taken through CMSSTOR ( HA8MEEXB )
	MEMDEP	Total times storage depleted through CMSSTOR ( HA8MEDEP )
	MEMDPBYT	Total bytes returned through CMSSTOR ( HA8MEDEB )

## Performance Variables of Trend Records X'FCA9'

Trend record X'FCA9' contains internal summary data for Linux guest systems.

The selection format for record A9 variables is:

►► *display(useridvarname)* ►►

### **display(**

Is the name of a performance display that contains the variable, shown under the heading 'Display' in the following table of selectable variables, and immediately followed by a left parenthesis.

### **userid**

Is the user identification of the Linux guest system for which data are to be extracted.

### **varname**

is the name of the selected variable, as shown under the **Variable Name** heading in the following table.

The variables shown in the following table have been pre-defined. Their description also includes the formula used to calculate the values from the original trend record data fields.

Display	Variable Name	Variable Description and Derivation
<b>Linux CPU Utilization Data</b>		
LXCPU	CPUS	The number of virtual CPUs used ( HA9CPCNT )
	TOTCPU	Total CPU utilization (all CPUs) ( (HA9CPUSR+HA9CPSYS+HA9CPNIC) * 100 * HA9CPCNT / (HA9CPUSR +HA9CPSYS+HA9CPNIC+HA9CPIDL) )
	USER	User mode CPU util. (all CPUs) ( HA9CPUSR * 100 * HA9CPCNT / (HA9CPUSR+HA9CPSYS +HA9CPNIC+HA9CPIDL) )
	KERNEL	Kernel mode CPU util. (all CPUs) ( HA9CPSYS * 100 * HA9CPCNT / (HA9CPUSR+HA9CPSYS +HA9CPNIC+HA9CPIDL) )
	NICE	'Nice' mode CPU util. (all CPUs) ( HA9CPNIC * 100 * HA9CPCNT / (HA9CPUSR+HA9CPSYS +HA9CPNIC+HA9CPIDL) )
	IDLE	CPU percent idle (all CPUs) ( HA9CPIDL * 100 * HA9CPCNT / (HA9CPUSR+HA9CPSYS +HA9CPNIC+HA9CPIDL) )

Display	Variable Name	Variable Description and Derivation
	RUNABL	Number of runnable processes ( HA9PROCR / HA9SAMPC )
	TOTAL	Total number of processes ( HA9PROCT / HA9SAMPC )
	AVPROC1	Average processes running last minute ( HA9AVR1 / 2048 )
	AVPROC5	Average proc. running last 5 minutes ( HA9AVR5 / 2048 )
	AVPROC15	Average proc. running last 15 minutes ( HA9AVR15 / 2048 )
<b>Linux Memory Data</b>		
LXMEM	M_TOTAL	Total main memory size in MB ( HA9TOTLM / HA9SAMPM / 1024 )
	%MUSED	Percent main memory in use ( (HA9TOTLM-HA9FREEM) * 100 / HA9TOTLM )
	H_TOTAL	Total high memory size in MB ( HA9TOTLH / HA9SAMPM / 1024 )
	%HUSED	Percent high memory in use ( (HA9TOTLH-HA9FREEH) * 100 / HA9TOTLH )
	SHARED	Shared memory size in MB ( HA9SHARD / HA9SAMPM / 1024 )
	BUFFERS	Memory reserved for buffers and free cache ( HA9TOTLC / HA9SAMPM / 1024 )
	S_TOTAL	Total swap area size in MB ( HA9TOTLS / HA9SAMPM / 1024 )
	%SUSED	Percent swap area in use ( (HA9TOTLS-HA9FREES) * 100 / HA9TOTLS )
	CACHE	Memory used for cache ( HA9USEDK / HA9SAMPM / 1024 )

Display	Variable Name	Variable Description and Derivation
	SWAPIN	Swap in rate, in 4kB pages per second ( HA9SWIN / TRNDRELT )
	SWAPOUT	Swap out rate, in 4kB pages per second ( HA9SWOUT / TRNDRELT )
	BLKREAD	Block I/O data read rate, in kB per second ( HA9BLKRD / TRNDRELT )
	BLKWRITE	Block I/O data write rate, in kB per second ( HA9BLKWR / TRNDRELT )

#### Linux Networking Activity Data

LXNETWRK	INTFACES	Average number of network interfaces ( HA9INTF / HA9SAMPN )
	PACKRECV	Packets received per second ( HA9PKRCV / TRNDRELT )
	PACKXMIT	Packets transmitted per second ( HA9PKSNT / TRNDRELT )
	BYTERECV	Bytes received per second ( HA9BYRCV / TRNDRELT )
	BYTEXMIT	Bytes transmitted per second ( HA9BYSNT / TRNDRELT )
	BADPACK	Bad packets received per second ( HA9PKBAD / TRNDRELT )
	XMITERR	Packet transmit errors per second ( HA9PKERR / TRNDRELT )
	NOBUFFER	No space found in buffers per second ( HA9NOSPB / TRNDRELT )
	NOSPACE	No space found in Linux per second ( HA9NOSPL / TRNDRELT )
	COLLISN	Transmit collisions per second ( HA9COLSN / TRNDRELT )

## Notices

---

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan Ltd.  
19-21, Nihonbashi-Hakozakicho, Chuo-ku  
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

#### COPYRIGHT LICENSE:

This information may contain sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

## IBM Online Privacy Statement

---

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see:

- The section entitled **IBM Websites** at [IBM Privacy Statement](https://www.ibm.com/privacy) (<https://www.ibm.com/privacy>)
- [Cookies and Similar Technologies](https://www.ibm.com/privacy#Cookies_and_Similar_Technologies) ([https://www.ibm.com/privacy#Cookies\\_and\\_Similar\\_Technologies](https://www.ibm.com/privacy#Cookies_and_Similar_Technologies))

## Trademarks

---

IBM, the IBM logo, and [ibm.com](https://www.ibm.com)® are trademarks or registered trademarks of International Business Machines Corp., in the United States and/or other countries. Other product and service names might be



trademarks of IBM or other companies. A current list of IBM trademarks is available on [IBM Copyright and trademark information](https://www.ibm.com/legal/copytrade) (<https://www.ibm.com/legal/copytrade>).

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a world-wide basis.



# Bibliography

---

This topic lists the publications in the z/VM library. For abstracts of the z/VM publications, see [z/VM: General Information](#).

## Where to Get z/VM Information

---

The current z/VM product documentation is available in [IBM Documentation - z/VM \(https://www.ibm.com/docs/en/zvm\)](https://www.ibm.com/docs/en/zvm).

## z/VM Base Library

---

### Overview

- [z/VM: License Information](#), GI13-4377
- [z/VM: General Information](#), GC24-6286

### Installation, Migration, and Service

- [z/VM: Installation Guide](#), GC24-6292
- [z/VM: Migration Guide](#), GC24-6294
- [z/VM: Service Guide](#), GC24-6325
- [z/VM: VMSES/E Introduction and Reference](#), GC24-6336

### Planning and Administration

- [z/VM: CMS File Pool Planning, Administration, and Operation](#), SC24-6261
- [z/VM: CMS Planning and Administration](#), SC24-6264
- [z/VM: Connectivity](#), SC24-6267
- [z/VM: CP Planning and Administration](#), SC24-6271
- [z/VM: Getting Started with Linux on IBM Z](#), SC24-6287
- [z/VM: Group Control System](#), SC24-6289
- [z/VM: I/O Configuration](#), SC24-6291
- [z/VM: Running Guest Operating Systems](#), SC24-6321
- [z/VM: Saved Segments Planning and Administration](#), SC24-6322
- [z/VM: Secure Configuration Guide](#), SC24-6323

### Customization and Tuning

- [z/VM: CP Exit Customization](#), SC24-6269
- [z/VM: Performance](#), SC24-6301

### Operation and Use

- [z/VM: CMS Commands and Utilities Reference](#), SC24-6260
- [z/VM: CMS Primer](#), SC24-6265
- [z/VM: CMS User's Guide](#), SC24-6266
- [z/VM: CP Commands and Utilities Reference](#), SC24-6268

- [z/VM: System Operation](#), SC24-6326
- [z/VM: Virtual Machine Operation](#), SC24-6334
- [z/VM: XEDIT Commands and Macros Reference](#), SC24-6337
- [z/VM: XEDIT User's Guide](#), SC24-6338

## Application Programming

- [z/VM: CMS Application Development Guide](#), SC24-6256
- [z/VM: CMS Application Development Guide for Assembler](#), SC24-6257
- [z/VM: CMS Application Multitasking](#), SC24-6258
- [z/VM: CMS Callable Services Reference](#), SC24-6259
- [z/VM: CMS Macros and Functions Reference](#), SC24-6262
- [z/VM: CMS Pipelines User's Guide and Reference](#), SC24-6252
- [z/VM: CP Programming Services](#), SC24-6272
- [z/VM: CPI Communications User's Guide](#), SC24-6273
- [z/VM: ESA/XC Principles of Operation](#), SC24-6285
- [z/VM: Language Environment User's Guide](#), SC24-6293
- [z/VM: OpenExtensions Advanced Application Programming Tools](#), SC24-6295
- [z/VM: OpenExtensions Callable Services Reference](#), SC24-6296
- [z/VM: OpenExtensions Commands Reference](#), SC24-6297
- [z/VM: OpenExtensions POSIX Conformance Document](#), GC24-6298
- [z/VM: OpenExtensions User's Guide](#), SC24-6299
- [z/VM: Program Management Binder for CMS](#), SC24-6304
- [z/VM: Reusable Server Kernel Programmer's Guide and Reference](#), SC24-6313
- [z/VM: REXX/VM Reference](#), SC24-6314
- [z/VM: REXX/VM User's Guide](#), SC24-6315
- [z/VM: Systems Management Application Programming](#), SC24-6327
- [z/VM: z/Architecture Extended Configuration \(z/XC\) Principles of Operation](#), SC27-4940

## Diagnosis

- [z/VM: CMS and REXX/VM Messages and Codes](#), GC24-6255
- [z/VM: CP Messages and Codes](#), GC24-6270
- [z/VM: Diagnosis Guide](#), GC24-6280
- [z/VM: Dump Viewing Facility](#), GC24-6284
- [z/VM: Other Components Messages and Codes](#), GC24-6300
- [z/VM: VM Dump Tool](#), GC24-6335

## z/VM Facilities and Features

---

### Data Facility Storage Management Subsystem for z/VM

- [z/VM: DFSMS/VM Customization](#), SC24-6274
- [z/VM: DFSMS/VM Diagnosis Guide](#), GC24-6275
- [z/VM: DFSMS/VM Messages and Codes](#), GC24-6276
- [z/VM: DFSMS/VM Planning Guide](#), SC24-6277

- *z/VM: DFSMS/VM Removable Media Services*, SC24-6278
- *z/VM: DFSMS/VM Storage Administration*, SC24-6279

## Directory Maintenance Facility for z/VM

- *z/VM: Directory Maintenance Facility Commands Reference*, SC24-6281
- *z/VM: Directory Maintenance Facility Messages*, GC24-6282
- *z/VM: Directory Maintenance Facility Tailoring and Administration Guide*, SC24-6283

## Open Systems Adapter

- Open Systems Adapter/Support Facility on the Hardware Management Console ([https://www.ibm.com/docs/en/SSLTBW\\_2.3.0/pdf/SC14-7580-02.pdf](https://www.ibm.com/docs/en/SSLTBW_2.3.0/pdf/SC14-7580-02.pdf)), SC14-7580
- Open Systems Adapter-Express ICC 3215 Support (<https://www.ibm.com/docs/en/zos/2.3.0?topic=osa-icc-3215-support>), SA23-2247
- Open Systems Adapter Integrated Console Controller User's Guide ([https://www.ibm.com/docs/en/SSLTBW\\_2.3.0/pdf/SC27-9003-02.pdf](https://www.ibm.com/docs/en/SSLTBW_2.3.0/pdf/SC27-9003-02.pdf)), SC27-9003
- Open Systems Adapter-Express Customer's Guide and Reference ([https://www.ibm.com/docs/en/SSLTBW\\_2.3.0/pdf/iaa2z1f0.pdf](https://www.ibm.com/docs/en/SSLTBW_2.3.0/pdf/iaa2z1f0.pdf)), SA22-7935

## Performance Toolkit for z/VM

- *z/VM: Performance Toolkit Guide*, SC24-6302
- *z/VM: Performance Toolkit Reference*, SC24-6303

The following publications contain sections that provide information about z/VM Performance Data Pump, which is licensed with Performance Toolkit for z/VM.

- *z/VM: Performance*, SC24-6301. See *z/VM Performance Data Pump*.
- *z/VM: Other Components Messages and Codes*, GC24-6300. See *Data Pump Messages*.

## RACF® Security Server for z/VM

- *z/VM: RACF Security Server Auditor's Guide*, SC24-6305
- *z/VM: RACF Security Server Command Language Reference*, SC24-6306
- *z/VM: RACF Security Server Diagnosis Guide*, GC24-6307
- *z/VM: RACF Security Server General User's Guide*, SC24-6308
- *z/VM: RACF Security Server Macros and Interfaces*, SC24-6309
- *z/VM: RACF Security Server Messages and Codes*, GC24-6310
- *z/VM: RACF Security Server Security Administrator's Guide*, SC24-6311
- *z/VM: RACF Security Server System Programmer's Guide*, SC24-6312
- *z/VM: Security Server RACROUTE Macro Reference*, SC24-6324

## Remote Spooling Communications Subsystem Networking for z/VM

- *z/VM: RSCS Networking Diagnosis*, GC24-6316
- *z/VM: RSCS Networking Exit Customization*, SC24-6317
- *z/VM: RSCS Networking Messages and Codes*, GC24-6318
- *z/VM: RSCS Networking Operation and Use*, SC24-6319
- *z/VM: RSCS Networking Planning and Configuration*, SC24-6320

## TCP/IP for z/VM

- [\*z/VM: TCP/IP Diagnosis Guide\*](#), GC24-6328
- [\*z/VM: TCP/IP LDAP Administration Guide\*](#), SC24-6329
- [\*z/VM: TCP/IP Messages and Codes\*](#), GC24-6330
- [\*z/VM: TCP/IP Planning and Customization\*](#), SC24-6331
- [\*z/VM: TCP/IP Programmer's Reference\*](#), SC24-6332
- [\*z/VM: TCP/IP User's Guide\*](#), SC24-6333

## Prerequisite Products

---

### Device Support Facilities

- Device Support Facilities (ICKDSF): User's Guide and Reference ([https://www.ibm.com/docs/en/SSLTBW\\_3.1.0/pdf/ickug00\\_v3r1.pdf](https://www.ibm.com/docs/en/SSLTBW_3.1.0/pdf/ickug00_v3r1.pdf)), GC35-0033

## Related Products

---

### XL C<sup>®</sup> ++ for z/VM

- [\*XL C/C++ for z/VM: Runtime Library Reference\*](#), SC09-7624
- [\*XL C/C++ for z/VM: User's Guide\*](#), SC09-7625

### z/OS<sup>®</sup>

IBM Documentation - z/OS (<https://www.ibm.com/docs/en/zos>)

# Index

## Special Characters

? (retrieve subcommand) [5](#)  
\* (comment) [3](#)  
#CP subcommand [103](#)  
= (re-execute subcommand) [4](#)

## A

active users [375](#), [380](#)  
ACTMSG argument of FControl cmd [13](#)  
Age List Log screen [809](#)  
AGELlog subcommand [130](#)  
allocation  
    pages allocated on disk (user) [436](#), [445](#)  
AUTOREFR argument of FControl cmd [14](#)  
auxiliary storage log screen [509](#)  
AUXlog subcommand [131](#)  
Available List Data Above 2G screen [806](#)  
Available List Data Below 2G screen [804](#)  
available list frame request [382](#)  
available list log screen [728](#)  
average  
    users in queue [374](#)  
AVERAGE subcommand [133](#)  
AVLA2glg subcommand [134](#)  
AVLB2glg subcommand [135](#)

## B

BACKWARD subcommand  
    perf. monitor mode [136](#)  
    redisplay mode [120](#)  
benchmark log displays  
    cache extended functions log [587](#)  
    control unit cache log [552](#)  
    CPOWNED device log [553](#)  
    DASD seeks distances log (per device) [557](#)  
    DASD seeks locations log (per device) [556](#)  
    general I/O device data log (per device) [548](#)  
    Linux CPU load log (per user) [718](#)  
    Linux memory utilization log (per user) [719](#)  
    Linux network activity log [720](#)  
    QDIO activity log (per device) [725](#)  
    user communication log (per user) [547](#)  
    user page data log (per user) [543](#)  
    user resources per transaction log (per user) [546](#)  
    user response time log (per user) [545](#)  
    user wait states log (per user) [544](#)  
benchmark log reports  
    user resource usage log (per user) [542](#)  
benchmark log selection menu [558](#)  
BENCHMRK argument of FControl cmd [15](#)  
BFS file activity log [599](#)  
BFS locking activity log [604](#)  
BFS PIPEs and linking activity log [601](#)  
BOTTOM subcommand

BOTTOM subcommand (*continued*)  
    perf. monitor mode [142](#)  
    redisplay mode [121](#)  
busy  
    channel busy [410](#), [541](#), [658](#)  
    device busy [423](#)  
    minidisk busy [425](#), [453](#)

## C

CACHdbse subcommand [143](#)  
CACHDLog subcommand [144](#)  
cache extended functions  
    performance monitor displays  
        cache extended functions [569](#)  
    screen identifiers  
        FCX177, cache extended functions [569](#)  
cache extended functions log (per device) [587](#)  
cached control units screen [566](#)  
CACHLog subcommand [145](#)  
CACHExt subcommand [146](#)  
channel load screen [410](#), [657](#), [691](#)  
CHannel subcommand [147](#)  
CLEAR subcommand [101](#)  
CMS subcommand [102](#)  
color and highlighting [18](#)  
COLOR argument of FControl cmd [18](#)  
command search order [79](#)  
commands [1](#)  
commands, basic mode  
    'ENTER' [118](#)  
    #CP [103](#)  
    CLear [101](#)  
    CMS [102](#)  
    CP [103](#)  
    Delete [104](#)  
    FCONAppc [105](#)  
    FCONRmt [106](#)  
    HISTdata [109](#)  
    MONitor [110](#)  
    MONScan [111](#)  
    QUIT [113](#)  
    REDisp [114](#)  
    Reply [115](#)  
    RETurn [116](#)  
    TRNDScan [117](#)  
commands, general  
    ? (retrieve subcommand) [5](#)  
    \* (comment) [3](#)  
    = (re-execute subcommand) [4](#)  
    AVAILLOG [132](#)  
    BASic [6](#)  
    CANCEl [7](#)  
    DEMNDLOG [167](#)  
    FControl ...  
        FControl ACTmsg [13](#)  
        FControl AUTOREfr [14](#)

commands, general (*continued*)

FControl ... (*continued*)

FControl BENCHmrk [15](#)  
FControl Color [18](#)  
FControl DEFLog [21](#)  
FControl DEFScrN [24](#)  
FControl DEVGROUp [26](#)  
FControl DUMPID [28](#)  
FControl EMERGENc [29](#)  
FControl FCXLevel [32](#)  
FControl FORCeusr [33](#)  
FControl GDDMspeC [38](#)  
FControl HMA [40](#)  
FControl LImit [42](#)  
FControl MAINTid [46](#)  
FControl Mxrec [47](#)  
FControl MINPaths [48](#)  
FControl Model [50](#)  
FControl MONCAche [51](#)  
FControl MONcoll [53](#)  
FControl MSGClear [63](#)  
FControl MSGWait [64](#)  
FControl NUmbers [65](#)  
FControl PFkey [66](#)  
FControl PRocess [70](#)  
FControl RELoad [75](#)  
FControl RMTlines [76](#)  
FControl SCroll [77](#)  
FControl SEarch [79](#)  
FControl SECuser [80](#)  
FControl SETEvent [81](#)  
FControl SETtings [83](#)  
FControl TIMStamp [88](#)  
FControl UCLass [89](#)  
FControl UPdtcms [90](#)  
FControl UPDTScrN [91](#)  
FControl USerbotl [92](#)  
FControl USERHdr [93](#)  
FControl USERVar [94](#)  
FControl USRlimit [97](#)  
MODLEVEL [108](#)

HELP [8](#)

QEBSM [266](#)  
QEBSMLOG [267](#)  
STORLOG [303](#)  
SXSAVAIL [306](#)  
SXsDEFER [307](#)  
SXSPAGE [308](#)  
SXsUTIL [309](#)  
UQDIO [342](#)  
UQDIOLOG [343](#)

commands, performance monitor

/... (locate) [364](#)

AGELlog [130](#)  
AUXlog [131](#)  
Average [133](#)  
AVLA2glg [134](#)  
AVLB2glg [135](#)  
Backward [136](#)  
BASMode [137](#)  
BENchmrk [138](#)  
BFSFiles [139](#)  
BFSLocks [140](#)  
BFSPipes [141](#)

commands, performance monitor (*continued*)

Bottom [142](#)  
CACHdbse [143](#)  
CACHDLog [144](#)  
CACHExt [145](#)  
CACHExt [146](#)  
CHannel [147](#)  
CMS [148](#)  
CORrel [149](#)  
CP [152](#)  
CPLACT [153](#)  
CPLCONF [154](#)  
CPLMENU [155](#)  
CPOWNLog [156](#)  
CPSAll [157](#)  
CPServ [158](#)  
CPSFails [159](#)  
CPSTOCp [160](#)  
CPSTovm [161](#)  
CPU [162](#)  
CPUMENU [163](#)  
CTLunit [164](#)  
Current [165](#)  
DASDlog [166](#)  
DEVConf [168](#)  
DEvIce [169](#)  
DEVLog [170](#)  
DEVMenu [171](#)  
Down [180](#)  
DSPACESh [172](#)  
DSVBKact [173](#)  
DSVClog [174](#)  
DSVSLog [175](#)  
ENTER key [365](#)  
EVNIC [176](#)  
EVSWITCH [177](#)  
EXcept [178](#)  
FCHannel [179](#)  
Forward [180](#)  
GLONact [181](#)  
GRAPHdet [182](#)  
GRAPHIcs [184](#)  
GRAPHSum [185](#)  
GRAPHVAr [189](#)  
GVNIC [192](#)  
GVSWITCH [193](#)  
HIPSock [194](#)  
HISTdata [195](#)  
HPALIAS [196](#)  
HPFLog [197](#)  
HPSHARE [198](#)  
I/O [169](#)  
IDLEuser [199](#)  
INTerim [200](#)  
IOChange [201](#)  
IOPqueue [202](#)  
IOPRoclg [203](#)  
ISFEAct [204](#)  
ISFEConf [205](#)  
ISFELog [206](#)  
ISFLAct [207](#)  
ISFLALog [208](#)  
ISFLConf [209](#)  
ISFLLog [210](#)



commands, performance monitor (*continued*)

[LCHannel 211](#)  
[Left 212](#)  
[LGRData 213](#)  
[LGRElog 214](#)  
[LINux 215](#)  
[LINUXMEN 216](#)  
[LOCate 217](#)  
[LOCKact 218](#)  
[LOCKLog 219](#)  
[LPAr 220](#)  
[LPARConf 221](#)  
[LPARLog 222](#)  
[LPARLOGM 223](#)  
[LPARMenu 224](#)  
[LSHaract Menu 225](#)  
[LXCPU 226](#)  
[LXCPULog 227](#)  
[LXFILsys 228](#)  
[LXMEM 229](#)  
[LXMEMLog 230](#)  
[LXNETLog 231](#)  
[LXNETwrk 232](#)  
[MDCache 233](#)  
[MDCStor 234](#)  
[MENu 235](#)  
[MONData 236](#)  
[MONSEt 237](#)  
[MTCORElg 238](#)  
[MTCTYPlg 239](#)  
[MTUsers 240](#)  
[MTUSRLog 241](#)  
[Next 180](#)  
[NEXTSamp 242](#)  
[NSS 243](#)  
[PAGelog 244](#)  
[PCIACT 245](#)  
[PCICONF 246](#)  
[PCILOG 247](#)  
[PCIMENU 248](#)  
[PHYSLog 249](#)  
[PINLOG 250](#)  
[PLOTdet 182](#)  
[PLOTSum 185](#)  
[PLOTVar 189](#)  
[PRCLog 251](#)  
[Print 252](#)  
[PRINTRmt 255](#)  
[PRIVop 257](#)  
[PROCCConf 258](#)  
[PROClog 259](#)  
[PROCmenu 260](#)  
[PROCSum 261](#)  
[PUCfglog 262](#)  
[PUOrglog 263](#)  
[QDIO 264](#)  
[QDIOLOG 265](#)  
[REDHist 268](#)  
[REDisp 269](#)  
[RESet 270](#)  
[RESPall 271](#)  
[RESUp 272](#)  
[RETURN 273](#)  
[RIght 274](#)

commands, performance monitor (*continued*)

[RSKDet 275](#)  
[RSKDisK 276](#)  
[RSKLog 277](#)  
[RSKmenu 278](#)  
[RSKStor 279](#)  
[SCHDATA 280](#)  
[SCHedlog 281](#)  
[SCSI 282](#)  
[SCSILog 283](#)  
[SEEKDEt 284](#)  
[SEEKdist 285](#)  
[SEEKDLog 286](#)  
[SEEKLLog 287](#)  
[SEEKLoc 288](#)  
[SFS 289](#)  
[SFSIOLOG 290](#)  
[SFSLog 291](#)  
[SFSMenu 292](#)  
[SFSReq 293](#)  
[SKIPsmp 294](#)  
[SOrt 295](#)  
[SPool 296](#)  
[SSICnf 297](#)  
[SSIMenu 298](#)  
[SSISChlg 299](#)  
[SSISMilg 300](#)  
[STEALlog 301](#)  
[STOrage 302](#)  
[STORMenu 304](#)  
[SUBPLOG 305](#)  
[SYSConf 310](#)  
[SYSLOAD 311](#)  
[SYSLog 312](#)  
[SYSMENU 313](#)  
[SYSSEt 314](#)  
[SYSSUmlg 315](#)  
[SYStem 316](#)  
[SYSTRans 317](#)  
[TCPActlg 318](#)  
[TCPBplog 319](#)  
[TCPConf 320](#)  
[TCPDatlg 321, 322](#)  
[TCPICmp 323](#)  
[TCPIolog 324](#)  
[TCPLinks 325](#)  
[TCPMenu 326](#)  
[TCPSess 327](#)  
[TCPUsers 328](#)  
[Top 329](#)  
[TOPolog 330](#)  
[UCOmm 331](#)  
[UCOMMLog 332](#)  
[UCONf 333](#)  
[UDEFMenu 334](#)  
[Up 136](#)  
[UPAge 335](#)  
[UPAGELog 336](#)  
[UPGAct 337](#)  
[UPGACTLg 338](#)  
[UPGMenu 339](#)  
[UPGUtl 340](#)  
[UPGUTLLg 341](#)  
[UREsp 344](#)

commands, performance monitor (*continued*)

- [URESPLog 345](#)
- [USER 346](#)
- [USERLog 347](#)
- [USRMPLOG 348](#)
- [USRPrctm 349](#)
- [USRTmlog 350](#)
- [USTat 351](#)
- [USTATG 351](#)
- [USTATLog 352](#)
- [USTLog 353](#)
- [USTMPLOG 354](#)
- [UTRANDet 355](#)
- [UTRANLog 356](#)
- [UTRans 357](#)
- [VDIsks 358](#)
- [VMRM 359](#)
- [VNIC 360](#)
- [VOLLOG 361](#)
- [VOLUME 362](#)
- [VSWITCH 363](#)
- commands, redisplay
  - ['ENTER'-key 128](#)
  - [/... \(locate\) 127](#)
  - [Backward 120](#)
  - [Bottom 121](#)
  - [Down 122](#)
  - [Forward 122](#)
  - [Left 123](#)
  - [Next 122](#)
  - [RETURN 124](#)
  - [Rlght 125](#)
  - [Top 126](#)
  - [Up 120](#)
- configuration
  - [I/O configuration changes screen 590](#)
  - [I/O device configuration screen 482](#)
- [connect time, I/O device 414, 423, 550, 567](#)
- [connect time, volume 866, 869](#)
- [console function wait 440, 448, 492](#)
- [Control unit cache log 552](#)
- [control unit queuing 414, 423, 550](#)
- [correlation coefficients 149, 607](#)
- [CP IUCV services log screens 500](#)
- [CP IUCV services screen 408](#)
- [CP subcommand 103](#)
- [CP-owned device screen 416](#)
- [CPLACT subcommand 153](#)
- [CPLCONF subcommand 154](#)
- [CPLMENU subcommand 155](#)
- [CPOWNED device log \(per device\) 553](#)
- [CPOWNLog subcommand 156](#)
- [CPSAll subcommand 157](#)
- [CPSERV subcommand 158](#)
- [CPSFails subcommand 159](#)
- [CPSTOCp subcommand 160](#)
- [CPSTovm subcommand 161](#)
- [CPU activity selection screen 858](#)
- [CPU Load and Transactions report 371](#)
- [CPU pool activity screen 835](#)
- [CPU pool configuration screen 833](#)
- [CPU pool menu screen 857](#)
- [CPU subcommand 162](#)

- [CPUMENU subcommand 163](#)
- [CTLunit subcommand 164](#)
- [CU-cached disks screen 425, 428](#)
- [CUMUL argument of GRAPHxxx commands 183, 186, 190, 968, 969](#)
- [cumulative graphic output 968, 969](#)
- [CURRENT subcommand 165](#)

## D

- [DASD performance log 575, 586](#)
- [DASD seeks details screen 536](#)
- [DASD seeks distances log \(per device\) 557](#)
- [DASD seeks distances screen 534](#)
- [DASD seeks locations log \(per device\) 556](#)
- [DASD seeks locations screen 532](#)
- [DASD volume performance log screen 868](#)
- [DASD volume performance screen 865](#)
- [DASDlog subcommand 166](#)
- [data completeness 370](#)
- [data consistency 369](#)
- [data field values 369](#)
- [data space activity 444](#)
- [DEFLOG argument of FControl cmd 21](#)
- [DEFSCRN argument of FControl cmd 24](#)
- [DELETE subcommand 104](#)
- [demand scan 382, 388](#)
- [demand scan log screen 736](#)
- [detailed shared file system screen 457](#)
- detection of
  - [missing I/O interrupts \(real\) 415](#)
  - [missing I/O interrupts \(virt.\) 454](#)
- [DEVConf subcommand 168](#)
- [DEVGROU argument of FControl cmd 26](#)
- [DEVICE subcommand 169](#)
- [DEVLog subcommand 170](#)
- [DEVMenu subcommand 171](#)
- [diagnose instructions 373, 379, 400](#)
- [disconnect time, I/O device 414, 423, 550, 567](#)
- [disconnect time, volume 866, 869](#)
- [Dispatch Vector Activity screen 818](#)
- [Dispatch Vector Configuration Change Log screen 816](#)
- [DOWN subcommand
  - \[perf. monitor mode 180\]\(#\)
  - \[redisplay mode 122\]\(#\)](#)
- [DSPACESh subcommand 172](#)
- [DSVBK Steals per Processor Log screen 822](#)
- [DSVBKact subcommand 173](#)
- [DSVClog subcommand 174](#)
- [DSVSLog subcommand 175](#)
- [DUMPID argument of FControl cmd 28](#)

## E

- [EMERGENC argument of FControl cmd 29](#)
- [ENTER key
  - \[operation in basic mode 118\]\(#\)
  - \[performance monitor mode 365\]\(#\)
  - \[redisplay mode 128\]\(#\)](#)
- [entering commands
  - \[Performance Toolkit subcommands from EXECs 1\]\(#\)](#)
- [error messages 879](#)

- EVNIC subcommand [176](#)
- EVSWITCH activity screen [753](#)
- EVSWITCH subcommand [177](#)
- EXCEPT subcommand [178](#)
- exception log display [524](#)
- extended history data log display [609](#)
- extended summary files [985](#)
- extended trend files [983](#)
- Extended Virtual Network Device Activity Screen [759](#)

## F

- FCHANNEL subcommand [179](#)
- FCONAPPC subcommand [105](#)
- FCONCMD [1](#)
- FCONRMT AUTHORIZ file
  - dynamically refreshing [75](#)
- FCONRMT menu, initial [629](#)
- FCONRMT subcommand [106](#)
- FCONRMT SYSTEMS file
  - dynamically refreshing [75](#)
- FCONX LINUXUSR file
  - dynamically refreshing [75](#)
- FCONX REPORTS file [252](#)
- FCONXREP summary history files screen [631](#), [972](#)
- FCXLEVEL argument of FControl cmd [32](#)
- features
  - display terminals [422](#)
- FICON channel load screen [657](#)
- field values [369](#)
- FORCEUSR argument of FControl cmd [33](#)
- FORWARD subcommand
  - perf. monitor mode [180](#)
  - redisplay mode [122](#)
- frame requests
  - from available list [382](#)
  - from XSTORE [383](#), [389](#)
- function key definitions [369](#)
- function key settings
  - initial, in performance monitor mode [367](#)
- function pending time [414](#), [423](#), [550](#), [567](#), [866](#), [869](#)

## G

- GDDM [956](#)
- GDDMSPEC argument of FControl cmd [38](#)
- general
  - EVSWITCH activity screen [753](#)
  - Extended Virtual Network Device Activity Screen [759](#)
  - General Virtual Network Device Description Screen [755](#)
  - General Virtual Switch Description Screen [750](#)
  - HPF I/O Device Performance Log Screen [782](#)
  - HPF I/O Device Screen [778](#)
  - I/O Device Data Selection Menu [785](#)
  - I/O device screen [412](#)
  - ISFC End Point Activity [764](#)
  - ISFC End Point Activity Log [775](#)
  - ISFC End Point Configuration [761](#)
  - ISFC Logical Link Activity [766](#)
  - ISFC Logical Link Activity By-Time Log [777](#)
  - ISFC Logical Link Activity Log [776](#)
  - ISFC Logical Link Configuration [768](#)
  - QDIO activity screen [723](#)

- general (*continued*)
  - seeks details screen [536](#)
  - seeks distances screen [534](#)
  - seeks locations screen [532](#)
  - SSI Configuration [770](#)
  - SSI Data Menu [760](#)
  - SSI State Change Synchronization Activity Log [772](#)
  - SSI State/Mode Information Log [773](#)
  - user report [431](#)
  - Virtual Network Device Activity Screen [757](#)
  - VM Resource Manager screen [709](#)
  - VSWITCH activity screen [707](#)
- general I/O device data log (per device) [548](#)
- general Linux CPU utilization [712](#)
- general Linux memory utilization [714](#)
- general Linux networking activity [716](#)
- general shared file system screen [455](#)
- general TCP/IP activity log screen [638](#)
- general TCP/IP data transfer log screen [640](#)
- General Virtual Network Device Description Screen [755](#)
- General Virtual Switch Description Screen [750](#)
- global networking object activity screen [849](#)
- GLONACT subcommand [181](#)
- GRAPHDET subcommand [182](#)
- graphical
  - daily performance profile [965](#)
  - history data display [955](#)
  - history data displays [955](#)
  - performance trend display [964](#)
  - user wait state screen [441](#)
  - variable correlation plots [966](#)
- graphical history data display [955](#)
- GRAPHICS subcommand [184](#)
- GRAPHSUM subcommand [185](#)
- GRAPHVAR subcommand [189](#)
- GVNIC subcommand [192](#)
- GVSITCH subcommand [193](#)

## H

- HELP subcommand [8](#)
- highlighting and color [18](#)
- Hipersocket channel activity screen [691](#)
- HIPSOCK subcommand [194](#)
- HISTDATA subcommand [195](#)
- HISTLOG file records [981](#)
- history data display [269](#)
- history data display, extended [268](#)
- history data displays [955](#)
- history data selection menu [539](#)
- history files
  - daily perf. profiles based on [956](#), [965](#)
  - perf. graphics based on history files [955](#), [964](#)
  - records [981](#)
  - trend graphics based on [956](#)
  - variable correl. graphics [966](#)
- HISTSUM file records [981](#)
- HMA argument of FControl cmd [40](#)
- HPALIAS subcommand [196](#)
- HPF I/O Device Performance Log Screen [782](#)
- HPF I/O Device Screen [778](#)
- HPFLog subcommand [197](#)
- HPSHARE subcommand [198](#)
- HyperPAV alias activity screen [861](#)

HyperPAV share activity screen [863](#)

## I

I/O assist [454](#)  
I/O configuration changes screen [590](#)  
I/O contention  
    device level [423](#)  
I/O device configuration screen [482](#)  
I/O Device Data Selection Menu [785](#)  
I/O device details screen [420](#)  
I/O device screen [412](#)  
I/O interrupts, missing [415](#)  
I/O priority queueing screen [702](#)  
I/O Processor Log screen [693](#)  
I/O rate  
    minidisks [426](#), [451](#)  
    real devices [413](#), [421](#), [429](#), [549](#), [571](#), [589](#)  
    system areas [427](#)  
I/O subcommand [169](#)  
I/O wait  
    system [372](#)  
    user [439](#), [448](#), [491](#)  
idle users screen [704](#)  
IDLEuser subcommand [199](#)  
initial FCONRMT menu [629](#)  
instruction wait [439](#), [448](#), [492](#)  
INTerim subcommand [200](#)  
interrupts, missing [415](#)  
interval [367](#)  
IOChange subcommand [201](#)  
IOPqueue subcommand [202](#)  
IOPRoclg subcommand [203](#)  
IORBKs queued [423](#), [447](#)  
ISFC End Point Activity [764](#)  
ISFC End Point Activity Log [775](#)  
ISFC End Point Configuration [761](#)  
ISFC Logical Link Activity [766](#)  
ISFC Logical Link Activity Log [776](#), [777](#)  
ISFC Logical Link Configuration [768](#)  
ISFEAct subcommand [204](#)  
ISFEConf subcommand [205](#)  
ISFELog subcommand [206](#)  
ISFLAct subcommand [207](#)  
ISFLALog subcommand [208](#)  
ISFLConf subcommand [209](#)  
ISFLLog subcommand [210](#)  
IUCV soft wait [440](#), [492](#)

## L

LCHANNEL subcommand [211](#)  
LEFT subcommand  
    perf. monitor mode [212](#)  
    redisplay mode [123](#)  
LGR Data screen [789](#)  
LGR Event Log screen [786](#)  
LGRData subcommand [213](#)  
LGRELog subcommand [214](#)  
LIMIT argument of FControl cmd [42](#)  
Linux CPU load log (per user) [718](#)  
Linux CPU utilization screen [689](#)  
Linux data selection menu [711](#)

Linux details selection screen [675](#)  
Linux filesystem usage [684](#)  
Linux memory utilization and activity [686](#)  
Linux memory utilization log (per user) [719](#)  
Linux network activity log (per user) [720](#)  
Linux Network Activity screen [682](#)  
LINUX subcommand [215](#)  
Linux systems selection screen [674](#)  
LINUXMEN subcommand [216](#)  
locate subcommand  
    perf. monitor mode [364](#)  
    redisplay mode [127](#)  
LOCATE subcommand [217](#)  
LOCKact subcommand [218](#)  
LOCKLog subcommand [219](#)  
Logical Core Organization Log screen [811](#)  
Logical Partition Activity menu [828](#)  
Logical Partition Logs menu [832](#)  
Logical Partition Share screen [829](#)  
LPAR channel load screen [540](#)  
LPAR Load Log report [632](#)  
LPAR Load screen [472](#)  
LPAR subcommand [220](#)  
LPARConf subcommand [221](#)  
LPARLog subcommand [222](#)  
LPARLOGM subcommand [223](#)  
LPARMenu subcommand [224](#)  
LSHaract subcommand [225](#)  
LXCPU subcommand [226](#)  
LXCPULOG subcommand [227](#)  
LXFILSYS subcommand [228](#)  
LXMEM subcommand [229](#)  
LXMEMLOG subcommand [230](#)  
LXNETLOG subcommand [231](#)  
LXNETWRK subcommand [232](#)

## M

MAINTID argument of FControl cmd [46](#)  
maximum user load [375](#), [384](#)  
MAXREC argument of FControl cmd [47](#)  
MDCache subcommand [233](#)  
MDCStor subcommand [234](#)  
MENU subcommand [235](#)  
message area [367](#)  
messages  
    error messages [879](#)  
    Performance Toolkit messages [879](#)  
    Performance ToolkitPerformance Toolkit 'action' messages [879](#)  
minidisk  
    cache eligible [426](#)  
    I/O rate [425](#), [426](#), [453](#)  
    load analysis [421](#)  
    minidisk extents [425](#)  
    status [425](#), [453](#)  
minidisk cache log screen [498](#)  
minidisk cache storage log [573](#)  
MINPATHS argument of FControl cmd [48](#)  
missing interrupt detection [415](#)  
MLOAD, smoothed CP I/O response time [427](#)  
MODEL argument of FControl cmd [50](#)  
MODLEVELsubcommand [108](#)

- MONCACHE argument of FControl cmd [51](#)
- MONCOLL argument of FControl cmd [53](#)
- MONData subcommand [236](#)
- monitor data statistics screen [530](#)
- monitor settings report [515](#)
- MONITOR subcommand [110](#)
- MONSCAN subcommand [111](#)
- MONSEt subcommand [237](#)
- MSGCLEAR argument of FControl cmd [63](#)
- MSGWAIT argument of FControl cmd [64](#)
- MTCORElg subcommand [238](#)
- MTCTYPlg subcommand [239](#)
- MTUsers subcommand [240](#)
- MTUSRLog subcommand [241](#)
- multitasking user log [598](#)
- multitasking users screen [596](#)

## N

- NEXT subcommand
  - perf. monitor mode [180](#)
  - redisplay mode [122](#)
- NEXTSamp subcommand [242](#)
- normalized CPU, setting threshold for [43](#)
- normalized logical CPU, setting threshold for [44](#)
- NSS subcommand [243](#)
- NUMBERS argument of FControl cmd [65](#)

## P

- PAGE area [426](#)
- page migration
  - age of blocks at [392](#)
  - blocking factor [392](#)
  - for user [436](#)
- page rate
  - system page rate [379](#)
  - user page rate [432](#), [435](#)
- page wait
  - system [372](#)
  - user [439](#), [448](#), [491](#)
- PAGelog subcommand [244](#)
- paging log screen [502](#)
- path information [420](#), [424](#)
- PCI function activity log screen, format 0 [842](#)
- PCI function activity log screen, format 1 [852](#)
- PCI function activity log screen, format 2 [854](#)
- PCI function activity log screen, format 3 [856](#)
- PCI function activity screen, format 0 [840](#)
- PCI function activity screen, format 1 [851](#)
- PCI function activity screen, format 2 [853](#)
- PCI function activity screen, format 3 [855](#)
- PCI function configuration screen [838](#)
- PCI function menu screen [837](#)
- PCIACT subcommand [245](#)
- PCICONF subcommand [246](#)
- PCILOG subcommand [247](#)
- PCIMENU subcommand [248](#)
- performance data selection menu [471](#)
- performance history data screens [407](#)
- performance log file
  - graphics, based on [185](#), [186](#), [912](#), [955–957](#), [963](#)
- performance monitor displays

performance monitor displays (*continued*)

- Age List Log screen [809](#)
- auxiliary storage log screen [509](#)
- Available List Data Above 2G screen [806](#)
- Available List Data Below 2G screen [804](#)
- benchmark log selection menu [558](#)
- BFS file activity log [599](#)
- BFS locking activity log [604](#)
- BFS PIPEs and linking activity log [601](#)
- cache extended functions log [587](#)
- cached control units screen [566](#)
- channel load screen [410](#), [657](#), [691](#)
- Control unit cache log [552](#)
- CP IUCV services log screens [500](#)
- CP IUCV services screen [408](#)
- CP-owned device screen [416](#)
- CPOWNED device log [553](#)
- CU-cached disks screen [425](#), [428](#)
- DASD performance log [575](#), [586](#)
- DASD seeks distances log (per device) [557](#)
- detailed shared file system screen [457](#)
- Dispatch Vector Activity screen [818](#)
- Dispatch Vector Configuration Change Log screen [816](#)
- DSVBK Steals per Processor Log screen [822](#)
- exception log display [524](#)
- general I/O device screen [412](#), [723](#)
- general Linux CPU utilization screen [712](#)
- general Linux memory utilization screen [714](#)
- general Linux networking activity screen [716](#)
- general shared file system screen [455](#)
- general TCP/IP activity log screen [638](#)
- general TCP/IP data transfer log screen [640](#)
- graphical history data display [955](#)
- graphical user wait state screen [441](#)
- I/O configuration changes screen [590](#)
- I/O device configuration screen [482](#)
- I/O priority queueing screen [702](#)
- I/O Processor Log screen [693](#)
- idle users screen [704](#)
- LGR Data screen [789](#)
- LGR Event Log screen [786](#)
- Linux CPU load log (per user) [718](#)
- Linux CPU utilization screen [689](#)
- Linux data selection menu [711](#)
- Linux details selection screen [675](#)
- Linux filesystem usage screen [684](#)
- Linux memory utilization log (per user) [719](#)
- Linux memory utilization screen [686](#)
- Linux network activity log [720](#)
- Linux network activity screen [682](#)
- Linux systems selection screen [674](#)
- Logical Core Organization Log screen [811](#)
- Logical Partition Activity menu [828](#)
- Logical Partition Logs menu [832](#)
- Logical Partition Share screen [829](#)
- LPAR channel load screen [540](#)
- LPAR configuration log screen [699](#)
- LPAR Load screen [472](#)
- minidisk cache log screen [498](#)
- minidisk cache storage log [573](#)
- monitor data statistics screen [530](#)
- multiprocessor user activity log screen [794](#)
- multiprocessor user wait states log screen [845](#)
- multitasking user log [598](#)



performance monitor displays (*continued*)

- multitasking users screen [596](#)
- paging log screen [502](#)
- privileged operations display [400](#)
- Processor Load and Configuration Logs Menu [701](#)
- Processor Log screen [504](#), [825](#)
- Processor Summary Log screen [705](#)
- Processor Unit (Core and Threads) Configuration Log screen [813](#)
- Processor/Core Configuration Log screen [697](#)
- QDIO activity log (per device) [725](#)
- Real Core Utilization Log screen [820](#)
- remote basic mode screen [593](#)
- remote CMS command execution [470](#)
- remote CP command execution [469](#)
- response time log (all) screen [494](#)
- response time log (UP) screen [496](#)
- RSK disk performance log screen [667](#)
- RSK performance data menu [659](#)
- RSK performance details screen [660](#)
- RSK performance log screen [665](#)
- RSK storage log screen [669](#)
- scheduler data log [564](#)
- scheduler queue log screen [507](#)
- SFS/BFS logs selection menu [671](#)
- shared data spaces screen [489](#)
- shared file system I/O log screen [520](#)
- shared file system log screen [518](#)
- shared file system requests log screen [522](#)
- Shared Segments screen [487](#)
- Steal Statistics screen [807](#)
- system facilities log [575](#)
- system settings screen [526](#)
- system summary log screen [677](#)
- System Topology Machine Organization screen [792](#)
- system transaction statistics screen [480](#)
- TCP/IP buffer pools log screen [648](#)
- TCP/IP denial of service log screen [695](#)
- TCP/IP I/O activity log screen [672](#)
- TCP/IP ICMP messages log screen [642](#)
- TCP/IP links activity log screen [646](#)
- TCP/IP performance data menu [636](#)
- TCP/IP server configuration screen [650](#)
- TCP/IP TCP and UDP sessions log screen [644](#)
- TCP/IP users activity log screen [655](#)
- user communication log (per user) [547](#)
- user communication log (system) [585](#)
- user communication screen [484](#)
- User Page Activity Log screen [800](#)
- User Page Activity screen [798](#)
- user page data log (per user) [543](#)
- user page data screen [435](#)
- User Page Utilization Data Log screen [803](#)
- User Page Utilization Data screen [801](#)
- User Paging Menu screen [797](#)
- user resources per transaction log (per user) [546](#)
- user resources per transaction screen [478](#)
- user response time log (per user) [545](#)
- user response time screen [475](#)
- user status screen [438](#)
- User Transaction Details screen [559](#)
- user wait state log screen [491](#)
- user wait states log (per user) [544](#)
- user-defined displays selection menu [654](#)

performance monitor displays (*continued*)

- user-defined performance data logs [594](#)
- user-defined performance screen [513](#)
- variable correlation plots / graphics [966](#)
- variable correlation screen [607](#)
- virtual disks in storage screen [511](#)
- performance monitor mode
  - initial function key settings [367](#)
- performance monitor report
  - user configuration report [679](#)
- performance monitor reports
  - CPU Load and Transactions [371](#)
  - general I/O device report [721](#), [722](#)
  - LPAR Load Log [632](#)
  - monitor settings [515](#)
  - remote performance log report [377](#)
  - storage utilization [392](#)
  - System Configuration [577](#)
  - system counters report [385](#)
  - User Resource Details [443](#)
  - user resource usage [431](#)
  - user resource usage log (per user) [542](#)
- Performance Toolkit messages [879](#)
- PFKEY argument of FControl cmd [66](#)
- PHYSLog subcommand [249](#)
- PINLOG subcommand [250](#)
- pinned storage log screen [843](#)
- PLOTDET subcommand [182](#)
- PLOTSUM subcommand [185](#)
- PLOTVAR subcommand [189](#)
- PRCLog subcommand [251](#)
- print file header layout [979](#)
- PRINT subcommand [252](#)
- printed reports, table of contents [979](#)
- printing performance data
  - FCONX REPORTS file [252](#)
  - GDDM history graphics [957](#)
  - general perf. data [252](#), [255](#)
- PRINTRmt subcommand [255](#)
- privileged operations [373](#), [379](#), [400](#)
- PRIVOP subcommand [257](#)
- PROCCnf subcommand [258](#)
- PROCESS argument of FControl cmd [70](#)
- Processor configuration log screen [699](#)
- Processor Load and Configuration Logs Menu [701](#)
- Processor Log screen [504](#), [825](#)
- Processor Summary Log screen [705](#)
- Processor Unit (Core and Threads) Configuration Log screen [813](#)
- Processor/Core Configuration Log screen [697](#)
- PROClog subcommand [259](#)
- PROCmenu subcommand [260](#)
- PROCSum subcommand [261](#)
- PSW wait [448](#)
- PUCfglog subcommand [262](#)
- PUOrglog subcommand [263](#)
- purpose of this document [xxvii](#)

**Q**

- QDIO activity log (per device) [725](#)
- QDIO activity screen [723](#)
- QDIO subcommand [264](#)
- QDIOLOG subcommand [265](#)

QEBSM activity log [732](#)  
QEBSM activity screen [730](#)  
queue activity [373](#), [374](#)  
QUIT subcommand  
    operation in basic mode [113](#)

## R

read hits (ctl unit cache) [429](#)  
read hits (minidisk cache) [392](#)  
read hits (minidisk cache). [383](#)  
Real Core Utilization Log screen [820](#)  
real storage [382](#), [392](#)  
record layout  
    HISTLOG files [981](#)  
    HISTSUM files [981](#)  
    Summary record header [985](#)  
REDISP subcommand  
    in basic mode [114](#)  
    perf. monitor mode [268](#), [269](#)  
RELOAD argument of FControl cmd [75](#)  
remote basic mode screen [593](#)  
remote CMS command execution [470](#)  
remote CP command execution [469](#)  
remote history data screen [630](#)  
remote performance log report [377](#)  
REPLY subcommand [115](#)  
report identifiers  
    FCX100, CPU Load and Transactions [371](#)  
    FCX101, remote performance log [377](#)  
    FCX102, system counters [385](#)  
    FCX103, storage utilization [392](#)  
    FCX112, user resource usage [431](#)  
    FCX115, User Resource Details [443](#)  
    FCX149, monitor settings [515](#)  
    FCX162, user resource usage log (per user) [542](#)  
    FCX180, System Configuration [577](#)  
    FCX202, LPAR Load Log [632](#)  
    FCX226, user configuration [679](#)  
rerunning commands [4](#)  
reserved pages [436](#), [445](#)  
RESET subcommand [270](#)  
resident pages [436](#)  
resource wait [375](#)  
RESPall subcommand [271](#)  
response time  
    C1ES, 85% class 1 [374](#), [380](#)  
    non-trivial [374](#), [380](#)  
    trivial [374](#), [380](#)  
response time log (all) screen [494](#)  
response time log (UP) screen [496](#)  
response time, I/O device [414](#), [423](#), [550](#), [567](#)  
RESPUp subcommand [272](#)  
retrieving commands [5](#)  
RETURN subcommand  
    perf. monitor mode [273](#)  
    redisplay mode [124](#)  
Reusable Server Kernel  
    disk performance log [667](#)  
    performance data menu [659](#)  
    performance details [660](#)  
    performance log screen [665](#)  
Reusable Server Kernel disk performance log screen [667](#)  
Reusable Server Kernel performance data menu [659](#)

Reusable Server Kernel performance details screen [660](#)  
Reusable Server Kernel performance log screen [665](#)  
Reusable Server Kernel storage log screen [669](#)  
RIGHT subcommand  
    perf. monitor mode [274](#)  
    redisplay mode [125](#)  
RMTLINES argument of FControl cmd [76](#)  
RSK disk performance log screen [667](#)  
RSK performance data menu [659](#)  
RSK performance details screen [660](#)  
RSK performance log screen [665](#)  
RSK storage log screen [669](#)  
RSKDet subcommand [275](#)  
RSKDisK subcommand [276](#)  
RSKLog subcommand [277](#)  
RSKmenu subcommand [278](#)  
RSKStor subcommand [279](#)

## S

saving on disk  
    GDDM history graphics [957](#)  
    selected general performance data [253](#)  
scale  
    setting with 'TRUNC' argument [190](#), [967](#)  
SCHData subcommand [280](#)  
SCHedlog subcommand [281](#)  
scheduler data log [564](#)  
scheduler queue log screen [507](#)  
screen identifiers  
    FCX104, privileged operations [400](#)  
    FCX105, performance history data screens [407](#)  
    FCX105, performance history plot [963](#)  
    FCX105, summary history plot [964](#)  
    FCX105, variable correlation plot [967](#)  
    FCX106, CP IUCV services [408](#)  
    FCX107, channel load screen [410](#)  
    FCX108, general I/O device screen [412](#)  
    FCX109, CP-owned device screen [416](#)  
    FCX110, I/O device details [420](#)  
    FCX111, CU-cached disks screen [428](#)  
    FCX113, user page data screen [435](#)  
    FCX114, user wait states screen [438](#)  
    FCX116, shared file system servers [455](#)  
    FCX117, SFS server details [457](#)  
    FCX119, SPOOL file summary [461](#)  
    FCX120, SPOOL file user summary [464](#)  
    FCX121, SPOOL file user details [466](#)  
    FCX122, remote CP command execution [469](#)  
    FCX123, remote CMS command execution [470](#)  
    FCX124, performance data selection menu [471](#)  
    FCX126, LPAR Load screen [472](#)  
    FCX127, user response time screen [475](#)  
    FCX128, graphics selection menu [958](#)  
    FCX129, user resources per transaction [478](#)  
    FCX130, system transaction statistics [480](#)  
    FCX131, I/O device configuration [482](#)  
    FCX132, user communication screen [484](#)  
    FCX133, Shared Segments screen [487](#)  
    FCX134, shared data spaces screen [489](#)  
    FCX136, response time log (all) [494](#)  
    FCX137, response time log (UP) [496](#)  
    FCX138, minidisk cache log [498](#)  
    FCX139, Total CP IUCV services log [500](#)

screen identifiers (*continued*)

FCX140, CP IUCV services log (users to CP) [500](#)  
FCX141, CP IUCV services log (CP to users) [500](#)  
FCX142, CP IUCV services log (failures) [500](#)  
FCX143, paging log [502](#)  
FCX144, Processor Log [504](#)  
FCX145, scheduler queue log [507](#)  
FCX146, auxiliary storage log screen [509](#)  
FCX147, virtual disks in storage [511](#)  
FCX148, user-defined performance screen [513](#)  
FCX150, shared file system log [518](#)  
FCX151, shared file system I/O log [520](#)  
FCX152, shared file system requests log [522](#)  
FCX153, exception log display [524](#)  
FCX154, system settings [526](#)  
FCX155, monitor data statistics [530](#)  
FCX156, DASD seeks locations [532](#)  
FCX157, DASD seeks distances [534](#)  
FCX158, DASD seeks details [536](#)  
FCX160, history data selection menu [539](#)  
FCX161, LPAR channel load and channel busy screen [540](#)  
FCX163, user page data log (per user) [543](#)  
FCX164, user wait states log (per user) [544](#)  
FCX165, user response time log (per user) [545](#)  
FCX166, user resources per transaction (per user) [546](#)  
FCX167, user communication log (per user) [547](#)  
FCX169, control unit cache log [552](#)  
FCX170, CPOWNER device log [553](#)  
FCX172, DASD seeks distances log (per device) [557](#)  
FCX173, benchmark log selection menu [558](#)  
FCX174, User Transaction Details [559](#)  
FCX175, scheduler data log [564](#)  
FCX176, cached control units [566](#)  
FCX178, minidisk cache storage log [573](#)  
FCX179, system facilities log [575](#)  
FCX182, user communication log (system) [585](#)  
FCX183, DASD performance log [575](#), [586](#)  
FCX184, cache extended functions log [587](#)  
FCX185, I/O configuration changes [590](#)  
FCX186, remote basic mode screen [593](#)  
FCX187, user-defined performance data logs [594](#)  
FCX188, multitasking users screen [596](#)  
FCX189, multitasking user log [598](#)  
FCX190, BFS file activity log [599](#)  
FCX191, BFS PIPEs and linking activity log [601](#)  
FCX192, BFS locking activity log [604](#)  
FCX193, system selection menu [606](#)  
FCX194, variable correlation display [607](#)  
FCX195, extended history data log [609](#)  
FCX196, cumulative GDDM history display [970](#)  
FCX196, GDDM history display (detailed) [962](#)  
FCX196, GDDM history display (profile) [966](#)  
FCX196, GDDM history display (trend) [965](#)  
FCX197, graphic user wait state display [441](#)  
FCX198, system load overview screen [629](#)  
FCX199, remote history data [630](#)  
FCX201, FCONXREP summary history files [631](#)  
FCX201, FCONXREP summary history files screen [972](#), [973](#)  
FCX203, TCP/IP performance data menu [636](#)  
FCX204, general TCP/IP activity log [638](#)  
FCX205, general TCP/IP data transfer log [640](#)  
FCX206, TCP/IP ICMP messages log [642](#)

screen identifiers (*continued*)

FCX207, TCP/IP TCP and UDP sessions log screen [644](#)  
FCX208, TCP/IP links activity log [646](#)  
FCX210, TCP/IP buffer pools log [648](#)  
FCX212, TCP/IP server configuration [650](#)  
FCX213, user-defined displays selection menu [654](#)  
FCX214, TCP/IP users activity log screen [655](#)  
FCX215, FICON channel load screen [657](#)  
FCX216, Reusable Server Kernel performance data menu [659](#)  
FCX217, Reusable Server Kernel performance details [660](#)  
FCX218, Reusable Server Kernel performance log screen [665](#)  
FCX219, Reusable Server Kernel disk performance log [667](#)  
FCX220, Reusable Server Kernel storage log [669](#)  
FCX221, SFS/BFS logs selection menu [671](#)  
FCX222, TCP/IP I/O activity log [672](#)  
FCX223, Linux systems selection [674](#)  
FCX224, Linux details selection [675](#)  
FCX225, system summary log [677](#)  
FCX227, Linux network activity [682](#)  
FCX228, Linux filesystem usage [684](#)  
FCX229, Linux memory utilization [686](#)  
FCX230, Linux CPU utilization [689](#)  
FCX231, Hipersocket channel activity screen [691](#)  
FCX232, I/O Processor Log [693](#)  
FCX233, TCP/IP denial of service log [695](#)  
FCX234, Processor/Core Configuration Log screen [697](#)  
FCX235, LPAR configuration log screen [699](#)  
FCX236, Processor Load and Configuration Logs Menu [701](#)  
FCX237, I/O priority queueing screen [702](#)  
FCX238, idle users [704](#)  
FCX239, Processor Summary Log [705](#)  
FCX240, VSWITCH activity screen [707](#)  
FCX241, VM Resource Manager screen [709](#)  
FCX242, Linux data selection menu [711](#)  
FCX243, general Linux CPU utilization [712](#)  
FCX244, general Linux memory utilization [714](#)  
FCX245, general Linux networking activity [716](#)  
FCX246, Linux CPU load log (per user) [718](#)  
FCX247, Linux memory utilization log (per user) [719](#)  
FCX248, Linux network activity log [720](#)  
FCX251, QDIO activity screen [723](#)  
FCX252, QDIO activity log (per device) [725](#)  
FCX253, storage utilization log screen [726](#)  
FCX254, available list log screen [728](#)  
FCX255, QEB SM activity screen [730](#)  
FCX256, QEB SM activity log [732](#)  
FCX257, user QDIO activity screen [733](#)  
FCX258, user QDIO activity log [735](#)  
FCX259, demand scan log screen [736](#)  
FCX260, storage management logs menu [739](#)  
FCX261, SXS available page queues management log screen [740](#)  
FCX262, SXS page management log screen [742](#)  
FCX263, SXS deferred tasks management log screen [744](#)  
FCX264, SXS storage utilization log screen [746](#)  
FCX265, Spin Lock Log Screen [748](#)  
FCX266, General Virtual Switch Description Screen [750](#)  
FCX267, Extended VSWITCH activity screen [753](#)



screen identifiers (*continued*)

FCX268, General Virtual Network Device Description Screen [755](#)  
 FCX269, Virtual Network Device Activity Screen [757](#)  
 FCX270, Extended Virtual Network Device Activity Screen [759](#)  
 FCX271, SSI Data Menu [760](#)  
 FCX272, ISFC End Point Configuration [761](#)  
 FCX273, ISFC End Point Activity [764](#)  
 FCX274, ISFC Logical Link Activity [766](#)  
 FCX275, ISFC Logical Link Configuration [768](#)  
 FCX276, SSI Configuration [770](#)  
 FCX277, SSI State Change Synchronization Activity Log [772](#)  
 FCX278, SSI State/Mode Information Log [773](#)  
 FCX279, ISFC End Point Activity Log [775](#)  
 FCX280, ISFC Logical Link Activity Log [776](#)  
 FCX281, ISFC Logical Link Activity By-Time Log [777](#)  
 FCX282, HPF I/O Device Screen [778](#)  
 FCX283, HPF I/O Device Performance Log Screen [782](#)  
 FCX284, I/O Device Data Selection Menu [785](#)  
 FCX285, LGR Event Log screen [786](#)  
 FCX286, LGR Data screen [789](#)  
 FCX287, System Topology Machine Organization [792](#)  
 FCX288, virtual processors usage log [794](#)  
 FCX289, User Paging Menu Screen [797](#)  
 FCX290, User Page Activity Screen [798](#)  
 FCX291, User Page Activity Log Screen [800](#)  
 FCX292, User Page Utilization Data screen [801](#)  
 FCX293, User Page Utilization Data Log screen [803](#)  
 FCX294, Available List Data Below 2G screen [804](#)  
 FCX295, Available List Data Above 2G screen [806](#)  
 FCX296, Steal Statistics screen [807](#)  
 FCX297, Age List Log screen [809](#)  
 FCX298, Logical Core Organization Log [811](#)  
 FCX299, Processor Unit (Core and Threads) Configuration Log [813](#)  
 FCX300, Dispatch Vector Configuration Change Log [816](#)  
 FCX301, Dispatch Vector Activity [818](#)  
 FCX302, Real Core Utilization Log [820](#)  
 FCX303, DSVBK Steals per Processor Log [822](#)  
 FCX304, Processor Log [825](#)  
 FCX305, Logical Partition Activity menu [828](#)  
 FCX306, Logical Partition Share screen [829](#)  
 FCX307, Logical Partition Logs menu [832](#)  
 FCX308, CPU pool configuration screen [833](#)  
 FCX309, CPU pool activity screen [835](#)  
 FCX310, PCI function menu screen [837](#)  
 FCX311, PCI function configuration screen [838](#)  
 FCX312, PCI function activity screen, format 0 [840](#)  
 FCX313, PCI function activity log screen, format 0 [842](#)  
 FCX314, pinned storage log screen [843](#)  
 FCX315, multiprocessor user wait states log [845](#)  
 FCX316, subpool storage log screen [848](#)  
 FCX317, global networking object activity screen [849](#)  
 FCX318, PCI function activity screen, format 1 [851](#)  
 FCX319, PCI function activity log screen, format 1 [852](#)  
 FCX320, PCI function activity screen, format 2 [853](#)  
 FCX321, PCI function activity log screen, format 2 [854](#)  
 FCX322, PCI function activity screen, format 3 [855](#)  
 FCX323, PCI function activity log screen, format 3 [856](#)  
 FCX324, CPU pool menu screen [857](#)  
 FCX325, CPU activity selection screen [858](#)  
 FCX326, spin lock collision activity screen [859](#)

screen identifiers (*continued*)

FCX327, HyperPAV alias activity screen [861](#)  
 FCX328, HyperPAV share activity screen [863](#)  
 FCX329, DASD volume performance screen [865](#)  
 FCX330, DASD volume performance log screen [868](#)  
 FCX331, SMT Metrics Per Core Type Log screen [871](#)  
 FCX332, SMT Metrics Per Core Log screen [873](#)  
 FCX333, User Processor Time screen [875](#)  
 FCX334, User Processor Time Log screen [877](#)  
 SCROLL argument of FControl cmd [77](#)  
 SCSI device report [721](#)  
 SCSI subcommand [282](#)  
 SCSI utilization log [722](#)  
 SCSILog subcommand [283](#)  
 SEARCH argument of FControl cmd [79](#)  
 SECUSER argument of FControl cmd [80](#)  
 SEEK  
   DASD seeks details screen [536](#)  
   DASD seeks distances screen [534](#)  
   DASD seeks locations screen [532](#)  
   last seek (real) [422](#)  
   last seek (virtual) [425](#)  
   SEEKs by minidisk [536](#)  
 SEEKDEt subcommand [284](#)  
 SEEKdist subcommand [285](#)  
 SEEKDLog subcommand [286](#)  
 SEEKLLog subcommand [287](#)  
 SEEKLoc subcommand [288](#)  
 service time, I/O device [414](#), [423](#), [550](#), [567](#)  
 SETEVENT argument of FControl cmd [81](#)  
 setting  
   color and highlighting [18](#)  
   PF-keys for basic mode [66](#)  
   thresholds for perf. monitoring [48](#), [51](#)  
   thresholds for performance monitoring [42](#)  
   thresholds for user monitoring [26](#), [33](#), [38](#), [94](#), [97](#)  
 setting timer events [81](#)  
 SETTINGS argument of FControl cmd [83](#)  
 severity code for exception messages [37](#), [45](#), [49](#), [52](#), [98](#)  
 SFS and BFS logs selection menu [671](#)  
 SFS I/O log screen [520](#)  
 SFS log screen [518](#)  
 SFS sequests log screen [522](#)  
 SFS subcommand [289](#)  
 SFSIOLOG subcommand [290](#)  
 SFSLog subcommand [291](#)  
 SFSMenu subcommand [292](#)  
 SFSReq subcommand [293](#)  
 shared data spaces screen [489](#)  
 Shared Segments screen [487](#)  
 SKIPsmp subcommand [294](#)  
 SLR, collecting summary data [985](#)  
 SMT Metrics Per Core Log screen [873](#)  
 SMT Metrics Per Core Type Log screen [871](#)  
 SORT subcommand [295](#)  
 special line processing [70](#)  
 spin lock collision activity screen [859](#)  
 Spin Lock Log Screen [748](#)  
 SPOOL file summary screen [461](#)  
 SPOOL file user details screen [466](#)  
 SPOOL file user summary screen [464](#)  
 SPOOL subcommand [296](#)  
 SSI Configuration [770](#)  
 SSI Data Menu [760](#)

- SSI State Change Synchronization Activity Log [772](#)
- SSI State/Mode Information Log [773](#)
- SSIConf subcommand [297](#)
- SSIMenu subcommand [298](#)
- SSISChlg subcommand [299](#)
- SSISMilg subcommand [300](#)
- status
  - virtual machines
    - status
      - of virtual machines [449](#)
- status of
  - GRAF devices [422](#)
  - I/O devices [421](#)
  - minidisks [425](#), [453](#)
  - virtual devices [451](#)
  - virtual machines [432](#), [438](#)
- Steal Statistics screen [807](#)
- STEALlog subcommand [301](#)
- storage management logs menu [739](#)
- STORAGE subcommand [302](#)
- storage utilization log screen [726](#)
- storage utilization report [392](#)
- storage wait [375](#), [381](#)
- SUBPLOG subcommand [305](#)
- subpool storage log screen [848](#)
- summary file records [985](#)
- SXS available page queues management log screen [740](#)
- SXS deferred tasks management log screen [744](#)
- SXS page management log screen [742](#)
- SXS storage utilization log screen [746](#)
- SYSConf subcommand [310](#)
- SYSLOAD subcommand [311](#)
- SYSLog subcommand [312](#)
- SYSMENU subcommand [313](#)
- SYSSet subcommand [314](#)
- SYSSumlg subcommand [315](#)
- system areas
  - DRCT area [417](#), [553](#)
  - DUMP area [417](#), [553](#)
  - PAGE area [417](#), [553](#)
  - SPOOL area [417](#), [553](#)
  - TDSK area [417](#), [553](#)
- System Configuration report [577](#)
- system facilities log [575](#)
- system load overview screen [629](#)
- system selection menu [606](#)
- system settings screen [526](#)
- SYStem subcommand [316](#)
- system summary log screen [677](#)
- System Topology Machine Organization screen [792](#)
- system transaction statistics screen [480](#)
- SYSTRans subcommand [317](#)

## T

- table of contents for printed reports [979](#)
- TCP/IP
  - activity [638](#)
  - buffer pools [648](#)
  - data transfer [640](#)
  - denial of service [695](#)
  - I/O activity [672](#)
  - ICMP messages activity [642](#)
  - links activity [646](#)

- TCP/IP (*continued*)
  - performance data [636](#)
  - server configuration [650](#)
  - sessions [644](#)
  - TCP/IP buffer pools log screen [648](#)
  - TCP/IP denial of service log screen [695](#)
  - TCP/IP I/O activity log screen [672](#)
  - TCP/IP ICMP messages log screen [642](#)
  - TCP/IP links activity log screen [646](#)
  - TCP/IP performance data menu [636](#)
  - TCP/IP server configuration screen [650](#)
  - TCP/IP TCP and UDP sessions log screen [644](#)
  - TCP/IP users activity log screen [655](#)
  - TCPActlg subcommand [318](#)
  - TCPBplog subcommand [319](#)
  - TCPCnf subcommand [320](#)
  - TCPDatlg subcommand [321](#), [322](#)
  - TCPIcmp subcommand [323](#)
  - TCPIolog subcommand [324](#)
  - TCPLinks subcommand [325](#)
  - TCPMenu subcommand [326](#)
  - TCPSEss subcommand [327](#)
  - TCPUUsers subcommand [328](#)
  - throttling indicator [413](#), [549](#)
  - time interval [367](#)
  - timer selection [88](#)
  - TIMStmp argument of FControl cmd [88](#)
  - TOP subcommand
    - perf. monitor mode [329](#)
    - redisplay mode [126](#)
  - TOPolog subcommand [330](#)
  - total I/O rate [379](#)
  - trademarks [1088](#)
  - trend analysis
    - performance trend display [964](#)
    - variable correlation display [967](#)
  - trend file record layout
    - Summary record header [985](#)
  - trend file records [983](#)
  - trend file variable selection [987](#)
  - TRNDSCAN subcommand [117](#)

## U

- UCLass argument of FControl cmd [89](#)
- UCOmm subcommand [331](#)
- UCOMMLog subcommand [332](#)
- UCONF subcommand [333](#)
- UDEFMENU subcommand [334](#)
- UP subcommand
  - perf. monitor mode [136](#)
  - redisplay mode [120](#)
- UPAGE subcommand [335](#)
- UPAGELog subcommand [336](#)
- UPDTCMS argument of FControl cmd [90](#)
- UPDTSCRN argument of FControl cmd [91](#)
- UPGAct subcommand [337](#)
- UPGACTLg subcommand [338](#)
- UPGMenu subcommand [339](#)
- UPGUtl subcommand [340](#)
- UPGUTLLg subcommand [341](#)
- UREsp subcommand [344](#)
- URESPLog subcommand [345](#)
- used

- used (*continued*)
  - space in system areas [417](#), [553](#)
  - SPOOL and PAGE areas [382](#), [416](#)
- user classes
  - defining user classes [89](#)
- user communication log (per user) [547](#)
- user communication log (system) [585](#)
- user communication screen [484](#)
- user configuration report [679](#)
- User Page Activity Log screen [800](#)
- User Page Activity screen [798](#)
- user page data log (per user) [543](#)
- user page data screen [435](#)
- User Page Utilization Data Log screen [803](#)
- User Page Utilization Data screen [801](#)
- User Paging Menu screen [797](#)
- User Processor Time Log screen [877](#)
- User Processor Time screen [875](#)
- user QDIO activity log [735](#)
- user QDIO activity screen [733](#)
- user report [431](#)
- User Resource Details report [443](#)
- user resource usage log (per user) [542](#)
- user resources per transaction log (per user) [546](#)
- user resources per transaction screen [478](#)
- user response time log (per user) [545](#)
- user response time screen [475](#)
- user status screen [438](#)
- USER subcommand [346](#)
- User Transaction Details screen [559](#)
- user wait state log screen [491](#)
- user wait state screen (GDDM) [441](#)
- user wait states log (per user) [544](#)
- user-defined displays selection menu [654](#)
- user-defined performance data logs [594](#)
- user-defined performance screen [513](#)
- USERBOTL argument of FControl cmd [92](#)
- USERHDR argument of FControl cmd [93](#)
- USERLog subcommand [347](#)
- USERVAR argument of FControl cmd [94](#)
- USRLIMIT argument of FControl cmd [97](#)
- USRMPLOG subcommand [348](#)
- USRPrctm subcommand [349](#)
- USRTmlog subcommand [350](#)
- USTAT subcommand [351](#)
- USTATG subcommand [351](#)
- USTATLog subcommand [352](#)
- USTLog subcommand [353](#)
- USTMPLOG subcommand [354](#)
- UTRANDet subcommand [355](#)
- UTRANLog subcommand [356](#)
- UTRans subcommand [357](#)

## V

- variable correlation display [607](#)
- variable correlation plot [966](#)
- variable selection for trend files [987](#)
- VDISks
  - overall storage utilization [392](#)
- VDIsks subcommand [358](#)
- virtual disks in storage screen [511](#)
- virtual I/O devices [451](#)
- virtual I/O rate [379](#)

- Virtual Network Device Activity Screen [757](#)
- VM Resource Manager Screen [709](#)
- VMCF [60](#)
- VMRM subcommand [359](#)
- VNIC subcommand [360](#)
- VOLLOG subcommand [361](#)
- VOLUME subcommand [362](#)
- VSWITCH activity screen [707](#)
- VSWItch subcommand [363](#)

## W

- wait states
  - system [372](#)
  - user [375](#), [381](#), [432](#), [438](#), [439](#), [448](#), [491](#)
- weight of exceptions
  - setting 'weight' value [37](#), [48](#), [52](#), [98](#)
  - setting weight value [44](#)
- write hits (ctl unit cache) [429](#), [567](#), [571](#), [589](#)

## X

- XSTORE frame requests [383](#), [389](#)







Product Number: 5741-A09

Printed in USA

SC24-6303-74

