

z/VM



General Information

version 6 release 2

z/VM



General Information

version 6 release 2

Note:

Before using this information and the product it supports, read the information in “Notices” on page 111.

This edition applies to version 6, release 2, modification 0 of IBM z/VM (product number 5741-A07) and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces GC24-6193-01.

© **Copyright IBM Corporation 1990, 2011.**

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

Contents

Tables	vii
About this document	ix
Intended audience	ix
Where to find more information	ix
How to send your comments to IBM	xi
If you have a technical problem	xi
Chapter 1. How z/VM can help you	1
Chapter 2. z/VM overview	3
z/VM components, facilities, and optional features	3
z/VM virtualization technology provides guest support	4
z/VM provides proven system integrity, security, availability, and reliability	5
z/VM supports application development and deployment	6
z/VM is accessible by people with disabilities	7
Chapter 3. What is new or changed in z/VM V6.2.	9
Virtualization technology and Linux enablement	9
z/VM single system image clusters	9
Introducing live guest relocation.	10
Technology exploitation	11
IBM zEnterprise 196 and IBM zEnterprise 114	11
OSX and OSM channel path	11
FICON Express8S and OSA-Express4S support	12
XRC time stamping	12
HyperSwap improvements.	12
IBM Extended Address Volumes	12
Direct attachment of IBM XIV Storage Systems	13
Dynamic emulated device path control	13
IBM System Storage TS1140 Tape Drive (3592 Model E07)	13
CPU-Measurement Counter Facility Host Support	13
Systems management	13
Support for the IBM zEnterprise Unified Resource Manager	14
Dynamic discovery of FCP disks	15
Enhanced systems management APIs	15
Enhanced Directory Maintenance Facility	15
Enhanced Performance Toolkit	16
Network virtualization	16
Support for IBM zEnterprise System networks	16
MPROUTE upgrade	16
Additional IPv6 support	16
FTP enhancements	17
Multiple access ports per guest	17
OSA Address Table support	17
LDAP server upgrade	17
Scalability	17
Memory constraint relief	17
Disabling the guest page reordering process	17
Improved page release serialization	18
Enhanced contiguous frame coalescing	18
Alternate method for hard limiting of scheduled system resources	18

Security	18
z/VM SSL server upgrade and FIPS support	18
ESM access control for real devices	18
ESM access control for SET SECUSER and SET OBSERVER	19
Enhanced RACF Security Server	19
Crypto Express3 support	19
Protected key CPACF support	19
EAL4+ certification for z/VM V6.1	19
Application enablement	20
z/Architecture CMS	20
Program Management Binder upgrade	20
Language Environment upgrade	20
Support for XL C/C++ Compiler for z/VM, V1.3	20
Installation, serviceability, and packaging	20
IBM z/VM Single System Image Feature	20
Installation process improvements	21
VMSES/E enhancements	21
Common service file pool	21
Additional changes	21
Withdrawal of support	21
Shipping the zMAP agents with z/VM discontinued	22
Statement of direction	22
Chapter 4. Technical information	25
Hardware and storage requirements	25
Server requirements	25
Installation requirements	25
Storage requirements	26
Workstation requirement	26
CSE hardware requirements	27
OSA/SF hardware requirements	27
Performance Toolkit hardware requirements	27
SSI clusters hardware requirements	27
TCP/IP hardware requirements	28
VM Guest RSA-Assist Support for Linux hardware requirements	28
Device support	28
Supported devices	28
Unsupported devices	29
Software (program) requirements	29
Prerequisite licensed programs	30
High Level Assembler	30
Levels of the z/VM base facilities and optional features	31
DFSMS/VM additional program requirements	31
DirMaint additional program requirements	32
HCD and HCM additional program requirements	32
IOCP additional program requirements	32
Language Environment additional program requirements	32
OpenExtensions additional program requirements	32
OSA/SF additional program requirements	33
Performance Toolkit additional program requirements	33
RACF additional program requirements	33
RSCS additional program requirements	34
Shared-DASD complex and CSE complex additional program requirements	34
SSI clusters program requirements	34
TCP/IP additional program requirements	34
Online books additional program requirements	35

POSIX application development program requirements	35
Program requirements for other products	35
Limitations and restrictions	35
General restriction for z/VM virtual machines	35
DirMaint restrictions	36
Performance considerations	36
Operating systems supported as guests.	36
Other programs supported on z/VM	36
Multicultural support	37
Packaging and ordering information	37
Object code only and limited source availability	38
Integrity and security.	38
Data integrity for guests	39
System integrity statement for z/VM	39
Security, auditability, and control	40
Chapter 5. z/VM base product.	45
CP	45
Virtual machine modes	45
Architecture compatibility	46
z/Architecture support	46
ESA/XC support	46
Cross system extensions	46
Guest LANs and virtual switches	46
Inter-System Facility for Communications	47
VM Dump Tool	47
Support for hardware facilities, features, and architectures	47
CMS.	55
Versions of CMS	55
Shared File System	55
CMS application programming	56
Systems management API	56
CMS Pipelines	57
CMS application multitasking.	57
OpenExtensions	57
Program Management Binder for CMS	58
Reusable Server Kernel	58
Virtual Machine Resource Manager	58
XEDIT	59
HELP facility.	59
TCP/IP for z/VM	59
Link protocols	60
Network protocols	60
Transport protocols	60
Applications and protocols.	61
AVS	62
Dump Viewing Facility	62
GCS.	62
HCD and HCM for z/VM	62
Language Environment	63
OSA/SF	63
REXX/VM.	64
TSAF	64
VMSES/E	64
Chapter 6. z/VM optional features	65

Data Facility Storage Management Subsystem for VM	65
Directory Maintenance Facility for z/VM	65
IBM z/VM Single System Image Feature	66
Performance Toolkit for VM	67
RACF Security Server for z/VM	68
RSCS Networking for z/VM	69
Chapter 7. z/VM library guide	71
Where to get z/VM information	71
z/VM Information Center	71
z/VM Internet Library	71
z/VM Collection DVD	72
IBM Publications Center	72
Documentation supplied with z/VM	72
Abstracts of the publications for the z/VM base product	72
System overview	73
Installation, migration, and service	73
Planning and administration	74
Customization and tailoring	77
Operation and use	77
Application programming	79
Diagnosis	84
Abstracts of the publications for the z/VM optional features	85
DFSMS/VM	86
Directory Maintenance Facility for z/VM	86
Performance Toolkit for VM	87
RACF Security Server for z/VM	87
RSCS Networking for z/VM	88
Appendix A. IBM servers supported by z/VM	91
Appendix B. IBM operating systems supported as guests of z/VM	95
Appendix C. IBM devices supported by z/VM	97
Direct access storage devices (DASD)	98
DASD control units and storage controls	102
Tape units and tape libraries	102
Tape control units	103
Printers	103
Card readers and card punches	104
Terminals, displays, and consoles	104
Display printers	105
Display control units	105
Communications controllers	105
Switches	105
Server adapters and miscellaneous devices	105
Notices	111
Trademarks	113
Index	115

Tables

1. IBM-translated parts of z/VM	37
2. Server support.	92
3. Guest support	95
4. Device support – DASD	98
5. Device support – DASD control units and storage controls	102
6. Device support – tape units and tape libraries.	102
7. Device support – tape control units.	103
8. Device support – printers	103
9. Device support – card readers and card punches	104
10. Device support – display control units.	105
11. Device support – communications controllers	105
12. Device support – switches	105
13. Device support – server adapters and miscellaneous devices	105

About this document

This document provides the following information about the IBM® z/VM® licensed program:

- How z/VM can help you
- z/VM overview
- What is new or changed in z/VM V6.2
- z/VM hardware and software requirements
- Guide to the z/VM product information
- IBM servers, guest operating systems, and devices supported by current z/VM releases

Intended audience

This information is intended for anyone who wants a general overview of z/VM. It is also useful for those who need to evaluate the capabilities of z/VM and determine the resources necessary to install and run it.

Where to find more information

You can obtain more information about z/VM from the documents described in Chapter 7, “z/VM library guide,” on page 71.

Note: This edition is published for the announcement of z/VM V6.2. Sources of additional information identified in this document might not be available until the general availability (GA) of z/VM V6.2.

Links to other online documents

The online version of this document contains links to other online documents. These links are to editions that were current when this document was published. However, due to the nature of some links, 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. Also, a link from this document to another document works only when both documents are in the same directory.

How to send your comments to IBM

We appreciate your input on this publication. Feel free to comment on the clarity, accuracy, and completeness of the information or give us any other feedback that you might have.

Use one of the following methods to send us your comments:

1. Send an email to mhvrcfs@us.ibm.com.
2. Go to IBM z/VM Reader's Comments (www.ibm.com/systems/z/os/zvm/zvmforms/webqs.html).
3. Mail the comments to the following address:
IBM Corporation
Attention: MHVRCFS Reader Comments
Department H6MA, Building 707
2455 South Road
Poughkeepsie, NY 12601-5400
U.S.A.
4. Fax the comments to us as follows:
From the United States and Canada: 1+845+432-9405
From all other countries: Your international access code +1+845+432-9405

Include the following information:

- Your name and address
- Your email address
- Your telephone or fax number
- The publication title and order number:
z/VM V6R2 General Information
GC24-6193-02
- The topic name or page number related to your comment
- The text of your comment

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

IBM or any other organizations will use the personal information that you supply only to contact you about the issues that you submit to IBM.

If you have a technical problem

Do not use the feedback methods listed above. Instead, do one of the following:

- Contact your IBM service representative.
- Contact IBM technical support.
- See IBM: z/VM Service Resources (www.ibm.com/vm/service/).
- Go to IBM Support Portal (www.ibm.com/support/entry/portal/Overview/).

Chapter 1. How z/VM can help you

Today's businesses are challenged to accelerate the deployment of new applications to address the needs of the business while also reducing cost, improving service, and managing risks. Operational costs need to be more than just contained. Businesses are looking to reduce complexity and achieve breakthrough productivity gains through virtualization, optimization, energy stewardship, and flexible sourcing. Ensuring high availability and the quality of existing services is expected, but meeting customer expectations for realtime, dynamic access to innovative new services such as cloud, places requirements for an innovative approach. Security, resiliency, and compliance challenges need to be managed while also preparing for the new risks posed by an even more connected and collaborative world. IBM System z® with z/VM delivers superior IT optimization and consolidation that offers a clear advantage over competitive alternatives by delivering better business value in these key areas.

One of the greatest strengths of z/VM is the ability to share physical resources among virtual machines. z/VM offers the highest levels of resource sharing, including the sharing of Linux program executables, overcommitment capabilities for processors and memory, cooperative memory management, I/O bandwidth, and system availability, resulting in nearly 100 percent utilization of the system resources nearly 100 percent of the time. In addition, z/VM can significantly overcommit those real resources, which allows users to create a set of virtual machines with assets that considerably exceed the amount of real hardware. z/VM offers several data-in-memory techniques that further enhance the scalability and performance of memory-intensive workloads. In-memory emulated storage, called virtual disks, allows Linux systems to achieve memory-speed data transfers for read and write I/O operations and has proven beneficial when used to store Linux swap (paging) data.

Enterprises that require multisystem server solutions will find that z/VM can help them address their IT infrastructure needs with concurrent supports for virtual machines running different operating systems in a secure isolated environment. z/VM supports z/OS®, z/VSE®, z/TPF, z/VM, and Linux on System z operating environments, and the CMS application environment. The ability to support multiple machine images and architectures enables z/VM to run multiple production and test versions of System z operating systems, all on the same System z server. z/VM provides a highly flexible test and production environment for enterprises that simplifies migration from one release to another, facilitates the transition to newer applications, provides a test system whenever one is needed, and consolidates several systems onto one physical server.

Existing z/VM customers will find z/VM V6.2 particularly attractive for its ability to help clients avoid outages for virtual servers when performing maintenance. Live guest relocation has the ability to move running Linux virtual servers to another LPAR or another mainframe server without disrupting the business. Multisystem virtualization allows up to 4 z/VM instances on other LPARs or other mainframes to be clustered and managed as a single system image, which can help clients avoid the virtual machine sprawl challenges of competitive systems and provides cluster management from any z/VM member in the single system image. The ability to provide a set of shared resources for the z/VM systems and their hosted virtual machines and manage those shared resources as a single resource pool simplifies systems management of a multi-z/VM environment and provides a more manageable infrastructure for cloud computing.

z/VM can help you

For the most current information on z/VM, see IBM: z/VM Operating System (www.ibm.com/vm/).

Chapter 2. z/VM overview

z/VM provides a highly flexible test and production environment on the IBM System z platform. The z/VM implementation of IBM virtualization technology provides the capability to run full-function operating systems such as Linux on IBM System z, z/OS, z/VSE, z/TPF, and z/VM as “guests” of z/VM. z/VM supports 64-bit IBM z/Architecture® guests and 31-bit IBM Enterprise Systems Architecture/390 (ESA/390) guests.

z/VM components, facilities, and optional features

The z/VM base product includes the following components and facilities:

Control Program (CP)

CP is a hypervisor and real-machine resource manager.

Conversational Monitor System (CMS)

CMS provides a high-capacity application and interactive user environment and provides the z/VM file systems.

TCP/IP for z/VM

TCP/IP for z/VM provides support for the TCP/IP networking environment.

Advanced Program-to-Program Communication/Virtual Machine (APPC/VM) Virtual Telecommunications Access Method (VTAM®) Support (AVS)

AVS provides connectivity in an IBM Systems Network Architecture (SNA) network.

Dump Viewing Facility

Dump Viewing Facility is a tool for interactively diagnosing z/VM system problems.

Group Control System (GCS)

GCS is a virtual machine supervisor that provides multitasking services and supports a native SNA network.

Hardware Configuration Definition (HCD) and Hardware Configuration Manager (HCM) for z/VM

HCD and HCM provide a comprehensive I/O configuration management environment.

Language Environment® for z/VM

Language Environment provides the runtime environment for z/VM application programs written in C/C++, COBOL, or PL/I.

Open Systems Adapter/Support Facility (OSA/SF)

OSA/SF is a tool for customizing the modes of operation of OSA hardware features.

REXX/VM

REXX/VM contains the interpreter for processing the REXX programming language.

Transparent Services Access Facility (TSAF)

TSAF provides communication services within a collection of z/VM systems without using VTAM.

Virtual Machine Serviceability Enhancements Staged/Extended (VMSES/E)

VMSES/E provides a tools suite for installing and servicing z/VM and other enabled products.

For more information about these components and facilities, see Chapter 5, “z/VM base product,” on page 45.

z/VM also offers the following optional features:

Data Facility Storage Management Subsystem for VM (DFSMS/VM)

DFSMS/VM controls data and storage resources.

Directory Maintenance Facility for z/VM (DirMaint™)

DirMaint provides interactive facilities for managing the z/VM user directory.

IBM z/VM Single System Image Feature (VMSSI)

VMSSI enhances the z/VM systems management, communications, disk management, device mapping, virtual machine definition management, installation, and service functions to enable up to four z/VM systems to share and coordinate resources in a single system image (SSI) cluster.

Performance Toolkit for VM

Performance Toolkit provides tools for analyzing z/VM and Linux performance data.

Remote Spooling Communications Subsystem (RSCS) Networking for z/VM

RSCS enables users to send messages, commands, files, and jobs to other users in a network, to send print jobs to local or remotely attached printers, and to receive remote print jobs.

Resource Access Control Facility (RACF®) Security Server for z/VM

RACF provides data security for an installation by controlling access to it.

For more information about these optional features, see Chapter 6, “z/VM optional features,” on page 65.

z/VM virtualization technology provides guest support

z/VM provides each user with an individual working environment known as a *virtual machine*. A virtual machine simulates the existence of a dedicated real machine, including processor functions, memory, networking, and input/output (I/O) resources. Operating systems and application programs can run in virtual machines as guests. As a result, development, testing, and production environments can share a single physical computer.

The virtual machine capability of z/VM allows you to:

- **Test programs that can cause abnormal termination of real machine operations.** The isolation that is provided for a virtual machine enables system-oriented programs and teleprocessing applications, for example, to be tested on the virtual machine while production work is in progress, because this testing cannot cause abnormal termination of the real machine.
- **Test a new operating system release.** A new release of an operating system can be generated and tested at the same time that the existing release is performing production work. This enables the new release to be installed and put into production more quickly. The ability to operate multiple operating systems concurrently under z/VM might enable an installation to continue running programs that operate only under a back-level release (programs that are release-sensitive and uneconomical to convert, for example) concurrently with the most current release.
- **Develop a new operating system or release.** A new operating system or a new release can be developed and tested without dedicating the real machine to this function.

- **Perform operating system maintenance concurrently with production work.** The installation and testing of program temporary fixes (PTFs) for an operating system can be done at the same time normal production operations are in progress.
- **Provide backup facilities for the primary operating system.** A generated z/VM system is not model-dependent and can operate on various server models as long as the minimum hardware requirements are present. This enables a smaller server model that has less real storage, fewer channels, fewer direct access devices, and fewer unit record devices than a larger server model to provide backup for the larger model (normally at a reduced level of performance).
- **Perform operator training concurrently with production work processing.** The real machine does not have to be dedicated to training additional or new operators or to providing initial training when a new operating system is installed. Operator errors cannot cause termination of real machine operations.
- **Simulate new system configurations before the installation of additional channels and I/O devices.** The relative load on channels and I/O devices can be determined using the simulated I/O configuration rather than the real I/O configuration. Experience with generating and operating an I/O configuration for multiple guests can be obtained using one real machine.
- **Test customer-written system exits.** Customer-written system exits can be tested without disrupting production work.

z/VM provides proven system integrity, security, availability, and reliability

z/VM is built on a foundation of system integrity and security, and incorporates many design features for availability and reliability:

- Integrity and security:
 - z/VM supports guest use of the cryptographic facilities provided by supported IBM servers.
 - z/VM supports drive-based data encryption with IBM System Storage® tape drives. z/VM supports native and guest use of the IBM Encryption Key Manager for encryption of tapes.
 - z/VM supports the use of DASD volumes that reside on data encryption drives (DEDs).
 - IBM will correct any integrity exposures introduced by unauthorized programs into the system.
 - Secure Sockets Layer (SSL) support is provided through TCP/IP for z/VM. The TCP/IP SSL server can operate in Federal Information Processing Standard (FIPS) mode.
 - Integrated access control and authentication services can be augmented with the addition of an external security manager (ESM), such as the RACF Security Server for z/VM. RACF can also be used to audit connections to z/VM real devices.
- Availability and reliability:
 - Application recovery: z/VM provides services which permit recovery of incomplete interactions with resource managers.
 - Automated operations: z/VM offers several levels of automated system management support. One example is the Programmable Operator. For a higher degree of automation, IBM Operations Manager for z/VM can be added.

Introduction

- z/VM provides duplexed data with transparent ongoing synchronization between the primary and backup copy, and automatic transparent switching to the backup copy in case of an error in the primary copy.
- Online configuration changes eliminate many previously required outages.
- z/VM systems can be connected for improved server and user availability.
- Fast restart reduces the end user impact of any outage.
- Setting up a z/VM single system image (SSI) cluster and using live guest relocation allows z/VM and hardware maintenance to be less disruptive to workloads and allows less disruptive workload balancing.

z/VM supports application development and deployment

z/VM supports a number of programming languages and environments, including:

- APL2[®]
- Assembler
- C
- C++
- COBOL
- FORTRAN
- Language Environment
- MQSeries[®] Client (WebSphere[®] MQ)
- PASCAL
- PL/I
- REXX
- Systems Application Architecture[®] (SAA) common programming interface (CPI)
- VisualAge[®] Generator

z/VM also provides a rich set of application development services, including:

- Integrated editors and support for code version management
- Trace and debug facilities

z/VM supports program development not only for z/VM applications, but also for operating systems supported as guests of z/VM.

z/VM application programming interfaces (APIs) include:

- CP DIAGNOSE instructions
- Certain z/VM control blocks
- Data record formats, such as accounting records, intended to be processed by application programs
- CP system services (such as *MSG)
- CP assembler macros (such as IUCV and APPCVM)
- TCP/IP APIs (such as C sockets, IUCV sockets, and Remote Procedure Calls (RPCs))
- GCS assembler macros

Special facilities are available to CMS applications, such as:

- Systems management API
- Enterprise Systems Architecture/Extended Configuration (ESA/XC) support
- CMS Pipelines

- Callable services library (CSL)
- CMS multitasking services
- OpenExtensions (POSIX interfaces and shell)
- Program Management Binder
- Reusable Server Kernel
- Assembler macros and functions
- OS/MVS simulation
- DOS/VSE support

Note: Although IBM VSE/VSAM, V6.1 (5686-081), was withdrawn from marketing on September 30, 2005, CMS DOS/BAM will continue to provide its current level of function in support of DOS simulation and VSE/VSAM.

In general, z/VM programming interfaces are designed to be used exclusively from programs (often using binary or other machine-level formats as parameter values) and are usually supported in a compatible manner from release to release.

z/VM might also ship other program materials (primarily macros), but these are internal facilities designed only for use between z/VM components and modules and are not intended to be used as programming interfaces.

z/VM is accessible by people with disabilities

The following features support use by people with disabilities:

- Operation by keyboard alone
- Optional font enlargement and high-contrast display settings
- Screen readers and screen magnifiers tested for use by people with visual impairment

Introduction

Chapter 3. What is new or changed in z/VM V6.2

This section provides an overview of the new functions, enhancements, and other changes included in z/VM V6.2.

Notes:

1. Some product changes described in this section might be provided to previous z/VM releases through program temporary fixes (PTFs) for authorized program analysis reports (APARs). See IBM: z/VM Service Resources (www.ibm.com/vm/service/).

Also refer to the appropriate hardware Preventive Service Planning (PSP) buckets at IBM Resource Link (www.ibm.com/servers/resourcelink/). A Resource Link® user ID and password are required.

2. Some product changes might require additional hardware or software support. See Chapter 4, “Technical information,” on page 25.

The changes are described under the following topics:

- “Virtualization technology and Linux enablement”
- “Technology exploitation” on page 11
- “Systems management” on page 13
- “Network virtualization” on page 16
- “Scalability” on page 17
- “Security” on page 18
- “Application enablement” on page 20
- “Installation, serviceability, and packaging” on page 20
- “Additional changes” on page 21

For planned future changes to z/VM, see “Statement of direction” on page 22.

For additional information about the changes to z/VM external interfaces in this release and previous releases, see *z/VM: Migration Guide*.

Virtualization technology and Linux enablement

This section describes extensions to z/VM virtualization technology in support of Linux on System z, z/OS, and other guests.

z/VM single system image clusters

z/VM V6.2 implements multisystem virtualization using a z/VM single system image (SSI) cluster composed of up to four z/VM systems. This multisystem virtualization technology for the mainframe extends the z/VM virtualization technology to a new level, allowing members of the cluster to share resources and synchronize with other nodes, together presenting the appearance of a single system.

Members of a z/VM SSI cluster are part of the same Inter-System Facility for Communications (ISFC) collection and use ISFC channel connections to communicate. All members of a cluster also share DASD for virtual machines and selected z/VM data, as well as LAN segments and IP subnets. The concept of a global virtual switch provides identical network connectivity across all active members within a cluster.

New or changed

Simplified z/VM systems management is realized when z/VM instances are members of an SSI cluster and can be serviced and administered as one system. Sharing of resources used by both CP and virtual machines is coordinated among all members. This allows Linux guests to access the same devices and networks regardless of which member they are logged on or relocated to. Shared resources include:

- User directory
- Minidisks
- Spool files
- Network device MAC addresses

Each member of a z/VM SSI cluster is able to communicate with all other active members. When a z/VM system is configured as a member of a cluster, it automatically joins the other members during system startup. Coordination of members joining and leaving the cluster, maintaining a common view of member and resource states, and negotiating access to shared cluster resources are all done seamlessly.

To use the functions that define and maintain a z/VM SSI cluster, the optional IBM z/VM Single System Image Feature (VMSSI) must be licensed and enabled. (See “IBM z/VM Single System Image Feature” on page 20.)

Servicing a z/VM SSI cluster is more efficient through the use of a single service stream for the cluster. Time will be saved by the common application of service to all members of the cluster, while flexibility is provided in allowing members to roll out that service on individual schedules, avoiding an outage for the entire cluster. This allows running Linux guests to be relocated to a different member prior to a planned service outage.

For more information about the z/VM SSI environment and setting up z/VM SSI clusters, see *z/VM: CP Planning and Administration*.

Introducing live guest relocation

Live guest relocation (LGR) provides the capability for a running Linux virtual machine to be moved without disruption from one z/VM system to another within a z/VM SSI cluster.

Live guest relocation provides continuity for virtual server workloads over planned z/VM and machine outages, such as service and hardware upgrades. LGR allows applications to remain available over such outages with less impact to the application and less setup required. Verification that needed resources and machine features are available on the destination system prior to the relocation is provided. This verification may also be performed on request to assess a guest's eligibility for relocation. In an SSI cluster comprising different machine models, the architecture level presented to each guest is tailored to the set of machine features common to the member systems to which the guest may be relocated.

With the introduction of live guest relocation among members of the SSI cluster, it is increasingly important to identify the level of Linux on System z running within each guest. Unsupported levels of Linux might not operate correctly following relocation. Additional information will be provided at IBM: z/VM V6.2 Resources (www.ibm.com/vm/zvm620/).

For more information about live guest relocation, see *z/VM: CP Planning and Administration*.

Technology exploitation

z/VM provides support for the newest IBM System z hardware technologies.

Note: Support and exploitation of hardware functions depends on hardware level and availability. For detailed information on hardware capabilities, requirements, and availability, see the appropriate Hardware Announcement materials. Also refer to the hardware Preventive Service Planning (PSP) buckets at IBM Resource Link (www.ibm.com/servers/resourceink/). A Resource Link user ID and password are required.

IBM zEnterprise 196 and IBM zEnterprise 114

z/VM V6.2 supports the IBM zEnterprise™ 196 (z196) and the IBM zEnterprise 114 (z114) and enables guests to exploit selected new features. In addition, z/VM will recognize and report processing capability reduction due to customer-initiated power-saving mode or an autonomic response to environmental conditions and will reflect the change and its cause in monitor and accounting data streams. Supporting guests will also be notified of capability changes. The output of the QUERY CAPABILITY command has been enhanced to include information about such machine capacity changes.

System topology information for the z196 and z114 (and the System z10™) is collected in the CP monitor records. A configuration monitor record is used to collect the initial system topology and an event monitor record is cut whenever CP sees the system topology as changed.

CP will also collect the first 180 bytes of the Basic-machine configuration, which includes the hardware model number, in a configuration monitor record and an event monitor record.

z196 and z114 driver 93 provides a new physical-system-management time (SysMgmttime) value. z/VM will provide this information within a CP monitor record.

In addition, the Unified Resource Manager has been enhanced with firmware driver level 93 to support life cycle management of individual members within a z/VM V6.2 SSI environment. Driver 93 is required for managing z/VM 6.2. However, the Unified Resource Manager does not support relocation of managed guests.

OSX and OSM channel path

z/VM includes support for two System z channel path types (CHPIDs) to enable access to the networks that interconnect the components of the IBM zEnterprise System: zEnterprise 196 and zEnterprise 114, zEnterprise BladeCenter® Extension (zBX) with its integrated optimizers and select IBM blades, and zEnterprise Unified Resource Manager running in a Hardware Management Console (HMC). The new CHPIDs are based on the existing OSA-Express3 OSD channel path types:

- OSA-Express for zBX provides connectivity and access control to the intraensemble data network (IEDN) from the z196 or z114 to Unified Resource Manager functions using an OSA-Express3 or OSA-Express4S 10 Gigabit Ethernet feature configured as CHPID type OSX for application workloads.
- OSA-Express for Unified Resource Manager provides connectivity to the intranode management network (INMN) from the z196 or z114 to Unified

New or changed

Resource Manager functions using an OSA-Express3 1000BASE-T Ethernet feature configured as CHPID type OSM for management related functions.

These subchannels can be brought online and exploited when the z/VM image is managed by the Unified Resource Manager. Dynamic I/O support, using either CP commands or HCD (when z/VM is the controlling LPAR), and CP QUERY capabilities are provided for these two channel path identifiers.

FICON Express8S and OSA-Express4S support

IBM has introduced a new I/O drawer and new form factor I/O cards for the z196 and z114 to support a direct Peripheral Component Interconnect Express Generation 2 (PCIe Gen2) infrastructure with increased capacity, granularity, and infrastructure bandwidth, as well as increased reliability, availability, and serviceability. z/VM V6.2 supports the following features:

- FICON® Express8S features for the storage area network (SAN) provide for single mode fiber (LX) and multimode fiber (SX) environments with two channels per feature and two channel path identifiers (CHPIDs).
- OSA-Express4S features for the local area network (LAN) provide for single mode fiber (LX, LR) and multimode fiber (SX, LR) environments. The 10 GbE features have one port per feature and one CHPID. The GbE features have two ports per feature and one CHPID shared by the two ports.

XRC time stamping

z/VM exploits the Server Time Protocol (STP) facility to generate time stamps for guest and system DASD write I/O operations, allowing these I/O operations to be synchronized with those of other systems. This support allows data used by z/VM and its guests to be replicated asynchronously over long distances by IBM System Storage z/OS Global Mirror (formerly known as Extended Remote Copy, or XRC). For example, this allows z/VM to participate in a Geographically Dispersed Parallel Sysplex® Metro/z/OS Global Mirror (GDPS/MzGM) environment. This satisfies the statement of direction in the XRC Services announcement dated March 23, 2010.

HyperSwap improvements

The z/VM HyperSwap® capability has been updated to provide higher reliability, availability, and serviceability (RAS). The majority of this support is internal to the z/VM hypervisor. However, a noteworthy external change is the addition of INCLUDE and EXCLUDE options on the HYPERSWAP command to provide finer control over the z/VM missing-interrupt automatic quiesce trigger. For example, devices can be excluded from the trigger, or the trigger for a device can be made a multiple of the current z/VM MITIME setting. This support can improve the robustness of environments that have deployed z/VM with GDPS®, such as with GDPS/PPRC or GDPS/MzGM.

IBM Extended Address Volumes

The Extended Address Volumes (EAV) function on the IBM DS8000® adds support for ECKD™ volumes up to 262,668 cylinders. z/VM supports using EAV volumes for:

- Devices dedicated to guests
- Fullpack minidisks for guests that support EAV
- DDR
- FlashCopy® SE

CMS has been enhanced to support volumes up to 65,520 cylinders for its own use. CP continues to support volumes up to 65,520 cylinders for its own use.

Direct attachment of IBM XIV Storage Systems

z/VM supports the direct attachment of IBM XIV[®] Storage Systems for system use (such as paging, spooling, and IPL). This support eliminates the need for z/VM to access XIV devices through the IBM SAN Volume Controller (SVC). This support provides the ability to define system and guest volumes as emulated devices on XIV devices.

Dynamic emulated device path control

z/VM provides the ability to add and remove paths from an emulated device (representing a real SCSI device) while the device is online and possibly in use. This is especially useful for making configuration changes for z/VM system volumes that cannot be varied offline while z/VM is running.

IBM System Storage TS1140 Tape Drive (3592 Model E07)

The IBM System Storage TS1140 Tape Drive (3592 Model E07) can provide IBM System z customers with FICON or ESCON[®] attachment of the highest capacity, highest performing tape drives offered by IBM. Its corresponding Enterprise Tape, the new Controller Model C07 (or the existing Controller Model C06) can potentially benefit customers who utilize tape strictly for backup and archive, as well as customers who use tape for other applications.

The Enterprise Tape Controller 3592 Model C07 is the follow-on to the highly successful Enterprise Tape Controller 3592 Model C06 and is designed for FICON-only attached server environments where customers depend on the mainframe-class reliability, performance, and capacity that the 3592 Model C07 can provide. Attachment environments for the 3592 Model C07 with the 3592 Model E07 include:

- IBM 3584 Tape Library
- IBM 3952 F05 Tape Frame
- Standalone rack

Note: The 3592 Model E07 tape drive is not supported in the IBM 3494 Tape Library.

z/VM V5.4 and V6.1 support for the 3592 Model E07 tape drive is provided with APAR VM64979. In addition, tape library support for these drives is provided by DFSMS/VM FL221 with APAR VM65005 and its prerequisite service. Attachment of these 3592 E07 drives under the Enterprise Tape Controller Model C07 does not require any additional software maintenance.

CPU-Measurement Counter Facility Host Support

CP can collect counter data generated by the CPU-Measurement Counter Facility for MONDATA and MONSET reports and include the counter data in the CP MONITOR data stream. Collection of the counter data can be enabled or disabled with the MONITOR SAMPLE command.

Systems management

This section describes z/VM systems management improvements that help to provide self-configuring, self-managing, and self-optimization facilities.

Support for the IBM zEnterprise Unified Resource Manager

The IBM zEnterprise System (zEnterprise) is a workload-optimized multiplatform computing environment that spans and tightly integrates mainframe and distributed technologies. The IBM zEnterprise System consists of an IBM zEnterprise 196 (z196) or zEnterprise 114 (z114), the IBM zEnterprise BladeCenter Extension (zBX), and the IBM zEnterprise Unified Resource Manager. The Unified Resource Manager, running in a Hardware Management Console (HMC), controls virtual server lifecycle management, which enables attached blades to help deliver end-to-end virtualization, and supports directed and dynamic virtual server provisioning across all hypervisors in the ensemble (PR/SM™, z/VM, and PowerVM™) from a single uniform point of control.

Augmenting the existing z/VM virtual server management functions available on System z10, this support now also permits virtual servers to be created and deleted, and allows real and virtual networking resources to be managed by the Unified Resource Manager.

Within a zEnterprise, a single z196 or z114 and an optional attached zBX is called a node. A collection of one or more zEnterprise nodes that are managed as a single logical virtualized system by the Unified Resource Manager is called an ensemble. An ensemble can consist of a single z196 or z114 with no zBX attached, or 2 to 8 CPCs where at least one of the z196 or z114 machines has a zBX attached.

The intranode management network (INMN) is a 1000BASE-T Ethernet network that allows the Unified Resource Manager to manage the resources within a zEnterprise node. The intraensemble data network (IEDN) is a 10 Gigabit Ethernet network for application data communications within an ensemble. Data communications for workloads can flow over the IEDN within and between nodes of an ensemble. All of the physical and logical resources of the IEDN are configured, provisioned, and managed by the Unified Resource Manager.

The z/VM Systems Management Application Programming Interface (SMAPI), the Performance Toolkit for VM, and the z/VM Directory Maintenance Facility (DirMaint) have been enhanced to enable the Unified Resource Manager to manage the z/VM hypervisor and the virtual servers it supports. The SMAPI enhancements provide new and updated APIs and function in the following functional areas:

- System configuration file management
- User directory management
- Disk management
- Network management
- Discontiguous saved segment (DCSS) management
- Guest facilities management
- Directory parsing
- API function level interrogation
- Performance and events management
- Image management

z/VM installation has been enhanced to create a system configuration which would allow management by the Unified Resource Manager or some other SMAPI client.

For information on how to configure z/VM to participate in an ensemble, see *z/VM: CP Planning and Administration*.

Dynamic discovery of FCP disks

The SCSIDISC utility, previously provided as a sample program, is now a fully supported part of z/VM. The SCSIDISC utility provides information about the DASD logical units (LUNs) accessible through virtual FCP devices available to the virtual machine. Using DASD that implement the Small Computer System Interface (SCSI) in emulated devices and directly attached FCP devices requires specific knowledge of the storage area network topology. The SCSIDISC utility provides a means to analyze this topology with no prior knowledge of it. For more information, see *z/VM: CP Commands and Utilities Reference*.

Enhanced systems management APIs

The z/VM systems management APIs have been enhanced to integrate the z/VM SSI cluster environment and other new functionality in z/VM V6.2, along with enhancements supporting the IBM zEnterprise Unified Resource Manager when at the driver 93 level. For more information, see *z/VM: Systems Management Application Programming*.

Enhanced Directory Maintenance Facility

The Directory Maintenance Facility for z/VM, function level 620, has been enhanced to support z/VM SSI clusters along with additional enhancements supporting the IBM zEnterprise Unified Resource Manager. New capabilities include support for the enhanced directory syntax for SSI clusters, conversion of directory contents to aid transition to an SSI cluster, and assistance in adjusting the directory when adding a member to the cluster.

These enhancements include:

- Several DirMaint commands have been updated to support new IDENTITY and SUBCONFIG directory entries and the new BUILD directory statement.
- DIRM ADD command has been updated for cloning SUBCONFIG entries. In support of this, a new :SSI_VOLUMES. section has been added to the EXTENT CONTROL file.
- DIRM DIRECTORY command has been updated to add, change or delete the SSI option and additional volume labels.
- DIRM GLOBALOPTS, OPTION and SETOPTN commands have been updated to add, delete, alter or query the new CHPIDVIRTUALIZATION option.
- New DIRM SSI command has been added to prepare a source directory for use within an SSI cluster.
- New DIRM UNDOSSI command has been added to reverse changes made to a source directory by the DIRM SSI command.
- New DIRM VMRELOCATE command has been added to support the new VMRELOCATE directory statement in a user or profile entry.
- New support has been added for the satellite server machines to act as a spool file bridge for commands to and responses from the DIRMAINT machine for users on remote nodes to the DIRMAINT machine in an SSI environment.

For more information, see:

- *z/VM: Directory Maintenance Facility Tailoring and Administration Guide*
- *z/VM: Directory Maintenance Facility Commands Reference*

Enhanced Performance Toolkit

The Performance Toolkit for VM, function level 620, has been enhanced to provide the following new information in its reports:

- SSI cluster data
- CPU-Measurement Facility host counters
- Channel subsystem ID values
- Memory constraint relief values
- Multiple access ports per guest values
- Capacity-Change Reason and Capacity-Adjustment Indication values
- Ensemble membership and ID values
- Server Time Protocol (STP) Facility Configuration values
- Virtual NIC types, including IEDN and INMN
- Reorder settings for virtual machines
- New commands: FCONTROL EMERGENC and FCONTROL DUMPID

In addition, previously the MONWRITE module was generated to load at storage location X'20000', which would prevent other programs generated at the same storage location from being executed if called using the EXEC exit support on the MONWRITE command. The MONWRITE module is now generated as relocatable, increasing the flexibility and usability of its EXEC exit support by allowing these additional programs to execute.

For more information, see *z/VM: Performance Toolkit Reference*.

Network virtualization

This section describes enhancements to z/VM network virtualization.

Support for IBM zEnterprise System networks

z/VM enables access to the IBM zEnterprise System intranode management network (INMN) and intraensemble data network (IEDN) through the real and virtual networking capabilities of z/VM. As a result, the deployment and management of z/VM network topology is integrated into the zEnterprise environment and can be managed by the Unified Resource Manager. When connected to a z196 or z114 through the INMN, the resources of the zBX are viewed as a logical extension of the z196 or z114. These resources are then controlled and managed through the Hardware Management Console (HMC) and Support Element (SE) of the z196 or z114. The INMN requires an OSA-Express3 1000BASE-T, while the IEDN requires a port on two separate OSA-Express3 or OSA-Express4S 10 GbE features. Managing z/VM V6.2 from the Unified Resource Manager requires the z196 or z114 to be at driver level 93.

MROUTE upgrade

The z/VM MROUTE server has been upgraded to provide functional equivalency to z/OS V1.12 OMROUTE. z/VM MPRPOUTE now supports RFC 4191 and RFC 5175. For more information, see *z/VM: TCP/IP Planning and Customization*.

Additional IPv6 support

- Support for IPv6 has been included in:
- z/VM TCP/IP FTP client and server
 - z/VM TCP/IP SMTP client and server

- CMS NOTE and SENDFILE commands

The combination of these networking enhancements enable agencies doing business with the US Federal Government to support IPv6-based transmissions.

FTP enhancements

The following enhancements have been made to z/VM FTP support:

- User passwords are now suppressed in trace output of the FTP server.
- File sizes reported in LIST output are now correct when using LISTFORMAT UNIX and the correct transfer mode was selected.
- The SIZE column in the LIST output when in LISTFORMAT UNIX was increased to be able to show large file sizes.
- Broken SFS ALIAS files are now correctly handled.
- Wildcard support was implemented for BFS file systems (the existing restriction was removed).

Multiple access ports per guest

z/VM now provides the ability for a guest to have multiple unique access ports connected to the same virtual switch instance. Each access port has its own attributes (promiscuous and OSDSIM authority) and VLAN ID. Methods to associate guest NICs and VSWITCH ports are also provided. With this support, configuring a server to provide connectivity to multiple VLANs is easier because the guest does not need to be VLAN-aware. In addition, customers can choose which port numbers will be used. Because the mapping of systems to switch ports is known ahead of time, SNMP monitoring of the switch status is more useful, allowing virtual switches to be monitored and managed more like real switches. For more information, see *z/VM: Connectivity*.

OSA Address Table support

The z/VM TCP/IP NETSTAT command has been enhanced with an option to display information such as IP addresses and MAC addresses from the OSA Address Table (OAT) for OSA devices and virtual switch controllers. Installation of OSA/SF to satisfy diagnostic requirements in service situations is no longer necessary for OSA-Express3 or OSA-Express4S devices.

LDAP server upgrade

The z/VM TCP/IP LDAP server and utilities have been upgraded to equivalency with z/OS V1.12. For more information, see *z/VM: TCP/IP LDAP Administration Guide*

Scalability

This section describes enhancements that can help support increased workloads on z/VM.

Memory constraint relief

Frame allocation management has improved system memory scalability and decreased memory and processor constraints for larger memory sizes.

Disabling the guest page reordering process

The SET REORDER command allows the guest page reordering process to be disabled for the entire z/VM system or for specific virtual machines, helping to

New or changed

improve performance for guests with large memory footprints. The QUERY REORDER command displays the system default reorder setting or the reorder setting for a specific virtual machine.

Improved page release serialization

The page release serialization process has been improved to help reduce system overhead and increase guest throughput.

Enhanced contiguous frame coalescing

Contiguous frame coalescing algorithms have been enhanced to increase the probability and reduce the overhead of z/VM finding multiple adjacent frames of real memory when necessary, helping to increase system throughput.

Alternate method for hard limiting of scheduled system resources

z/VM supports the LIMITHARD operand on the SET SRM command. This operand selects the method the CP scheduler will use to enforce the hard limiting of scheduled system resources. This setting affects only users with absolute maximum shares defined with the SET SHARE command or the SHARE directory statement. The LIMITHARD CONSUMPTION option selects a method of hard limiting that is based on a guest's measured processor utilization. The LIMITHARD DEADLINE option selects the traditional method of hard limiting, which enforces limits based on a calculated deadline. The CONSUMPTION method, which is the default, is generally more accurate than the DEADLINE method.

Security

This section describes enhancements to the security characteristics of z/VM.

z/VM SSL server upgrade and FIPS support

The z/VM SSL server has been upgraded to use z/OS V1.12 System SSL technology. The z/VM SSL server includes support for the Federal Information Protection Standard (FIPS) 140-2. FIPS is a government required procedure for authenticating the security of a certificate database that requires signing of the database and using separate cipher suites for traffic encryption and decryption.

Additional enhancements to the z/VM SSL server improve the ability to provide concurrent secure connectivity by increasing its overall capacity while decreasing the amount of system resources it requires:

- A new threading model has been implemented to improve the handling of multiple worker threads.
- Support for multiple SSL servers allows for a pool of SSL servers to work in conjunction with a single TCP/IP stack.

ESM access control for real devices

The CP ATTACH and GIVE commands have been enhanced to enable use of an external security manager (ESM) to provide mandatory access control and discretionary access control for real devices. Who can connect to z/VM real devices can be controlled by using profiles in the new VMDEV resource class. z/VM can call the ESM for an authorization check in the VMDEV class when connecting a real device to a virtual machine for exclusive use or connecting a tape device to a virtual machine for shared use.

ESM access control for SET SECUSER and SET OBSERVER

The CP SET SECUSER command, which specifies a single console image facility (SCIF) secondary user ID for a virtual machine, and the SET OBSERVER command, which specifies a virtual machine to receive console output for another virtual machine, can now be used when ESM mandatory access control (security label checking) is active.

Enhanced RACF Security Server

The RACF Security Server for z/VM, function level 620, has been enhanced to provide the following support:

- Support for z/VM SSI clusters

When RACF is used in a z/VM SSI cluster, all member systems will share the RACF database, providing z/VM guests with a common and consistent security image.
- Support for defining protected user IDs

Protected user IDs are protected from being used to log on to the system and from being revoked through inactivity or unsuccessful attempts to access the system using incorrect passwords and password phrases.
- Mandatory access control and discretionary access control for z/VM real devices (CP ATTACH and GIVE commands)
- Mandatory access control for the CP SET SECUSER and SET OBSERVER commands

For more information, see *z/VM: RACF Security Server Security Administrator's Guide*.

Crypto Express3 support

z/VM provides guest support for the Crypto Express3 feature on supported IBM System z servers. Like the Crypto Express2, the Crypto Express3 can be defined as either a coprocessor or an accelerator and supports both encrypted key and clear key applications. In addition to the guest support, the QUERY CRYPTO command has been enhanced to provide information about users who have virtual crypto queues defined for sharing.

Protected key CPACF support

The CP Assist for Cryptographic Functions (CPACF) hardware feature supports high-performance bulk encryption using clear cryptographic keys, and also supports encrypted keys that are enabled for guest use. z/VM supports machine generated cryptographic wrapping keys and generates unique wrapping key masks for each guest virtual machine. This ensures that the encrypted keys can be used only by the guest for which the keys are intended.

EAL4+ certification for z/VM V6.1

z/VM V6.1 with the RACF Security Server optional feature, including labeled security, is being evaluated and certified for conformance to the Operating System Protection Profile (OSPP) of the Common Criteria standard for IT security, ISO/IEC 15408, at Evaluation Assurance Level 4+ (EAL4+). This satisfies the statement of direction made in the IBM Software Announcement dated July 22, 2010.

All of the z/VM V6.1 service items that were included in the EAL4+ certification are included in the z/VM V6.2 base product.

Application enablement

This section describes new or enhanced z/VM application development facilities and support provided by other products to support application development and deployment on z/VM.

z/Architecture CMS

z/Architecture CMS (z/CMS), previously shipped as a sample program, is now a fully supported part of z/VM. z/CMS runs in IBM z/Architecture 31-bit addressing mode in an ESA or XA virtual machine and enables CMS programs to use z/Architecture instructions, including those that operate on 64-bit registers, while permitting existing ESA/390 architecture CMS programs to continue to function without change.

z/CMS protects customers' existing investment in CMS-based applications and infrastructure, and enables new applications to be developed which exploit features and instructions available only in z/Architecture.

For more information, see *z/VM: CMS Planning and Administration*.

Program Management Binder upgrade

The Program Management Binder for CMS has been upgraded to the Binder level shipped with z/OS V1.12 MVS™ Program Management. The new level of the Binder includes the following enhancements:

- The COMPAT option of the BIND command includes three new suboptions: ZOSV1R10, ZOSV1R11, and ZOSV1R12.
- The RMODE option of the BIND command includes new suboptions.
- The IEWPARMS DDNAME file can be used to set Binder options.
- The Binder includes a new C/C++ API.

Language Environment upgrade

The z/VM Language Environment runtime libraries have been upgraded to z/OS V1.12 equivalency.

Support for XL C/C++ Compiler for z/VM, V1.3

z/VM V6.2 supports the IBM XL C/C++ Compiler for z/VM, V1.3, which is a z/VM-enabled version of the z/OS V1.12 XL C/C++ compiler. For more information, see *XL C/C++ for z/VM: User's Guide*.

Installation, serviceability, and packaging

This section describes changes to functions or processes to improve the installation and serviceability of z/VM and changes to how z/VM is packaged (what is included or available with z/VM).

IBM z/VM Single System Image Feature

The new IBM z/VM Single System Image Feature (VMSSI) enhances the z/VM systems management, communications, disk management, device mapping, virtual machine definition management, installation, and service functions to enable multiple z/VM systems to share and coordinate resources within a single system image (SSI) structure. This combination of enhanced functions provides the

foundation that enables live guest relocation, which is the ability for a running Linux guest to be moved from one system to another within the SSI cluster.

This optional feature must be ordered through the z/VM System Delivery Offering (SDO). For information about the license requirements for this feature, see *z/VM: License Information*.

This support satisfies the statement of direction made in the IBM Software Announcements dated July 7, 2009, October 20, 2009, and July 22, 2010.

Installation process improvements

A number of enhancements have been made to the z/VM installation process. The process supports installation of either a non-SSI (traditional) z/VM system or an SSI cluster consisting of 1-4 members. The installation procedure has been restructured so that all planning information is gathered at one time and installation is initiated with a single command. This change minimizes the chance of errors, as the planning information is validated before the actual installation begins. Additionally, customers can now specify labels for all DASD volume labels, including the system residence volume. For more information, see *z/VM: Installation Guide*.

VMSES/E enhancements

VMSES/E has been enhanced to support z/VM SSI clusters. Product service disks and inventory are shared by all member systems in an SSI cluster. Each member of the cluster has its own production disks, allowing flexibility for placing new service into production in a staged fashion. With the VMSSI feature, Linux guests can be moved from one SSI member to another ahead of most planned outages, such as those required for service and hardware upgrades. Without the relocation opportunity, an outage for the guest would be necessary.

VMSES/E has been enhanced to record the serviced objects and copy only the serviced objects to the appropriate production disks. Previously, entire disks were copied. This change is less disruptive to the running system. A new release-specific maintenance user ID is supplied for the cluster, and this user ID owns the products service disks, including the test build disks.

For more information, see *z/VM: VMSES/E Introduction and Reference*.

Common service file pool

A new common service file pool, VMPSFS, contains the service directories for any products installed to SFS. This common file pool allows members of an SSI cluster to share information stored within the file pool.

Additional changes

This section describes some additional changes included in or effective with z/VM V6.2.

Withdrawal of support

With z/VM V6.2, IBM has withdrawn support for the following functions:

- Kerberos Authentication System. This satisfies the statement of direction made in the IBM Software Announcement dated August 5, 2008, and July 22, 2010.
- Native Domain Name System (DNS) server (NAMESRV). IBM will continue to support the use of DNS servers on other platforms for TCP/IP host name

New or changed

resolution. This satisfies the statement of direction made in the IBM Software Announcements dated July 7, 2009, and July 22, 2010.

- RESOURCE function of the VMSES/E VMFINS command. The NORESOURCE option will continue to be supported on the VMFINS command and is the default. This satisfies the statement of direction made in the IBM Software Announcement dated July 22, 2010.

Shipping the zMAP agents with z/VM discontinued

The z/VM Manageability Access Point (zMAP) Agent and Platform Agent for the IBM Systems Director Server for Linux on System z, previously shipped with z/VM V6.1, are not shipped with z/VM V6.2.

Statement of direction

- **Support for HiperSockets™ Completion Queue**

IBM plans to support transferring HiperSockets messages asynchronously, in addition to the current synchronous manner on z196 and z114. This file transfer method can be especially helpful in burst situations. The Completion Queue function is designed to allow HiperSockets to transfer data synchronously if possible and asynchronously if necessary, thus combining ultra-low latency with more tolerance for traffic peaks. HiperSockets Completion Queue is planned to be supported in a future z/VM deliverable.

- **Support for HiperSockets integration with the intraensemble data network (IEDN)**

Within a zEnterprise environment, it is planned for HiperSockets to be integrated with the intraensemble data network (IEDN), extending the reach of the HiperSockets network outside of the central processor complex (CPC) to the entire ensemble, appearing as a single Layer 2 network. HiperSockets integration with the IEDN is planned to be supported in a future z/VM deliverable.

- **Support for High Performance FICON (zHPF) guest exploitation**

IBM intends to provide support for guest virtual machines utilizing the High Performance FICON for System z (zHPF) I/O protocol. z/VM will support guest operating systems that issue single track as well as multiple track zHPF I/O.

- **Stabilization of z/VM Performance Toolkit rmpfms support**

z/VM Performance Toolkit support for **rmpfms** has been stabilized and may cease to function as the underlying Linux system evolves. Support for the Linux **rmpfms** agent has already been withdrawn and is available on an as-is basis.

- **Withdrawal of HMC z/VM Tower systems management**

z/VM V6.2 is intended to be the last release supported by the HMC z/VM Tower function management support originally introduced with System z9®. The alternative implementation of virtual server and virtual resource management for z/VM V6 continues to be supported by the zEnterprise Unified Resource Manager on a zEnterprise server or later.

- **Withdrawal of support for A220 and CLAW device drivers**

z/VM V6.2 is intended to be the last release to support HYPERchannel Processor Adapter 42990007 and HYPERchannel Series A devices and HYPERchannel Series DX devices which are functioning as Series A devices.

- **Withdrawal of z/VM Dynamic Host Configuration Protocol (DHCP) server**

z/VM V6.2 is intended to be the last release to support a native Dynamic Host Configuration Protocol (DHCP) server.

- **Withdrawal of z/VM Line Printer Daemon (LPD) server support**

z/VM V6.2 is intended to be the last release to support a native Line Printer Daemon (LPD) link driver. IBM recommends using RSCS for remote printing instead, but plans to continue to support the Line Printer Remote (LPR) client for use with LPD on other platforms.

- **Withdrawal of the OVERRIDE utility and UCR support**

z/VM V6.2 is intended to be the last release to support the OVERRIDE utility and User Class Restructuring (UCR). IBM supports native CP commands for the creation and modification of privilege classes and recommends the use of these commands instead of OVERRIDE today.

- **Withdrawal of CSE**

z/VM V6.2 is intended to be the last release to support the Cross System Extensions (CSE) environment. The z/VM VMSSI feature, introduced in z/VM V6.2, includes most of the functionality in CSE, and extends it with new capabilities such as live guest relocation and autonomic minidisk cache management. VMSSI is the new technology for clustering z/VM systems.

All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

Chapter 4. Technical information

This section provides technical information about z/VM V6.2. It contains the following topics:

- “Hardware and storage requirements”
- “Device support” on page 28
- “Software (program) requirements” on page 29
- “Limitations and restrictions” on page 35
- “Performance considerations” on page 36
- “Operating systems supported as guests” on page 36
- “Other programs supported on z/VM” on page 36
- “Multicultural support” on page 37
- “Packaging and ordering information” on page 37
- “Object code only and limited source availability” on page 38
- “Integrity and security” on page 38

Note: Some technical information might have changed since the publication of this edition. There might be a later edition of this publication. For the latest product updates, you should contact your IBM Support Center or see IBM: z/VM Operating System (www.ibm.com/vm/).

Hardware and storage requirements

This section identifies the hardware and storage required to install and operate z/VM V6.2. It includes the following topics:

- “Server requirements”
- “Installation requirements”
- “Storage requirements” on page 26
- “Workstation requirement” on page 26
- “CSE hardware requirements” on page 27
- “OSA/SF hardware requirements” on page 27
- “Performance Toolkit hardware requirements” on page 27
- “TCP/IP hardware requirements” on page 28
- “VM Guest RSA-Assist Support for Linux hardware requirements” on page 28

Server requirements

z/VM V6.2 operates on the IBM zEnterprise 196, the IBM zEnterprise 114, the IBM System z10 Enterprise Class, and the IBM System z10 Business Class. See Appendix A, “IBM servers supported by z/VM,” on page 91.

Specific server facilities are also required. See IBM: z/VM V6.2 - Architecture Level Set (www.ibm.com/vm/zvm620/architecture/).

Installation requirements

z/VM V6.2 is available in the following formats:

- DASD Dump/Restore (DDR) image on magnetic tape
- Binary image on DVD

Technical information

To install z/VM:

- You need a local non-SNA 3270 terminal to use as a console, such as:
 - IBM Open Systems Adapter Integrated Communications Controller (OSA-ICC)
 - Integrated 3270 console provided through the Hardware Management Console (HMC)
 - IBM 2074 Console Support Controller
 - IBM 3174 Control Unit
- Installation from magnetic tape requires a tape drive capable of reading IBM 3590 or 3592 tape cartridges.
- Installation from DVD requires the HMC, V1.8 or later.
- Installation of an SSI cluster (1 - 4 z/VM images) requires DASD defined to emulate an IBM 3390 Model 3 or Model 9.

Note: Post-installation use of SCSI devices by guests is permitted.

- Installation of a non-SSI system to ECKD DASD requires DASD defined to emulate an IBM 3390 Model 3 or Model 9.
- Installation of a non-SSI system to FCP-attached SCSI LUNs requires IBM System Storage SCSI disks or equivalent.

Note: A z/VM system installed to SCSI disks cannot be converted to an SSI cluster.

For specific DASD volume requirements, see *z/VM: Installation Guide*.

Storage requirements

The amount of memory (storage) required to install and operate z/VM depends on several factors, such as the settings in the system configuration file, the environment where you plan to run, and the workload you want to run. Some guidelines for installation are:

- Installation first level from tape requires up to 256 MB of real memory.
- Installation first level from DVD requires that a minimum of 576 MB of real memory is assigned to the LPAR where z/VM will be installed.
- Installation second level (from tape or DVD) requires a minimum of 64 MB of virtual memory for the installation user ID on your first level system.

z/VM can support up to 256 GB of real memory and more than 1 TB of total virtual memory in use by guests. The actual amount of usable real and virtual memory is dependent on the amount of real memory in the z/VM logical partition; the hardware server model, firmware level, and configuration; and the number of guests and their workload characteristics.

If you use a nonshared copy of CMS, your virtual machine size must be at least 20 MB. If you use CMS as a named saved system (NSS), your virtual machine size must be at least 2 MB, provided that the CMS callable services libraries (CSLs) have also been saved as a saved segment.

Workstation requirement

Many z/VM administration tasks require access to a 3270-equivalent workstation. A hardware management console (HMC) is required for managing system resources if the z/VM system is part of a zEnterprise ensemble.

CSE hardware requirements

The cross system extensions (CSE) function requires certain hardware, as follows:

- The z/VM systems in the CSE complex must be connected by channel-to-channel adapters or an IBM 3088 Multisystem Channel Communications Unit (MCCU).
- DASD shared by systems in the complex must be ECKD devices. Fixed-block architecture (FBA) devices are supported by the CSE function for spooling only.

OSA/SF hardware requirements

The Open Systems Adapter Support Facility (OSA/SF) requires general disk space and a minidisk for each user ID. See the OSA/SF program directory for more information.

To use the OSA/SF graphical user interface, you also need a workstation with sufficient processing and memory capabilities.

A full list of OSA and OSA/SF hardware requirements is beyond the scope of this document. See the following documents:

- *zEnterprise System, System z10, System z9 and eServer zSeries: Open Systems Adapter-Express Customer's Guide and Reference*
- *System z9 and eServer zSeries 890 and 990: Open Systems Adapter-Express Integrated Console Controller User's Guide*
- *System z: Open Systems Adapter-Express Integrated Console Controller 3215 Support*

Performance Toolkit hardware requirements

The Performance Toolkit for VM requires general disk space and additional space on each user's A-disk, depending on the size of the log files to be written. A workstation is required to use the browser interface.

Extended color and highlighting can be used when supported by the hardware. Graphic performance data displays can be shown only on terminals with graphics capability which are also supported by the IBM Graphical Data Display Manager (GDDM®).

SSI clusters hardware requirements

The following hardware is required to support a z/VM single system image (SSI) cluster:

- The z/VM systems must be installed on ECKD DASD. Post-installation use of SCSI devices by guests is supported.
- The SSI persistent data record (PDR) must be created on a shared 3390 DASD volume. All other DASD shared by the members of the cluster can be ECKD or FBA devices.
- The z/VM systems require the following connectivity:
 - FICON channels to support 1 to 16 direct ISFC links from each member to each of the other members of the cluster
 - FICON channels to the shared DASD
 - OSA access to the same LAN segments
 - (If needed) FCP access to the same storage area networks (SANs) with the same storage access rights

TCP/IP hardware requirements

TCP/IP for z/VM requires at least one of the following:

- IBM Open Systems Adapter-Express2 (OSA-Express2)
- IBM Open Systems Adapter-Express3 (OSA-Express3)
- HYPERchannel A220 Processor Adapter 42990007

TCP/IP for z/VM supports the HYPERchannel Series A devices (and HYPERchannel Series DX devices that are functioning as Series A devices). For additional information, see the appropriate Network Systems Corporation documentation.

- IBM RS/6000[®] Channel Attachment using the Block Multiplexer Channel or ESCON Adapter
- IBM 3088 Multi-system Channel Communication Unit

TCP/IP for z/VM supports direct connection to another TCP/IP for z/VM or to z/OS using the IBM 3088.

- IBM ESCON Channel-to-Channel Adapter

TCP/IP for z/VM supports direct connection to another TCP/IP for z/VM or to z/OS using the IBM ESCON Channel-to-Channel Adapter.

- IBM FICON Channel-to-Channel Adapter

TCP/IP for z/VM supports direct connection to another TCP/IP for z/VM or to z/OS using the IBM FICON Channel-to-Channel Adapter.

- IBM HiperSockets

z/VM programs using TCP/IP can communicate using HiperSockets with other z/VM programs using TCP/IP and with guest operating systems using TCP/IP in other logical partitions.

VM Guest RSA-Assist Support for Linux hardware requirements

In order for the VM Guest RSA-Assist Support for Linux to operate, the following are required:

- Corresponding function in Linux on System z
- One of the following:
 - IBM Crypto Express2 (CEX2) configured as a coprocessor (CEX2C) or accelerator (CEX2A)
 - IBM Crypto Express3 (CEX3) configured as a coprocessor (CEX3C) or accelerator (CEX3A)

Device support

z/VM handles real devices as supported or unsupported.

Supported devices

A supported device is one of those listed in Appendix C, “IBM devices supported by z/VM,” on page 97. The use of such a device is fully supported by IBM through the service support available with z/VM. There are two types of device support:

- Supported for CP and guest use

The device is attached to the system and can be used by CP and virtual machines. CP provides system services for the device, including error recovery for guest DIAGNOSE I/O requests, and a full command set (that is, you can use all of the device-oriented CP commands for the device). Such a device can also be shared among multiple guests if appropriate (as, for example, in the case of a DASD), or it can be dedicated to the exclusive use of a single guest.

- Supported for dedicated use by a single guest
The device can be logically attached to only one guest virtual machine at any one time. The guest must be capable of running with the device. CP cannot use the device itself, and DIAGNOSE I/O services are not available to the guest for such a device.

Unsupported devices

In addition to the devices listed in Appendix C, “IBM devices supported by z/VM,” on page 97, you can also connect other devices to z/VM if they are equivalent or similar to any of the supported devices. Such other devices are referred to as unsupported devices, and proper operation of such devices with z/VM and guest operating systems is your responsibility. IBM does not guarantee that unsupported devices run properly with z/VM, and service support for such device attachments is not supplied.

If a device is absolutely equivalent to a supported device, you can define it as such and use it in any way you would use the corresponding supported device. You are responsible for the determination of equivalence.

If a device is not equivalent to any supported device but is more or less similar to one, you can use the RDEVICE system configuration statement or the SET RDEVICE command to define it as an unsupported device. An unsupported device must be dedicated to a single guest; that is, an unsupported device cannot be a system-owned or CP-owned device.

Software (program) requirements

This section identifies the additional licensed programs and other program materials required by z/VM V6.2. It includes the following topics:

- “Prerequisite licensed programs” on page 30
- “High Level Assembler” on page 30
- “Levels of the z/VM base facilities and optional features” on page 31
- “DFSMS/VM additional program requirements” on page 31
- “DirMaint additional program requirements” on page 32
- “HCD and HCM additional program requirements” on page 32
- “IOCP additional program requirements” on page 32
- “Language Environment additional program requirements” on page 32
- “OpenExtensions additional program requirements” on page 32
- “OSA/SF additional program requirements” on page 33
- “Performance Toolkit additional program requirements” on page 33
- “RACF additional program requirements” on page 33
- “RSCS additional program requirements” on page 34
- “Shared-DASD complex and CSE complex additional program requirements” on page 34
- “SSI clusters program requirements” on page 34
- “TCP/IP additional program requirements” on page 34
- “Online books additional program requirements” on page 35
- “POSIX application development program requirements” on page 35
- “Program requirements for other products” on page 35

Prerequisite licensed programs

z/VM V6.2 requires the following additional IBM licensed programs:

- Environmental Record Editing and Printing (EREP/VM), V3.5 (5654-260)

The PTFs for the following APARs have been applied to the copy of EREP/VM preinstalled with z/VM V6.2:

- APAR VM63624 provides support for large disk sizes on the DS6000™ and DS8000.
- APAR VM64367 provides support for z10 servers.
- APAR VM64475 provides support for the z10 BC server.
- APAR VM64807 provides support for the z196.
- APAR VM64928 provides support for the z114.

- Device Support Facilities: ICKDSF for VM, V1.17 (5684-042) or later.

Any fixes requested for ICKDSF should specify the PTF number of the VMSES/E formatted update. VMSES/E installable PTFs have the release number H14. 'H' indicates release 17 of ICKDSF and '14' indicates VMSES/E format.

The PTFs for the following APARs have been applied to the copy of ICKDSF preinstalled with z/VM V6.2:

- APAR PQ84848 provides support for large FBA devices.
- APAR PQ95319 provides support for large disk sizes on the DS6000.
- APAR PQ96706 provides support for large disk sizes on the DS8000.
- APAR PM11220 provides support for the direct attachment of XIV Storage Systems.
- APAR PK87899 provides support for the CP volume ownership function of z/VM SSI cluster support.

EREP/VM and ICKDSF are preinstalled with z/VM. That is, they are included on the z/VM System Image and automatically installed. Therefore, when ordering z/VM, you must be licensed for these programs or you must place orders for them to establish licenses. Both the CMS version of ICKDSF and an IPLable standalone version (5747-DS1) are supplied with z/VM.

High Level Assembler

Although not a prerequisite for z/VM V6.2, the IBM High Level Assembler, V1.6 (5696-234) or later, or an equivalent product, is required for several tasks, such as:

- Creating a new DMSTRT for system languages (NLS)
- Creating image libraries for system printers (FCBs)
- Creating GCS application segments (CONTENTS macro)
- Accessing major CMS application interfaces (CMSCALL)
- Accessing most CP application interfaces (DIAGNOSE)
- Modifying the AVS tuning control module (AGWTUN)
- Using RAS tools (such as MDCHECK, FS2SFSE, AFTCHAIN, PRINTBLK, and PRINTFST)
- Using the API for data compression
- Using the CMS Pipelines assembler macro interface
- Customizing Language Environment or compiling assembler routines used in mixed-language user applications
- Creating applications that exploit the IEEE Floating Point hardware facility
- Adding devices that cannot be sensed without using SET RDEVICE (that is, by updating HCPRIO ASSEMBLE)
- Applying local modifications to modules written in assembler

- Servicing the CP Loader (HCPLDR)
- Assembling any CP modules
- Assembling exit routines for CP, DirMaint, or RSCS
- Migrating (converting) a file containing I/O Configuration Program (IOCP) statements into an existing HCD I/O Definition File (IODF)

The High Level Assembler, V1.6, has following additional program requirements:

- APAR PK97799 provides support for the new instructions provided with the z196 and z114.

Levels of the z/VM base facilities and optional features

z/VM V6.2 supports the following base facilities and optional features:

- Base facilities:
 - HCD and HCM for z/VM, function level 620 (preinstalled)
 - Language Environment, level 620 (preinstalled)
Language Environment, level 620, is equivalent to z/OS V1.12 Language Environment.
 - OSA/SF, function level 440 (preinstalled)
 - TCP/IP for z/VM, level 620 (preinstalled)
- Optional features:
 - DFSMS/VM, function level 221 (available with the z/VM SDO V6.2)
 - Directory Maintenance Facility for z/VM, function level 620 (preinstalled disabled)
 - IBM z/VM Single System Image Feature, function level 620 (available disabled with the z/VM SDO V6.2)
 - Performance Toolkit for VM, function level 620 (preinstalled disabled)
 - RACF Security Server for z/VM, function level 620 (preinstalled disabled)
 - RSCS Networking for z/VM, function level 620 (preinstalled disabled)

Previous levels of these facilities and features are not supported by z/VM V6.2. Additional PTFs might be required to support some z/VM V6.2 functions, as indicated in this section.

DFSMS/VM additional program requirements

DFSMS/VM, function level 221, has the following additional program requirements:

- To use the Interactive Storage Management Facility (ISMF) functions:
 - Interactive System Production Facility (ISPF), V3.2 (5684-043) or later

Note: ISPF is not required if you are using only the 3495 Tape Library Dataserver support of DFSMS/VM.
- To use the minidisk management functions:
 - Directory Maintenance Facility for z/VM, function level 620
- If an external security manager is needed for authorizations:
 - RACF Security Server for z/VM, function level 620, or an equivalent product
- To access the 3495 Tape Library Dataserver for servers capable of 3490/3490E tape I/O but incapable of 3495 Tape Library Data Server control:
 - TCP/IP for z/VM, level 620
- For remote operations:

Technical information

- RSCS Networking for z/VM, function level 620
- If the compiled REXX installation-wide exit or a compiled ACS REXX exit is desired:
 - Compiler for SAA REXX/370 (5695-013)
 - Library for SAA REXX/370 (5695-014)
- Additional enhancements and support:

Note: PTFs for these APARs have been applied to the copy of DFSMS/VM available with the z/VM SDO V6.2.

- APAR VM63004 provides support for more than 10017 cylinders on an emulated 3390 Model 9 on System Storage DASD subsystems.
- APAR VM63664 provides support for very large (near 1 TB) SCSI disks on the DS8000 and DS6000.
- APAR VM63746 provides support for DFSMS/VM RMS to share tape drive devices using the MULTIUSER option of the CP ATTACH command. For example, this support allows z/VM 3590 install media to be inserted into, and z/VM to be installed from, an Automated Tape Library (ATL).
- APAR VM64062 provides support for locating encryption-capable 3592 tape drives in an Enterprise Automated Tape Library. The DFSMS/VM support provides tape-encryption capabilities for a z/VSE guest running on z/VM.
- APAR VM64458 provides support for the 3592 Model E06.
- APAR VM65005 provides support for the 3592 Model E07.

DirMaint additional program requirements

Directory Maintenance Facility for z/VM, function level 620, has no additional program requirements.

HCD and HCM additional program requirements

HCD and HCM for z/VM, function level 620, has the following additional program requirements:

- To migrate (convert) a file containing I/O configuration program (IOCP) statements into an existing HCD I/O definition file (IODF):
 - IBM High Level Assembler, V1.6 (5696-234) or later

IOCP additional program requirements

The input/output configuration program (IOCP) preinstalled with z/VM V6.2 includes support for the z196, the z114, and the z10, and has no additional program requirements.

Language Environment additional program requirements

Language Environment, level 620, has the following additional programming requirements:

- To customize Language Environment, or to compile assembler routines included in mixed-language user applications:
 - IBM High Level Assembler, V1.6 (5696-234) or later

OpenExtensions additional program requirements

To develop OpenExtensions applications, you need:

- One of the following compilers:

- IBM XL C/C++ for z/VM, V1.1 (5654-A22) or later
- IBM C/C++ for z/VM, V1.1 (5654-A22)
- IBM C for VM/ESA, V3.1 (5654-033)
- Language Environment, level 620

OSA/SF additional program requirements

OSA/SF can be accessed by a CMS user ID, a REXX call to the OSA/SF API, or a Java-based graphical user interface (GUI).

Access through the OSA/SF GUI requires:

- TCP/IP for z/VM on the host
- TCP/IP on the workstation
- Java 1.4 runtime and help on the workstation

OSA/SF, function level 440, has the following additional program requirements:

- Additional enhancements and support:

Note: PTFs for these APARs have been applied to the copy of OSA/SF preinstalled with z/VM V6.2.

- APAR OA15170 provides support for Layer 2 and for displaying information for OSN features on the server.
- APAR OA23824 provides support for the z10 EC server.
- APAR OA26286 provides support for the z10 BC server and the enhanced z10 EC server.
- APAR OA37060 is required for z196 and z114 support for the OSA-Express4S.
- To support an OSA in TCP/IP Passthru mode if access to one of the OSA's ports is being shared by more than one host program:
 - TCP/IP for z/VM, level 620

For more information about the OSA/SF program requirements for your OSA type, see the following documents:

- *zEnterprise System, System z10, System z9 and eServer zSeries: Open Systems Adapter-Express Customer's Guide and Reference*
- *System z9 and eServer zSeries 890 and 990: Open Systems Adapter-Express Integrated Console Controller User's Guide*
- *System z: Open Systems Adapter-Express Integrated Console Controller 3215 Support*

Performance Toolkit additional program requirements

Performance Toolkit for VM, function level 620, has the following additional program requirements:

- To generate Performance Toolkit graphics on the console of a virtual machine:
 - IBM Graphical Data Display Manager (GDDM), 5684-007 or 5684-168

RACF additional program requirements

RACF Security Server for z/VM, function level 620, has no additional program requirements.

RSCS additional program requirements

RSCS Networking for z/VM, function level 620, has the following additional program requirements:

- Group Control System (GCS) must be configured and activated.
- To communicate with users in an SNA network:
 - VTAM for VM/ESA®, V4.2 (5654-010)
- To communicate within an IP network:
 - TCP/IP for z/VM, level 620

Shared-DASD complex and CSE complex additional program requirements

In a shared-DASD complex, a single DirMaint server with a single source directory can manage the object directory on up to 16 z/VM systems, if the DirMaint executable code disks and source directory disk are shared among all the systems. The following is required to support the shared-DASD complex:

- DirMaint, function level 620
Within the shared-DASD complex, all z/VM systems must be licensed for the DirMaint FL620 feature and must be running the same DirMaint service level.

The following are required to support the CSE function:

- The same DirMaint requirements as for the shared-DASD complex described above.
- For shared spool and cross-system messaging:
 - VM/Pass-Through Facility, V2.1.1 (5684-100) or later

SSI clusters program requirements

To enable the creation of a z/VM single system image (SSI) cluster, the IBM z/VM Single System Image Feature (VMSSI) must be licensed for the server processors that will host the members of the SSI cluster. For more information, see *z/VM: License Information*.

TCP/IP additional program requirements

TCP/IP for z/VM, level 620, has the following additional program requirements:

- If an external security manager is needed:
 - RACF Security Server for z/VM, function level 620, or an equivalent product
- To develop programs in C:
 - One of the following compilers:
 - IBM XL C/C++ for z/VM, V1.1 (5654-A22) or later
 - IBM C/C++ for z/VM, V1.1 (5654-A22)
 - IBM C for VM/ESA, V3.1 (5654-033)
 - Language Environment, level 620
- To develop programs in Pascal:
 - IBM VS Pascal, V1.2 (5668-767), Compiler and Library

For more information about program requirements for TCP/IP for z/VM, see *z/VM: TCP/IP Planning and Customization*.

Online books additional program requirements

z/VM online books are provided in Adobe Portable Document Format (PDF) and IBM BookManager®.

PDF books

To open and view PDF books on your workstation, the Adobe Reader is required. You can also use the Adobe Reader to print PDF books or sections of PDF books. The Adobe Reader is available free from Adobe (www.adobe.com).

BookManager books

To view BookManager books on z/VM, IBM BookManager READ/VM, Version 1 (5684-062) Release 3 or later, is required. This program might have additional program requirements. See the BookManager announcement documentation.

To run BookManager READ/VM Public Library in non-370 mode, the PTF for BookManager APAR GC05366 is required.

To view and search BookManager files and bookshelves on a Windows workstation, the IBM Softcopy Reader is required. This program is included on the *IBM Online Library: z/VM Collection DVD*.

POSIX application development program requirements

Developing POSIX applications on z/VM requires one of the following compiler programs:

- IBM XL C/C++ for z/VM, V1.1 (5654-A22) or later
- IBM C/C++ for z/VM, V1.1 (5654-A22)
- IBM C for VM/ESA, V3.1 (5654-033)

Program requirements for other products

- Linux on IBM System z
Linux is not provided with z/VM. In order to operate Linux as a guest of z/VM, Linux must be obtained from Linux distribution partners. For specific function and for the most current information on Linux distributions, see IBM System z: Linux on System z - Support (www.ibm.com/systems/z/os/linux/dist.html).
- Other IBM licensed products might also have additional program requirements to run on z/VM V6.2. See the product documentation and product websites. Also see IBM: z/VM and VM-related products and features (www.ibm.com/vm/related/).

Limitations and restrictions

z/VM has the following limitations:

- z/VM operates only on zEnterprise or z10 servers. See Appendix A, “IBM servers supported by z/VM,” on page 91.
- z/VM supports only ESA/390, LINUX only, and z/VM mode LPARs.
- z/VM may run on IFL processors only if, on the IFL processors, z/VM is being used exclusively to run (1) Linux or OpenSolaris workloads and, (2) if required, z/VM applications in support of those Linux or OpenSolaris workloads.

General restriction for z/VM virtual machines

Dynamically modified channel programs cannot run in a virtual machine unless the channel program modification facility that the DIAGNOSE interface provides can be used to ensure correct operation of such channel programs.

DirMaint restrictions

The following restrictions apply to DirMaint, function level 620:

- User IDs defined for DirMaint must conform to the CMS file naming convention. Each profile name and user ID must be unique. The names \$DIRCTL\$ and \$DIRGRP\$ are reserved for DirMaint's use. Names in the form \$DIRxxx\$ should be avoided because they are reserved for possible future use by DirMaint.
- The display of menus requires a 3270-compatible terminal with a minimum of 24 lines and at least an 80 column line length.
- Exits might require tailoring for use with RACF or other external security managers.
- Although minimal, some changes are made to the system directory as a side effect of using DirMaint:
 - All addresses not already four digits are expanded to four digits with a leading zero when the source directory is converted to DirMaint's internal format. When the internal format is converted back to the CP or DIRECTXA external format, the four-digit address expansion is maintained.
 - When converted to DirMaint's internal format, most directory statements are uppercased and extraneous blanks are removed between the statement operands. When the internal format is converted back to the external format, the uppercasing and single spacing are maintained. Comments, POSIX, and optionally the ACCOUNT statements are excluded from this conversion.
 - Only one copy of NOPDATA per System Affinity group is retained. If multiple copies of NOPDATA exist in the user's directory, DirMaint strips out the extraneous copies during initialization.
- Volume IDs on MDISK directory statements defined for DirMaint must also conform to the CMS file naming convention. DirMaint does not allow the use of &SYSRES for a volume ID. The value of +VMRES is supported, with some restrictions. The use of +VMRES is reserved by CP and should not be used as the real volume label of a physical DASD volume. The value of the synonym may be changed by including the &SYSRES parameter on the DIRECTXA_OPTIONS entry in the CONFIG* DATADVH files.

For example, DIRECTXA_OPTIONS= MIXED NOMIXMSG &SYSRES VM:RES tells DirMaint, DIRECTXA, and CP to use VM:RES as the synonym rather than +VMRES.

Performance considerations

System performance depends on the hardware resources allocated to z/VM and on the level of activity within each guest. For more information, see IBM: VM Performance Resources (www.ibm.com/vm/perf/).

For assistance in understanding the performance implications for a particular situation, contact your IBM representative or your IBM Business Partner.

Operating systems supported as guests

See Appendix B, "IBM operating systems supported as guests of z/VM," on page 95.

Other programs supported on z/VM

For information about other IBM licensed programs supported on z/VM, see IBM: z/VM and VM-related products and features (www.ibm.com/vm/related/).

For information about non-IBM Solution Developer products that run on z/VM, see IBM: Vendor (ISV) Products that Run on z/VM (www.ibm.com/vm/vendor/).

Multicultural support

The system default language for z/VM is mixed-case American English. However, you can install other languages on your z/VM system, and any installed language can be set as the system default language. Each z/VM user can select to receive messages and z/VM HELP information in any installed language.

In addition to mixed-case American English, z/VM includes uppercase American English and provides the IBM-translated components, facilities, and features listed in the following table.

Table 1. IBM-translated parts of z/VM

Language	IBM-Translated Parts ¹
Japanese (KANJI)	<ul style="list-style-type: none"> • CP messages • CMS messages • REXX/VM messages • DirMaint messages • DFSMS/VM panels and messages • HELP files for CP commands and messages • HELP files for CMS commands and messages • HELP files for REXX/VM statements and messages • HELP files for DFSMS/VM messages
¹ Some items in the listed parts might not be translated.	

Packaging and ordering information

To order z/VM V6.2, order the z/VM System Delivery Offering (SDO) V6.2, program number 5741-A08. The SDO is the IBM program offering for z/VM which includes a single deliverable containing the z/VM operating system and a group of fully serviced IBM Licensed Program Products.

The z/VM SDO V6.2 package contains:

- z/VM V6.2 System Image and RSU (recommended service upgrade)
- z/VM SDO optional features and licensed products, if ordered

For more information, see IBM: SDO - System Delivery Offering for IBM z/VM (www.ibm.com/vm/sdo/).

The z/VM V6.2 System Image is available as:

- Multivolume 3390 DASD Dump Restore (DDR) image on 3590 or 3592 tape cartridges
- Binary image on DVD

The z/VM V6.2 System Image contains:

- Prerequisite IBM licensed programs:
 - EREP/VM, V3.5 (5654-260)
 - ICKDSF for VM, V1.17 (5684-042) (includes ICKDSF Standalone release 17)
- z/VM V6.2 base product:
 - AVS
 - CP

Technical information

- CMS
- Dump Viewing Facility
- GCS
- HCD and HCM for z/VM, function level 620
- Language Environment, level 620
- OSA/SF, function level 440
- REXX/VM
- TCP/IP for z/VM, level 620 (includes source)
- TSAF
- VMSES/E
- SFS and CRR file pools
- DASD Dump Restore (DDR) program
- Input/output configuration program (IOCP)
- Translated CP and CMS message repositories
- English and translated z/VM HELP files
- 3800 printer image library object code
- 3800 printer image library source
- z/VM optional features:
 - DirMaint, function level 620
 - Performance Toolkit for VM, function level 620
 - RACF Security Server for z/VM, function level 620
 - RSCS Networking for z/VM, function level 620

Notes:

1. ICKDSF and EREP/VM are not part of the z/VM product. They are separate prerequisite IBM licensed programs that are included in the z/VM System Image for convenience. Before using ICKDSF and EREP/VM you must be licensed for them, or you must place orders for them to establish licenses.
2. The DirMaint, Performance Toolkit, RACF, and RSCS optional features are included in the z/VM System Image for convenience. They are installed disabled. If you decide to use any of these features, you must place orders for them to establish licenses so that you can enable them.
3. The DFSMS/VM optional feature is not included in the z/VM System Image. If desired, DFSMS/VM must be ordered as a no-cost optional feature of the z/VM SDO.
4. The IBM z/VM Single System Image Feature is not included in the z/VM System Image. If desired, this feature must be ordered through the z/VM SDO.

Object code only and limited source availability

Some z/VM components, facilities, and features are distributed in object code only (OCO) form. No source program materials are available for the following:

- AVS
- Dump Viewing Facility
- DFSMS/VM
- Performance Toolkit

CP, CMS, DirMaint, and RACF are distributed partially in OCO form, which means that some modules will not have source program materials available but others will.

Integrity and security

This section discusses facilities of z/VM that deal with the integrity and security of the system.

Data integrity for guests

Operating system failures that occur in virtual machines do not normally affect z/VM running on the real processor. If the error is isolated to a virtual machine, only that virtual machine fails, and the user can re-IPL without affecting the work running in other virtual machines.

System integrity statement for z/VM

System integrity is an important characteristic of z/VM. IBM has implemented specific design and coding guidelines for maintaining system integrity in the development of z/VM. Procedures have also been established to make the application of these design and coding guidelines a formal part of the design and development process.

However, because it is not possible to certify that any system has perfect integrity, IBM will accept APARs that describe exposures to the system integrity of z/VM or that describe problems encountered when a program running in a virtual machine not authorized by a mechanism under the customer's control introduces an exposure to the system integrity of z/VM, as defined in the following "z/VM system integrity definition" section.

IBM will continue its efforts to enhance the integrity of z/VM and to respond promptly when exposures are identified.

z/VM system integrity definition

Unless authorized by a z/VM control program (CP) mechanism under the customer's control or a guest operating system mechanism under the customer's control, a program running in a virtual machine cannot:

- Circumvent or disable the control program real or auxiliary storage protection.
- Access a resource protected by an external security manager (ESM), such as RACF. Protected resources include virtual machines, minidisks, and terminals.
- Access a control program password-protected resource.
- Obtain control in real supervisor state or with privilege class authority or directory capabilities greater than those it was assigned.
- Circumvent the system integrity of any guest operating system that itself has system integrity as the result of an operation by any z/VM control program facility.

Real storage protection refers to the isolation of one virtual machine from another. CP accomplishes this by hardware dynamic address translation, start interpretive-execution guest storage extent limitation, and the Set Address Limit facility.

Auxiliary storage protection refers to the disk extent isolation implemented for minidisks and virtual disks through channel program translation.

Password-protected resource refers to a resource protected by CP logon passwords and minidisk passwords.

Privilege class authority refers to the authorization of a virtual machine to use specific IBM-defined or customer-defined classes of CP system functions.

Directory capabilities refer to those directory options that control functions intended to be restricted by specific assignment, such as those that permit system integrity controls to be bypassed or those not intended to be generally granted to users.

Technical information

Guest operating system refers to a control program that operates under the z/VM control program.

Customer responsibilities

While protection of the customer's data remains the customer's responsibility, data security continues to be an area of vital importance to IBM. IBM is committed to continually improving the system integrity of the z/VM environment to help customers protect their data.

Product documentation, subject to change, describes the actions that must be taken and the facilities that must be restricted to complement the system integrity support provided by z/VM. Such actions and restrictions might vary depending on the system, configuration, or environment. The customer is responsible for the selection, application, adequacy, and implementation of these actions and restrictions, and for appropriate application controls.

DirMaint system integrity

DirMaint uses standard z/VM system facilities to:

- Protect the DirMaint service machines (DIRMAINT, DATAMOVE, DIRMSATs) from subversion
- Protect files from outside interference or contamination
- Isolate users from each other and from the system
- Exploit hardware protection mechanisms
- Identify the originating user ID (and node ID), for all incoming requests
- Record auditable information

The DIRMAINT and DIRMSAT service machines require the appropriate CP privilege class to use CP commands and DIAGNOSE codes. These machines benefit from use of the OPTION D84NOPAS directory statement, and security is enhanced with the D8ONECMD FAIL LOCK directory statement. Data integrity is enhanced when the optional DATAMOVE service machines have LNKSTABL and LNKEXCLU specified on the OPTION statement in their directory entries.

Security, auditability, and control

z/VM includes several facilities to enhance or improve the security and integrity of the system:

- Each guest and CMS user runs in a unique virtual machine definition which, in combination with hardware features, prohibits one user from accessing another's data in storage (unless specifically allowed through shared segments, communication vehicles such as IUCV and APPC/VM, or ESA/XC data sharing services).
- z/VM, in combination with hardware features, provides protection against channel programs accessing another user's virtual addresses.
- A password facility provides minidisk security to control both read-only and read-write access.
- Both user ID and password checking are provided to minimize unauthorized system access.
- User class restructure provides customers with the ability to control access to commands and DIAGNOSE codes more precisely through customer-defined classes.
- Journaling is supported on z/VM. In addition, the optional RACF feature provides customers with many of these facilities, as well as other security capabilities.

- Directory control statements and system configuration file statements provide controls for certain POSIX-related functions, such as the ability to change another virtual machine's POSIX security values.
- Authorizations for z/VM guest LANs and virtual switches can be handled by the RACF feature, or by any equivalent external security manager that supports this function.
- Working together, z/VM and the RACF feature support the use of passwords longer than eight characters, called password phrases, that may contain mixed-case letters, numbers, blanks, and special characters.
- z/VM supports drive-based data encryption with IBM System Storage tape drives. z/VM supports native and guest use of the IBM Encryption Key Manager for encryption of tapes.
- z/VM supports the use of DASD volumes that reside on data encryption drives (DEDs).

Support for cryptographic facilities on IBM servers

z/VM supports use of the cryptographic facilities provided by IBM servers:

- CP Assist for Cryptographic Functions (CPACF) is supported for use by Linux and z/OS guests.
- Crypto Express2 is supported for use by Linux and z/OS guests.
- Crypto Express3 is supported for use by Linux and z/OS guests.

CMS file pool security

CMS file pools include the following features to aid data security for SFS data and BFS data stored in them:

- To access a file pool, you must be authorized (enrolled) by someone with administrator authority for that file pool, or PUBLIC must be enrolled.
- If an administrator gives you an SFS file space in a file pool, you are the only one (other than an administrator) who can create files in that file space, unless you specifically grant this authority to another user.
- You can control access to your SFS files and directories by granting and revoking authority to other users.
- Only the owner of an SFS directory or an administrator can delete the directory.
- Implicit and explicit locks prevent simultaneous updates.
- An auditing facility is available that documents:
 - Attempts to access file pool resources
 - Use of CRR recovery server operator commands and file pool server operator commands, which erase CRR and SFS log data in the intervention of CRR activity.

In addition, an external security manager (ESM), such as RACF Security Server for z/VM, can replace file pool authorizations for those objects protected by the ESM.

File pools can exploit ESM services through documented interfaces including the use of the RACROUTE programming interface.

User management is responsible for evaluation, selection and implementation of these features, for administrative procedures, and for appropriate controls in application systems and communications facilities.

TCP/IP security

TCP/IP for z/VM offers the following security facilities:

- Secure Sockets Layer (SSL) support
SSL provides secure (encrypted) communication between a remote client and a TCP/IP application server. Under SSL protocol, the application server is always authenticated. To participate in an SSL session, an application server must provide a certificate from a Certifying Authority to prove its identity.
The TCP/IP SSL server can operate in Federal Information Processing Standard (FIPS) mode.
- Exits to an ESM
TCP/IP servers can be configured to use RACF Security Server for z/VM or another ESM to:
 - Authenticate and authorize access to resources managed by the server
 - Validate client user IDs and passwords
- LDAP server and client
The Lightweight Directory Access Protocol (LDAP) server and associated client utilities provide user authentication, authorization, and auditing capabilities.

DirMaint security

DirMaint supports the z/VM security strategy:

- Access to user IDs is password controlled. DirMaint maintains user passwords, with customer choice of administration control or user control. DirMaint also supports the use of an ESM for password control.
- Access to minidisks is controlled by either passwords or explicit link authorization, as determined by the minidisk owner. Minidisk passwords are optional for controlling minidisk directory links. DirMaint also supports control of minidisk links by an ESM.
- VM system services are used to identify the originating user ID (and node ID) for all requests, whether local (those originating from within the scope of a single system or a multiple system cluster) or remote (those originating from beyond the local system cluster). By default, all local requests must be authenticated by providing the correct logon password for each DirMaint transaction. Remote requests, other than surrogate requests, never require authentication. Unless prohibited by the system administrator, users may request suspension of authentication for local requests (other than surrogate requests). Surrogate requests (whether local or remote) always require authentication, unless the system administrator specifically allows the requesting user ID to make unauthenticated surrogate requests.
- All DirMaint commands involving the DirMaint service machines (DIRMAINT, DATAMOVE, DIRMSATs) are auditable. A few DirMaint commands (CHECK, DEFAULTS, EXECDROP, GLOBALV, HELP) are completely processed in the user's virtual machine, and are therefore not auditable by DirMaint. They may, however, be auditable by an ESM or some other security program.
All messages generated by the DirMaint service machines are auditable.
An exit routine will allow customers to tailor the filtering of unnecessary audit details.

Security provided by RACF Security Server for z/VM

RACF Security Server for z/VM provides additional comprehensive security capabilities. RACF controls user access to the system, checks authorization for use of system resources, and audits the use of system resources. RACF verifies logon passwords and checks access to minidisks, data in spool files, and RSCS nodes.

You can use RACF commands to audit security-relevant events and prevent users from entering the CP DIAL and MSG commands before they log on. The events you can audit include:

- Any CP command or DIAGNOSE code (including privileged commands and DIAGNOSE codes)
- The creation, opening, and deletion of spool files
- The dumping and loading of spool files through SPXTAPE commands
- IUCV CONNECT and SEVER operations and certain VMCF functions
- APPC/VM CONNECT and SEVER operations
- The creation and deletion of logical devices.
- Connections to z/VM real devices (use of the CP ATTACH and GIVE commands).

Chapter 5. z/VM base product

This section provides an overview of the components and additional facilities included in the z/VM base product:

- “CP”
- “CMS” on page 55
- “TCP/IP for z/VM” on page 59
- “AVS” on page 62
- “Dump Viewing Facility” on page 62
- “GCS” on page 62
- “HCD and HCM for z/VM” on page 62
- “Language Environment” on page 63
- “OSA/SF” on page 63
- “REXX/VM” on page 64
- “TSAF” on page 64
- “VMSES/E” on page 64

CP

The z/VM Control Program (CP) is primarily a hypervisor and real-machine resource manager. CP provides each user with an individual working environment known as a *virtual machine*. Each virtual machine is a functional equivalent of a real system, sharing the real processor function, storage, console, and input/output (I/O) device resources.

When you first log on to z/VM, CP controls the working environment. Many of the facilities of z/VM are immediately available to you. For example, you can use CP commands to do various system management tasks. However, most of the work done on z/VM requires the Conversational Monitor System (CMS) or a guest operating system, such as z/OS, to help with data processing tasks and to manage work flow.

CP provides connectivity support that allows application programs to exchange information with each other and to access resources residing on the same z/VM system or on different z/VM systems.

Virtual machine modes

z/VM supports the following virtual machine modes (architectures):

- **ESA** virtual machines process according to ESA/390 architecture (31-bit). ESA virtual machines are also capable of processing according to z/Architecture (64-bit) if switched into that mode by a guest operating system.
- **XA** virtual machines are supported for compatibility and are functionally equivalent to ESA virtual machines. Some CMS applications might require CMS to be running in an XA virtual machine.
- **XC** virtual machines process according to ESA/XC architecture.

Notes:

1. Although 370-mode (IBM System/370 architecture) virtual machines are not supported, the 370 Accommodation Facility allows many CMS applications

Base Product

written for 370 virtual machines to run in ESA, XA, and XC virtual machines. For more information about the 370 Accommodation Facility, see *z/VM: CP Programming Services*.

2. All z/VM virtual machines are type V=V (virtual=virtual), which means the virtual machine's guest real storage is paged into and out of host real storage.

Architecture compatibility

Because ESA/390 architecture is upwardly compatible from ESA/370 architecture, and ESA/370 architecture is upwardly compatible from 370-XA architecture, in most cases an ESA/370 or 370-XA guest can run in an ESA/390 virtual machine. CP does not distinguish between the 370-XA, ESA/370, and ESA/390 architectures.

z/Architecture support

A guest in an ESA or XA virtual machine might have the capability to issue an instruction to switch the virtual machine from ESA/390 mode to z/Architecture mode. z/Architecture guests can use z/Architecture facilities in z/VM, such as 64-bit arithmetic and addressing.

ESA/XC support

z/VM uses extensions to the interpretive-execution facility to provide the Enterprise Systems Architecture/Extended Configuration (ESA/XC) virtual machine architecture. ESA/XC is an architecture unique to virtual machines. Because it exists mainly to provide services to application programs in virtual machines, ESA/XC architecture does not have a native-mode equivalent.

ESA/XC architecture lets virtual machines share multiple data spaces. An XC virtual machine can access one or more data spaces of another virtual machine if so authorized. This is extremely useful for applications that require one or more virtual machines to serve many users.

For more information, see *z/VM: Enterprise Systems Architecture/Extended Configuration Principles of Operation*.

Cross system extensions

The cross system extensions (CSE) function allows two, three, or four z/VM systems to be linked together in a cluster that provides cross-system functions:

- Cross-system link, which extends the existing CP minidisk access protocols (for minidisk linking) across the CSE cluster
- Cross-system spool, which extends CP spooling
- Cross-system message, query, and other commands
- Support for printer spooling

Note: The IBM z/VM Single System Image Feature offers a more comprehensive clustering solution than CSE and adds functionality not available in CSE, such as multisystem installation, single point of maintenance, autonomic minidisk cache management, and virtual server mobility. For more information, see “IBM z/VM Single System Image Feature” on page 66.

Guest LANs and virtual switches

z/VM supports a network of virtual adapters connecting guests within a z/VM system:

- Virtual adapters —

CP offers a virtual network interface card (NIC) that simulates either a HiperSockets device or an OSA-Express QDIO device. The guest can operate the virtual adapter using the same software that would be used to drive the equivalent hardware.

- Connectivity for virtual adapters —
CP provides commands that enable z/VM users to connect virtual network adapters to an emulated local area network (LAN) segment, known as a guest LAN. Guests connected by a guest LAN can communicate using the same software that would be used to communicate over an equivalent physical LAN segment.

A guest LAN can be bridged through a z/VM virtual Internet Protocol (IP) switch, known as a virtual switch, to an associated real LAN connected by an OSA-Express QDIO adapter. A virtual switch allows connectivity to external LAN segments without requiring a router.

For more information, see *z/VM: Connectivity*.

Inter-System Facility for Communications

The Inter-System Facility for Communications (ISFC) enables channel-to-channel (CTC) communication between programs written to the APPC/VM, CPI Communications, or IUCV programming interfaces. A group of interconnected z/VM systems that use ISFC to communicate with each other is known as a communication services (CS) collection. Programs on systems in the CS collection can use ISFC to access, manage, and share resources defined in the collection. ISFC also enables programs in a CS collection to communicate with APPC programs on systems in the SNA network.

VM Dump Tool

The VM Dump Tool assists in analyzing dump data from a dump file created by the DUMpload utility. The VM Dump Tool provides a variety of subcommands and macros that allow you to display, locate, and format dump data interactively. This tool can process CP standalone dumps, CP ABEND dumps, and virtual machine dumps of a CP system. For more information, see *z/VM: VM Dump Tool*.

Note: The Dump Viewing Facility interactively diagnoses virtual machine dumps. See “Dump Viewing Facility” on page 62.

Support for hardware facilities, features, and architectures

This section lists *some* of the hardware facilities, features, and architectures that CP supports.

Notes:

1. Support for hardware depends on hardware level and availability. See Chapter 3, “What is new or changed in z/VM V6.2,” on page 9. For detailed information on hardware capabilities, requirements, and availability, see the appropriate Hardware Announcements.
2. In some cases, the CP support is available only to guests.

Collaborative Memory Management Assist

CP support for the Collaborative Memory Management Assist (CMMA) machine feature, in conjunction with CMMA exploitation in z/Architecture guest operating systems such as Linux on System z, allows the CP host and its guests to communicate attributes for specific 4 KB blocks of guest memory. This exchange of

information allows both host and guest to optimize their use and management of memory. For more information, see *z/VM: Performance*.

Concurrent-sense facility

CP provides guest support for the concurrent-sense facility, which is designed for use on channel paths that support the ESCON I/O interface. This facility allows the channel subsystem to present I/O status information and sense information together, in one I/O operation. This eliminates the wait for sense information whenever status information is received.

CPU-Measurement Counter Facility

CP can collect counter data generated by the CPU-Measurement Counter Facility and include the counter data in the CP MONITOR data stream. Collection of the counter data can be enabled or disabled with the MONITOR SAMPLE command. IBM might request that a customer use MONITOR SAMPLE to collect counter data for problem analysis, benchmarking, or capacity planning purposes.

Cryptographic facilities

CP supports guest use of server cryptographic facilities (where available):

- Crypto Express2 (CEX2), which can be configured as a coprocessor (CEX2C) or an accelerator (CEX2A)
- Crypto Express3 (CEX3), which can be configured as a coprocessor (CEX3C) or an accelerator (CEX3A)
- CP Assist for Cryptographic Functions (CPACF), which supports high-performance bulk encryption using clear keys or encrypted keys.

z/VM provides clear-key RSA support for a considerably large number of Linux guests, which helps enable hardware SSL acceleration on the servers.

Dynamic storage reconfiguration

The dynamic storage reconfiguration (DSR) capability of the server allows the amount of main storage available for use in an LPAR to be changed while the LPAR is active. z/VM exploits this system architecture for its own use and also virtualizes it for guest use. LPAR customization can assign storage to an LPAR as initial or reserved. The reserved storage in the LPAR is further separated by z/VM into standby and reserved. z/VM can dynamically increase the size of its main storage by bringing designated amounts of standby storage online, and z/VM guests that support DSR can dynamically increase or decrease their amounts of main storage.

Enterprise Systems Connection Architecture

Enterprise Systems Connection Architecture (ESCON) is built around fiber optic transmission technology. Fiber optic cables reduce cable bulk, allow for increased distance between the processor and attached devices, and improve data transfer rates.

Expanded storage

Expanded storage is an optional integrated high-speed storage facility that allows for the rapid transfer of 4 KB blocks between itself and real storage.

Extended Channel Measurement Data Support

Extended Channel Measurement Data Support (ECMDS) provides an extended I/O-measurement block for each subchannel and an extended measurement word at the completion of each I/O operation.

Extended-TOD-Clock Facility

The Extended-TOD-Clock Facility is a hardware facility that provides a 128-bit time-of-day (TOD) clock. CP supports the use of the Extended-TOD-Clock Facility from virtual machines.

Fibre Channel Protocol

The Fibre Channel Protocol (FCP) hardware channel allows a server to connect to a fibre-channel (FC) fabric. This support allows z/VM system facilities and guest operating systems to access selected devices on small computer system interface (SCSI) controllers connected to the fabric.

Fibre Connection channels

Fibre Connection (FICON) is a fiber optic I/O architecture that coexists with ESCON equipment and infrastructure, but offers improved link performance and enhanced distance connectivity. Each FICON channel provides the equivalent of eight ESCON channels.

- A FICON Express channel can provide up to 170 MBps.
- A FICON Express2 channel is designed to deliver up to 270 MBps.
- A FICON Express4 channel is designed to deliver up to 4 Gbps.
- A FICON Express8 or Express8S channel is designed to deliver up to 8 Gbps.

FICON Channel-to-Channel Adapter

CP supports guest use of the FICON Channel-to-Channel Adapter (FICON CTCA), which provides channel-to-channel communications between servers. This enables more reliable and higher bandwidth host-to-host communication than is available with ESCON channels.

FlashCopy

A native CP user or a guest can initiate the FlashCopy feature on IBM System Storage and IBM TotalStorage DASD subsystems to make an instant copy of a disk or data set from a source device to a target device. CP also supports certain FlashCopy Version 2 functions, which provide additional flexibility.

z/VM supports the FlashCopy Space-Efficient feature of the DS8000 by providing space-efficient volumes for temporary data duplication for backup or test provision scenarios, while minimizing the amount of required backing storage to accomplish these tasks.

Guest ASCII console

z/VM provides guest access to the system ASCII console, which is a facility presented by the Hardware Management Console (HMC). Dedicating the system ASCII console to a Linux guest can facilitate recovery of the guest during an emergency situation, using an environment that provides tools (such as vi and emacs) that are familiar to Linux support staff. This can be particularly useful when normal network access to a guest operating system is not available. Use of the console can be transferred from guest to guest as required.

HiperSockets

CP supports the z/Architecture HiperSockets function for high-speed TCP/IP communication among virtual machines and logical partitions (LPARs) within the same server. This function uses an adaptation of the queued direct I/O (QDIO) architecture.

CP supports HiperSockets for use by guest operating systems and by the TCP/IP for z/VM server virtual machine. z/VM programs using traditional TCP/IP socket

connections can communicate through HiperSockets with other z/VM programs, guest operating systems, and other logical partitions that are also using TCP/IP.

HiperSockets channels can be defined as spanning channels, which can allow them to be configured as part of multiple logical channel subsystems and to be transparently shared among any subset or all of the configured logical partitions. For example, spanned HiperSockets channels can help enable communication among Linux guests running on z/VM images in different logical partitions, when the z/VM images are using different logical channel subsystems.

IEEE Floating Point

CP supports guest use of the IEEE Floating Point hardware. This support allows multiple levels of guests to use basic floating point extensions, floating point support extensions, hexadecimal floating point extensions, and binary floating point.

Integrated 3270 console

CP supports this device, which is provided through a Hardware Management Console, to be used as the system operator console for installing and maintaining z/VM.

Integrated systems management

z/VM integrated systems management support uses the Hardware Management Console (HMC) to help enable administration of z/VM guests without having to establish additional network connections and reducing complex configuration of the system. The HMC will automatically detect z/VM images and provide integrated GUI-based basic management of z/VM guests. The z/VM integrated systems management capability supports the following image management functions: activate, deactivate, and display guest status.

z/VM provides a local management interface to allow basic z/VM systems management functions to be managed from the HMC. The z/VM SCLP system service (*SCLP) receives and transmits HMC events. A proxy server directs requests to the Systems Management API server to perform the desired function and send the results back to the HMC using *SCLP. The VM event system service (*VMEVENT) provides notification about certain events that occur in the VM system, such as some virtual machine status changes. The proxy server receives notification when these events occur and reports them to the HMC through the *SCLP interface.

Additional systems management enhancements are provided by the HMC and Support Element (SE) 2.10.1, exploiting the z/VM System Management APIs to allow selected virtual resources to be defined and managed.

Also see “zEnterprise Unified Resource Manager” on page 54.

Logical channel subsystem

The logical channel subsystem (LCSS) structure on the server allows the definition of more than one channel subsystem, providing channel-path and subchannel controls for configuring channel-subsystem images. Each channel-subsystem image can be configured with up to 256 channel paths, and each logical partition has access to one channel-subsystem image. z/VM dynamic-I/O configuration allows channel paths, control units, and devices to be dynamically added, changed, and deleted in multiple LCSSs when z/VM is running on a server with the applicable hardware support.

MIDAWs

z/VM supports guest use of Modified Indirect Data Address Words (MIDAWs). This server feature can allow more flexibility and performance in certain channel programs as an alternative to data-chained channel-command words (CCWs). MIDAWs accommodate noncontiguous data areas that cannot be handled by the predecessor indirect-data-address words (IDAWs). z/VM support for guest use of MIDAWs can allow operating systems such as z/OS to use this new aspect of z/Architecture without regard to whether the operating systems are running in a logical partition or a virtual machine. This allows guest operating systems to exercise their code-paths just as they would on the real machine during, for example, preproduction testing of z/OS systems. Likewise, the provision of the function in a virtual machine allows guest operating systems to benefit from the real machine's added-value function just as though the guests were running directly on the machine.

Move-Page Facility

CP supports the Move-Page Facility for ESA/XC and z/Architecture guests. A page of data can be moved from main storage to main storage, from main storage to expanded storage, or from expanded storage to main storage.

N_Port ID virtualization

N_Port ID virtualization (NPV) for FCP channels is designed to allow the sharing of a single physical FCP channel among multiple operating-system images, whether in LPARs or as z/VM guests in virtual machines. z/VM exploits this hardware capability to define multiple virtual FCP channels, each with its own worldwide-unique Fibre Channel port name (WWPN). By assigning different hardware-defined virtual port names to different guests, the guests can use the virtual FCP channels as if they were dedicated physical FCP channels.

NPV is applicable to all of the FICON features supported on the server. NPV is supported for CP and guest use.

Open Systems Adapter

CP supports guest use of the Open Systems Adapter (OSA) server features, which provide connectivity to various types of networks.

CP also supports:

- OSA-Express Integrated Console Controller (OSA-ICC), a function of the OSA-Express 1000BASE-T Ethernet feature that integrates 3270 emulation for console session connections into the server
- Open Systems Adapter for the Network Control Program (OSA NCP), a function of the OSA-Express2 Gigabit Ethernet and 1000BASE-T Ethernet features that provides channel connectivity from System z operating systems to the IBM Communication Controller for Linux (CCL) on the server.
- Virtual-switch-controlled link aggregation (IEEE 802.3ad), which allows you to dedicate an OSA port to z/VM when the port is participating in an aggregated group. Link aggregation (trunking) is designed to allow you to combine multiple physical OSA-Express2 and OSA-Express3 ports into a single logical link for increased throughput and for nondisruptive failover in the event that a port becomes unavailable.
- All four ports on the OSA-Express3 Gigabit Ethernet (GbE) and 1000BASE-T Ethernet features on the z10 EC. There are two PCI Express (PCI-E) adapters per feature, with two ports per adapter, for a total of four ports per feature. The OSA-Express3-2P GbE and 1000BASE-T features on the z10 BC provide two

ports per feature. Activating all of the available ports on an OSA-Express3 feature provides more physical connectivity to service the network and reduces the number of required resources.

- Port isolation security, which provides the ability to restrict guest-to-guest communications within a virtual switch, and OSA-Express QDIO data connection isolation, which provides the ability to isolate QDIO data connection communications on a shared OSA port in support of multitier security zones.

z/VM includes the Open Systems Adapter Support Facility (OSA/SF) for customizing an OSA's modes of operation. For more information, see "OSA/SF" on page 63.

Parallel Access Volumes

CP provides guest support for the Parallel Access Volumes (PAV) feature on IBM DASD subsystems, which allows configuration of base and alias DASD volumes. The alias volumes are logical volumes that map the physical space occupied by the base. However, each alias volume has a unique subchannel ID, which permits concurrent I/O to a base volume and all of its associated alias volumes.

PAVs enable a single server and applicable storage controllers to simultaneously process multiple I/O operations to the same logical volume, which can help to significantly improve I/O response times by reducing device queueing delays.

CP supports PAVs as linkable minidisks for guests such as z/OS that exploit the PAV architecture. In addition, CP provides the potential benefit of PAVs for I/O issued to minidisks owned or shared by guests that do not support native exploitation of PAVs, such as CMS.

CP also provides guest support for the HyperPAV function, which potentially reduces the number of alias-device addresses needed for parallel I/O operations, because HyperPAVs are dynamically bound to a base device for each I/O operation instead of being bound statically like basic PAVs. CP supports HyperPAVs as linkable minidisks for guests, such as z/OS, that exploit the HyperPAV architecture. In addition, CP transparently provides the potential benefit of HyperPAVs for I/O issued to minidisks owned or shared by guests that do not specifically exploit HyperPAVs, such as Linux and CMS.

Peer-to-Peer Remote Copy

CP supports the Peer-to-Peer Remote Copy (PPRC) feature on System Storage and TotalStorage DASD subsystems, a synchronous copy option that provides realtime mirroring of logical volumes within a device or to another device.

CP supports guest use of the PPRC Extended Distance (PPRC-XD) feature if the guest also supports PPRC-XD. PPRC-XD is an asynchronous copy option that can operate at distances well beyond the capability of PPRC. Native z/VM support for PPRC-XD is provided by ICKDSF, V1.17, with ICKDSF running in a CMS virtual machine.

CP also supports guest use of PPRC Version 2, including the Continuous Protection Option, which can offer an Asynchronous Cascading solution. Asynchronous Cascading provides a two-volume synchronous PPRC implementation with a non-synchronous third volume serving as a remote backup device that can provide a multisite, long distance solution.

Queued Direct I/O

CP supports guest use of Queued Direct I/O (QDIO), which allows a program to directly exchange data with an I/O device without performing traditional I/O

instructions. To exchange data, both the I/O device and the program reference main storage directly through a set of data queues.

The QDIO architecture is used by OSA-Express, HiperSockets, and FCP channels.

In HiperSockets, the QDIO architecture includes a type of high-performance I/O interruption known as an adapter interruption. Adapter interruptions are also used by OSA-Express and FCP channels.

In addition, the server can provide a performance assist for the virtualization of adapter interruptions being given to guest operating systems running on z/VM. This hardware performance assist can be used by guests that support QDIO.

The server also provides QDIO Enhanced Buffer-State Management (QEBSM) and Host Page-Management Assist (HPMA), which can allow a cooperating guest operating system to initiate QDIO operations directly to the applicable channel, without interception by z/VM, thereby helping to provide additional performance improvements.

SAN Volume Controller

The IBM System Storage SAN Volume Controller is implemented in the storage network between the hosts and the physical storage, providing a common volume-management “utility” shared by all hosts. z/VM supports the SAN Volume Controller through the generic SCSI device driver of z/VM. z/VM support for the SAN Volume Controller allows CP and guest operating systems that use SCSI devices (such as Linux on System z) to access IBM System Storage disk subsystems, as well as disk subsystems from other manufacturers supported by the SAN Volume Controller. The SAN Volume Controller can be used to provide SCSI devices as emulated FBA devices for use by CP and guest operating systems. Guests can also use SCSI devices accessed through the SAN Volume Controller by dedicated FCP subchannels.

Server Time Protocol

z/VM exploits the Server Time Protocol (STP) facility to generate time stamps for guest and system DASD write I/O operations, allowing these I/O operations to be synchronized with those of other systems. This support allows data used by z/VM and its guests to be replicated asynchronously over long distances by IBM System Storage z/OS Global Mirror (formerly known as Extended Remote Copy, or XRC). For example, this allows z/VM to participate in a Geographically Dispersed Parallel Sysplex Metro/z/OS Global Mirror (GDPS/MzGM) environment.

Specialty processors

z/VM provides guest support for IBM mainframe specialty processors:

IBM Integrated Facility for Linux (IFL)

IFL processors provide additional processing capacity for Linux or OpenSolaris workloads. IFLs can be allocated only in a LINUX only or z/VM mode LPAR.

IBM Internal Coupling Facility (ICF)

ICF processors allow multiple LPARs running z/OS to manage data and distribute workload in a Parallel Sysplex clustered system.

IBM System z Application Assist Processor (zAAP)

zAAPs provide an economical Java execution environment under z/OS.

IBM System z Integrated Information Processor (zIIP)

zIIPs are designed to help improve resource optimization and lower the cost

for eligible workloads. z/OS exploits zIIPs to offload software system overhead from standard central processors (CPs). This includes certain DB2® processing, enhancing the role of the mainframe as the data hub of the enterprise.

You can define a mix of these specialty processors and CPs in the same z/VM-mode logical partition.

z/VM provides two types of guest support for specialty processors:

- Simulation support
z/VM guest virtual machines can create virtual specialty processors on processor models that support the same types of specialty processor but do not necessarily have them installed. Virtual specialty processors are dispatched on real CPs. Simulating specialty processors provides a test platform for z/VM guests to exploit mixed-processor configurations. For example, this allows users to assess the operational and CPU utilization implications of configuring a z/OS system with zIIPs or zAAPs without requiring the real specialty processor hardware.
- Virtualization support
z/VM can create virtual specialty processors for virtual machines by dispatching the virtual processors on corresponding specialty processors of the same type in the real configuration. For example, guest support for zAAPs and zIIPs might help improve your total cost of ownership by allowing available zAAP and zIIP capacity not being used by z/OS LPARs to be allocated to a z/VM LPAR hosting z/OS guests running Java and DB2.

For more information about specialty processor support, see *z/VM: Running Guest Operating Systems*.

zEnterprise Unified Resource Manager

The IBM zEnterprise Unified Resource Manager runs in a Hardware Management Console (HMC) in an IBM zEnterprise System. The IBM zEnterprise System (zEnterprise) is a multiplatform computing environment that consists of an IBM zEnterprise 196 (z196) or zEnterprise 114 (z114), the IBM zEnterprise BladeCenter Extension (zBX), and the IBM zEnterprise Unified Resource Manager.

Within a zEnterprise, a single z196 or z114 and an optional attached zBX is called a node. A collection of one or more zEnterprise nodes that are managed as a single logical virtualized system by the Unified Resource Manager is called an ensemble.

The Unified Resource Manager controls virtual server lifecycle management and supports directed and dynamic virtual server provisioning across all hypervisors in the ensemble (PR/SM, z/VM, and Power® VM) from a single uniform point of control. The zEnterprise intranode management network (INMN) allows the Unified Resource Manager to manage the resources within a zEnterprise node. The zEnterprise intraensemble data network (IEDN) is used for application data communications within an ensemble.

z/VM can be configured to participate in an ensemble. z/VM enables access to the IEDN and INMN through the real and virtual networking capabilities of z/VM. As a result, the deployment and management of z/VM network topology is integrated into the zEnterprise environment and can be managed by the Unified Resource Manager. Through the z/VM Systems Management Application Programming Interface (SMAPI), the Performance Toolkit for VM, and the z/VM Directory Maintenance Facility (DirMaint), z/VM enables the Unified Resource Manager to manage the z/VM hypervisor and the virtual servers it supports.

For information about configuring z/VM to participate in an ensemble, see *z/VM: CP Planning and Administration*.

CMS

The z/VM Conversational Monitor System (CMS) provides a high-capacity application environment that can support large numbers of interactive users. CMS can help you perform a wide variety of tasks:

- Write, test, and debug application programs for use on CMS or guest systems
- Run application programs developed on CMS or guest systems
- Create and edit data files
- Process jobs in batch mode
- Share data between CMS and guest systems
- Communicate with other system users

For general information about using CMS, see *z/VM: CMS User's Guide*.

Versions of CMS

z/VM provides two versions of CMS:

- ESA/390 CMS

ESA/390 CMS (generally referred to simply as CMS) runs in an ESA/390 architecture (ESA or XA mode) or ESA/XC architecture (XC mode) virtual machine.

ESA/390 CMS is supplied as a predefined named saved system called CMS and as an IPLable nucleus on the MAINT 190 minidisk.

- z/Architecture CMS (z/CMS)

z/CMS runs in z/Architecture 31-bit addressing mode in an ESA or XA virtual machine and enables CMS programs to use z/Architecture instructions, including those that operate on 64-bit registers, while permitting existing ESA/390 architecture CMS programs to continue to function without change. Although it does not directly exploit storage above 2 GB, z/CMS can be IPLed in a virtual machine with more than 2 GB of storage, and programs running on z/CMS can allocate storage above 2 GB.

z/CMS is supplied as a predefined named saved system called ZCMS and as an IPLable nucleus on the MAINT 990 minidisk.

For more information about z/CMS, see *z/VM: CMS Planning and Administration*.

Unless otherwise indicated in the z/VM publications, “CMS” means either version, and descriptions of CMS functions apply to both CMS and z/CMS.

Shared File System

The Shared File System (SFS) is an extension of the CMS file system that offers you additional file management and file sharing functions:

- Files are stored in file pools.
- A user can be given an amount of file space in a file pool.
- The files in a file space are organized in directories.
- A file can be placed in more than one directory.
- Users can grant each other authorities on files or directories.
- Multiple users can have concurrent access to the same file or directory.
- Locks on files and directories ensure data integrity among multiple users.

Base Product

- You can share files and directories with users in other systems.

A file pool is a collection of minidisks assigned to a single virtual machine called a *file pool server machine*. Because the minidisks in the file pool are shared by many users, using SFS can save DASD space. Certain SFS directories can be placed into VM data spaces, providing an additional DASD savings. Using VM data spaces might also provide a performance improvement.

For more information, see *z/VM: CMS File Pool Planning, Administration, and Operation*.

For information about how DFSMS/VM can provide storage management functions for file pools, see “Data Facility Storage Management Subsystem for VM” on page 65.

CMS application programming

CMS supports a wide range of high-level languages and application environments. CMS also provides many special application programming facilities, including:

- Systems management API
- ESA/XC support
- CMS Pipelines
- Callable services library (CSL)
- Multitasking services
- OpenExtensions
- Program Management Binder
- Reusable Server Kernel
- Assembler macros and functions
- OS/MVS simulation
- DOS/VSE support

Note: Although IBM VSE/VSAM, V6.1 (5686-081), was withdrawn from marketing on September 30, 2005, CMS DOS/BAM will continue to provide its current level of function in support of DOS simulation and VSE/VSAM.

For general information about CMS application programming facilities, see *z/VM: CMS Application Development Guide*.

Systems management API

The z/VM systems management application programming interface (API) provides a standard, platform-independent client interface that reduces the amount of VM-specific programming skills required to manage resources for virtual systems (guests). The API includes functions for creating new virtual images, allocating and managing their resources, and changing their configurations. They can be used to activate and deactivate images individually or in groups. Security and directory management functions are also provided.

The API functions are invoked by a client through a sockets interface. For more information, see *z/VM: Systems Management Application Programming*.

CMS Pipelines

CMS Pipelines provides a rich and efficient set of functions that you can use to solve large problems by breaking them up into smaller, less complex programs. These smaller programs are called *stages*. Many stages are included with CMS Pipelines. Some stages read data from system sources, such as disk files, tape files, or the results of z/VM commands. Other stages filter and refine that data in some way. You can combine many stages within a single *pipeline* to create the results you need. You can also write your own stages. For more information, see *z/VM: CMS Pipelines User's Guide*.

CMS application multitasking

CMS application multitasking services provide an execution environment for high-performance applications and servers. With CMS multitasking, an application can divide itself into multiple units of execution and provide the ability for these units, called threads, to run on multiple CPUs simultaneously. The multitasking facilities are available only at the application programming level. The CMS user still runs one application at a time, but these applications can split themselves into multiple execution units, or threads. These multitasking facilities allow applications to harness the power of the underlying multiprocessor complex and to overlap operations to achieve high performance. For more information, see *z/VM: CMS Application Multitasking*.

OpenExtensions

OpenExtensions includes the z/VM implementation of four POSIX standards:

- POSIX 1003.1 (known as POSIX.1) - System Interfaces
- POSIX 1003.1a (known as POSIX.1a) - Extensions to POSIX.1
- POSIX 1003.1c (known as POSIX.1c) - Threads
- POSIX 1003.2 (known as POSIX.2) - Shell and Utilities

The POSIX.1, POSIX.1a, and POSIX.1c interfaces are provided as C/C++ library routines in the C/C++ runtime library included with Language Environment. For programs written in other languages, a language-neutral version of the POSIX functions is provided as a set of CMS callable services library (CSL) routines. These CSL routines are called by the C/C++ runtime routines to provide the functions, but they are also available to other applications. The CSL routines can be invoked as REXX functions through a REXX subcommand environment, ADDRESS OPENVM.

OpenExtensions includes a POSIX-compliant file system known as the Byte File System (BFS). BFS is a companion to the CMS Shared File System (SFS) that provides a byte-stream view of files. BFS allows data to be organized and used in a UNIX style and format.

Like SFS files, BFS files are organized in a hierarchical directory structure and stored in CMS file pools. While supporting the POSIX file system functions and rules, BFS also takes advantage of administration and system management facilities that it shares with SFS. These include space allocation, backup, and DFSMS/VM file migration, as well as other administrative functions.

CMS provides a set of commands, known as the OPENVM commands, that allow users to manage their BFS directories and files and control their related permission and ownership attributes. CMS Pipelines additionally provides the ability to use BFS from pipeline programs.

Base Product

The OpenExtensions shell and utilities provide a UNIX-like interactive environment in support of the POSIX application environment and provide tools that aid in program development and in porting applications from other open systems. Users of the shell environment have access to both the shell command set (built-in commands and utilities) and the full CP and CMS command sets, as well as both OpenExtensions and non-OpenExtensions applications. Each of the OpenExtensions POSIX.2 utilities additionally conforms to the X/Open Portability Guide, issue 4 (XPG4) for Commands and Utilities.

For an expanded introduction to OpenExtensions, see *z/VM: CMS Application Development Guide*.

Program Management Binder for CMS

The Program Management Binder for CMS (CMS Binder) is a CMS-enabled version of the z/OS MVS Program Management Binder. The CMS Binder converts the output of language translators and compilers into an executable program unit that can either be read directly into virtual storage for execution or stored in a program library. For more information, see *z/VM: Program Management Binder for CMS*.

Reusable Server Kernel

The Reusable Server Kernel enables vendors and application programmers to write multithreaded server programs that heavily exploit VM technologies. These servers can be constructed without knowledge of data transport mechanisms, multithreaded APIs, or I/O performance boosters and without reinventing API suites necessary in one server after another.

The Reusable Server Kernel is an “empty” server program that server writers can use as a starting point for developing and executing server programs on CMS. The Reusable Server Kernel consists of a text library of routines and a macro library of function prototypes and constant definitions. To construct an actual server program, the server author attaches application-specific code to a set of interfaces in the Reusable Server Kernel.

For more information, see *z/VM: Reusable Server Kernel Programmer's Guide and Reference*.

Virtual Machine Resource Manager

The Virtual Machine Resource Manager (VMRM) provides functions to:

- Manage guest performance
A service virtual machine (SVM) accepts customer-defined workload definitions, goal specifications, and associations between them. The SVM then adjusts virtual machine CPU and I/O performance controls based on actual performance measurements to attempt to achieve the goals associated with each workload.
- Exploit I/O Priority Queueing
A virtual equivalent of the hardware I/O Priority Queueing facility allows virtual machines running guest operating systems such as z/OS that exploit I/O Priority Queueing to determine the priority of their I/O operations within bounds that can be defined on z/VM. z/VM will automatically set a priority for I/O operations initiated by virtual machines that do not exploit this function.

VMRM Cooperative Memory Management (VMRM-CMM) between a z/VM system and Linux guests assists in managing memory constraint in the system. Based on several variables obtained from the system and storage domain CP monitor data,

VMRM detects when there is such constraint, and notifies specific Linux virtual guests when this occurs. The guests can then take the appropriate action to adjust their memory utilization in order to relieve this constraint on the system, such as issuing a CP DIAGNOSE X'10' instruction to release pages of storage.

For more information about VMRM, see *z/VM: Performance*.

XEDIT

XEDIT is a full-screen editing facility that runs under CMS. XEDIT creates and modifies CMS files and BFS files. System macros and user-written procedures are performed from the XEDIT environment. For more information, see *z/VM: XEDIT User's Guide*.

HELP facility

The HELP facility runs under CMS and provides online assistance for various z/VM functions in the form of menus and panels. HELP information is available for:

- Tasks
- Commands and subcommands
- Messages
- CMS Pipelines stages
- CMS callable routines
- REXX/VM, EXEC 2, and EXEC statements
- Assembler language macros

Additional facilities and features of z/VM, as well as other licensed programs that run on CMS, might provide HELP information for display through the HELP facility. You can also write your own HELP information. For more information, see *z/VM: CMS User's Guide*.

TCP/IP for z/VM

TCP/IP for z/VM brings the power and resources of your mainframe server to the Internet. TCP/IP for z/VM enables z/VM to participate in a multivendor, open networking environment using the TCP/IP protocol suite for communications and interoperability. The applications included in TCP/IP for z/VM provide the ability to transfer files, send mail, log on a remote host, allow access from any other TCP/IP node in the network, and perform other network client and server functions. Applications can be shared transparently across z/VM, z/OS, Linux, and other environments.

TCP/IP consists of a layered structure of protocols and functions. TCP/IP for z/VM provides the following types of functions:

Connectivity and gateway functions

Handle the physical interfaces and routing of data.

Server functions

Provide a service to a client (that is, send or transfer a file).

Client functions

Request a certain service from a server anywhere in the network.

Network status and management functions

Detect and solve network problems.

Base Product

Application programming interfaces (APIs)

Allow you to write your own client/server applications.

Link protocols

Various network protocols compose the network layer available in TCP/IP. Network protocols define how data is transported over a physical network. These network protocols are not defined by TCP/IP. After a TCP/IP packet is created, the network protocol adds a transport-dependent network header before the packet is sent out on the network.

Network protocols

Protocols in the network layer provide connection services for TCP/IP. These protocols connect physical networks and transport protocols. The network protocols include Internet Protocol (IP), Internet Control Message Protocol (ICMP), Address Resolution Protocol (ARP), and Internet Group Management Protocol.

Internet protocol

The Internet Protocol (IP) provides the interface from the transport layer (host-to-host, TCP, or UDP) protocols to the physical-level protocols. IP is the basic transport mechanism for routing IP packets to the next gateway, router, or destination host.

In TCP/IP for z/VM, a single stack provides support for static routing of IP version 6 (IPv6) traffic and IPv4 traffic. IPv6 uses 128-bit addresses. However, the z/VM TCP/IP stack cannot be configured as a tunnel endpoint for tunneling IPv6 traffic over IPv4 networks. Also, the z/VM TCP/IP stack does not support IPv6 security or IPv6 in the upper application layers, such as FTP and SMTP.

Transport protocols

The transport layer of TCP/IP consists of transport protocols, which allow communication between application programs.

Transmission Control Protocol

The Transmission Control Protocol (TCP) provides a reliable vehicle for delivering packets between hosts on an internet. TCP takes a stream of data, breaks it into datagrams, sends each one individually using Internet Protocol (IP), and reassembles the datagrams at the destination node. If any datagrams are lost or damaged during transmission, TCP detects this and re-sends the missing datagrams. The received data stream is a reliable copy of the transmitted data stream.

User Datagram Protocol

The User Datagram Protocol (UDP) provides an unreliable mode of communication between source and destination hosts. UDP is a datagram-level protocol built directly on the IP layer. UDP is used for application-to-application programs between TCP/IP hosts.

Like IP, UDP does not offer a guarantee of datagram delivery or duplication protection. UDP does provide checksums for both the header and data portions of a datagram. However, applications that require reliable delivery of streams of data should use TCP.

Applications and protocols

TCP/IP applications allow users to use network services. These applications are included in the application layer of TCP/IP. The application layer is built on the services of the transport layer. TCP/IP for z/VM supports the following applications and protocols:

- **Domain Name System (DNS)** is a hierarchical naming system for naming hosts.

Note: The native z/VM DNS server is no longer supported, but z/VM still supports the DNS protocol using non-z/VM servers.

- **Dynamic Host Configuration Protocol (DHCP)** provides a framework for passing configuration information to hosts on a TCP/IP network.
- **File Transfer Protocol (FTP)** allows you to transfer data between local and foreign hosts or between two foreign hosts.
- **GDDMXD** is an interface that allows graphics from the IBM Graphical Data Display Manager/VM to be displayed on workstations that support the X Window System.
- **Internet Message Access Protocol (IMAP)** provides the processing that allows a client to access electronic mail that is kept in an IMAP Mailstore server.
- **Lightweight Directory Access Protocol (LDAP)** provides user authentication, authorization, and auditing capabilities.
- **MPRoute** uses either Open Shortest Path First (OSPF) or Routing Information Protocol (RIP), or both, to dynamically create and maintain network routing tables.
- **Network File System (NFS)** allows you to manipulate files on different TCP/IP hosts as if they reside on your host.
- **Remote Execution Protocol (REXEC)** allows you to execute a command on a foreign host and receive the results on the local host.
- **Remote Printing (LPR and LPD)** provides both client and server support for remote printing.
- **Remote Procedure Call (RPC)** is a programming interface that calls subroutines to be executed on a foreign host.
- **Simple Mail Transfer Protocol (SMTP)** is an electronic mail protocol that is implemented with the CMS NOTE and SENDFILE commands.
- **Simple Network Management Protocol (SNMP)** provides a means for managing an internet environment.
- **Secure Sockets Layer (SSL)** protocol provides privacy between two communicating applications — a client and a server.
- **Socket interfaces** allow you to write your own applications to supplement those supplied by TCP/IP.
- **Telnet Protocol** provides a standard method to interface terminal devices and terminal-oriented processes with each other.
- **X Toolkit** is a collection of basic C language routines for developing a variety of application environments.
- **X Window System** supports network transparent windowing and graphics.

For more information, see *z/VM: TCP/IP User's Guide*.

AVS

APPC/VM VTAM Support (AVS) is a Virtual Telecommunications Access Method (VTAM) application that provides advanced program-to-program communication (APPC) services between z/VM and non-z/VM systems in an SNA network. AVS and VTAM run in the same GCS group on a z/VM system. Together, AVS and VTAM enable APPC/VM application programs in a TSAF or CS collection to communicate with:

- Other APPC/VM applications residing in other z/VM systems within the SNA network
- APPC applications residing in non-z/VM systems in the SNA network

For more information, see *z/VM: Connectivity*.

Dump Viewing Facility

The Dump Viewing Facility helps you interactively diagnose system problems. Using this facility, you can display, format, and print data interactively from virtual machine dumps, as well as display and format recorded trace data. The BLOCKDEF utility lets you display, format, and print control block information. The VIEWSYM command lets you display symptom records, making it easier to identify duplicate problems when they occur. For more information, see *z/VM: Dump Viewing Facility*.

Note: Analysis of CP standalone dumps, CP ABEND dumps, and virtual machine dumps of a CP system, functions formerly performed by the Dump Viewing Facility, are now performed by the VM Dump Tool. See “VM Dump Tool” on page 47.

GCS

The Group Control System (GCS) runs in an XA or XC virtual machine in place of CMS. It is a virtual machine supervisor, providing multitasking services that allow numerous tasks to remain active in the virtual machine at one time. One of the functions of GCS is to support a native Systems Network Architecture (SNA) network. The SNA network relies on ACF/VTAM, VTAM SNA Console Support (VSCS), and other network applications to manage its collection of links between terminals, controllers, and processors. GCS provides services for ACF/VTAM, VSCS, and the others, which eliminates your need for VTAM Communications Network Application (VM/VCNA) and a second operating system like VSE. For more information, see *z/VM: Group Control System*.

HCD and HCM for z/VM

Hardware Configuration Definition (HCD) and Hardware Configuration Manager (HCM) provide a comprehensive I/O configuration management environment, similar to that available with the z/OS operating system.

HCM runs on a Windows-based personal computer connected to the z/VM system through a TCP/IP network connection. HCM provides a graphical user interface as well as commands to help you configure your system. You supply the needed I/O configuration information to HCM, which processes the information and passes it to HCD.

HCD runs in a z/VM server virtual machine and performs the work of actually creating and changing the hardware and software aspects of your I/O configuration.

While HCM provides the primary user interface to HCD, HCD also provides a backup user interface on your z/VM host for certain I/O configuration tasks, in case HCM is not available.

The original dynamic I/O configuration capabilities of z/VM are still valid. These consist of a set of system operator commands for changing the hardware server's I/O configuration while the system continues to run, or for managing the hardware I/O configuration of all of the logical partitions in your server. You now have the choice of either using these commands or else using HCM and HCD to manage your I/O configuration. Note, however, that the use of HCM and HCD is incompatible with the original dynamic I/O configuration capabilities. You should select one method to use for the duration of any given IPL of your z/VM system.

For more information, see:

- *z/VM: I/O Configuration*
- *z/OS and z/VM: Hardware Configuration Manager User's Guide*

Language Environment

Language Environment provides the runtime environment for programs written in C/C++, COBOL, or PL/I. Language Environment helps you create mixed-language applications and gives you a consistent method of accessing common, frequently-used services.

Language Environment consists of:

- Basic routines that support starting and stopping programs, allocating storage, communicating with programs written in different languages, and indicating and handling conditions.
- Common library services, such as math services and date and time services, that are commonly needed by programs running on the system. These functions are supported through a library of callable services.
- Language-specific portions of the runtime library. Because many language-specific routines call Language Environment services, behavior is consistent across languages.

For more information, see:

- *z/VM: Language Environment User's Guide*
- *z/OS: Language Environment Concepts Guide*

OSA/SF

Open Systems Adapter-Express (OSA-Express) is a family of integrated hardware features that allow the System z platform to provide industry-standard connectivity directly to clients on local area networks (LANs) and wide area networks (WANs). The Open Systems Adapter Support Facility (OSA/SF) is a host-based tool supplied with z/VM that allows you to customize an OSA's modes of operation. You can access OSA/SF by a CMS user ID, by a REXX call to the OSA/SF API, or through a Java-based graphical user interface (GUI). For more information, see *zEnterprise System, System z10, System z9 and eServer zSeries: Open Systems Adapter-Express Customer's Guide and Reference*.

REXX/VM

REXX/VM contains the REXX/VM Interpreter, which processes the English-like REXX programming language. It also contains the z/VM implementation of the SAA REXX programming language. REXX/VM provides a single source base for the REXX/VM Interpreter in the CMS and GCS components. The REXX/VM Interpreter exploits 31-bit addressing.

The REXX/VM Interpreter helps improve the productivity of your organization. Using REXX, you can write customized application programs and command procedures, tailor CMS commands, and create new XEDIT macros. For more information, see *z/VM: REXX/VM User's Guide*.

TSAF

The Transparent Services Access Facility (TSAF) provides communication services within a collection of z/VM systems without using VTAM. TSAF runs in a CMS virtual machine. A group of up to eight z/VM systems that each have TSAF installed and running can form a TSAF collection. APPC/VM programs on one z/VM system in the TSAF collection can communicate with other APPC/VM programs on the other z/VM systems in the collection. The routing is transparent to the application programs. Communications between the applications proceed as if the applications were running on the same system. For more information, see *z/VM: Connectivity*.

VMSES/E

Virtual Machine Serviceability Enhancements Staged/Extended (VMSES/E) helps you install z/VM and other VMSES/E-enabled products and apply code changes that correct or circumvent reported problems. VMSES/E handles both source code and object code.

VMSES/E can also help you define, build, and manage saved segments. The VMFSGMAP command provides a saved segment mapping interface that lets you modify saved segment definitions and view saved segment layouts prior to actually building them on your system.

For more information, see *z/VM: VMSES/E Introduction and Reference*.

Chapter 6. z/VM optional features

This section provides overviews of the optional features of z/VM:

- “Data Facility Storage Management Subsystem for VM”
- “Directory Maintenance Facility for z/VM”
- “IBM z/VM Single System Image Feature” on page 66
- “Performance Toolkit for VM” on page 67
- “RACF Security Server for z/VM” on page 68
- “RSCS Networking for z/VM” on page 69

Data Facility Storage Management Subsystem for VM

Data Facility Storage Management Subsystem for VM (DFSMS/VM) allows you to control your data and storage resources more efficiently. DFSMS/VM provides:

- Space Management

DFSMS/VM improves DASD utilization by automatically managing space in SFS file pools. As the SFS administrator, DFSMS/VM allows you to:

 - Convert SFS storage to DFSMS/VM-managed storage by assigning *management classes* to files and directories. Each management class tells DFSMS/VM how to treat its members in the course of its management of the file pool.
 - Automatically manage files based on the criteria in each management class. This management may consist of deletion of files, automatic migration of files, or both.
 - *Migrate* (or move) files from DFSMS/VM-managed storage to DFSMS/VM-owned storage by using the assigned *management class*. This function also compresses the data. The files can be automatically recalled when referenced (opened and browsed), or they can be explicitly recalled.
- Minidisk Management

Using DFSMS/VM for minidisk management allows you to check the integrity of CMS minidisks and move them from one location to another. DFSMS/VM helps you migrate CMS minidisks to new DASD quickly, efficiently, and with minimal impact to users.
- Interactive Storage Management Facility (ISMF)

DFSMS/VM uses the ISMF to provide a consistent user interface for storage management tasks.
- IBM Tape Library Dataserver Support

DFSMS/VM Removable Media Services (RMS) provides native z/VM support for the IBM 3494 and 3495 Tape Library Dataservers.

For more information, see *z/VM: DFSMS/VM Planning Guide*.

Directory Maintenance Facility for z/VM

Directory Maintenance Facility for z/VM (DirMaint) provides efficient and secure interactive facilities for maintaining your z/VM system directory. Directory management is simplified by DirMaint's command interface and automated facilities. DirMaint provides a corresponding command for every z/VM directory statement, including cross system extensions (CSE) cluster directory statements. DirMaint's

Optional Features

error checking ensures that only valid changes are made to the directory, and that only authorized personnel are able to make the requested changes.

Program highlights include:

- DirMaint operates as a CMS application and uses CMS interfaces for CMS and CP services. As a CMS application, DirMaint is not dependent on specific hardware, although it does verify that the device types specified in DirMaint commands are only those supported by the z/VM host.
- DirMaint functions are accomplished by two disconnected virtual machines equipped with an automatic restart facility. The use of virtual machines takes advantage of the inherent reliability, availability, and serviceability of the system architecture.
- Any transaction requiring the allocation or deallocation of minidisk extents can be handled automatically.
- All user-initiated transactions can be password-controlled and can be recorded for auditing purposes.
- Command authorization is controlled by assigning DirMaint commands to privileged command sets. Users may be authorized to issue commands from multiple command sets. DirMaint provides nine predefined command sets, but up to 36 sets are supported.
- User exit routines enable centralized directory maintenance of remote systems. Some exit routines also enable DirMaint to interact with other facilities, such as RACF.
- The open command structure allows you to replace any and all commands with your own user-written commands.
- An automated process for copying CMS minidisk files minimizes the possibility of human error. This process optionally formats the old (source) minidisk before returning it to the available minidisk pool.
- The integrity of CMS files is ensured by preventing new minidisk space from being inadvertently allocated over existing extents.
- DirMaint improves overall system efficiency by minimizing the number of DIRECTXA utility runs required. The update-in-place facility (DIAGNOSE code X'84') can be used to place many of the changes online immediately.
- System security is enhanced by providing the ability to enforce regular password changes. When changing the password, the user is required to enter the new password twice to guard against typographical errors.
- An additional level of security can be implemented by requiring that a password be entered for every user transaction. This is the default.

For more information, see:

- “DirMaint restrictions” on page 36
- “DirMaint system integrity” on page 40
- “DirMaint security” on page 42
- *z/VM: Directory Maintenance Facility Tailoring and Administration Guide*

IBM z/VM Single System Image Feature

The IBM z/VM Single System Image Feature (VMSSI) enhances the z/VM systems management, communications, disk management, device mapping, virtual machine definition management, installation, and service functions to enable up to four z/VM systems to share and coordinate resources within a single system image (SSI) cluster.

The member systems in a z/VM SSI cluster are part of the same ISFC collection and use ISFC channel connections to communicate. All members of a cluster also share DASD for virtual machines and selected z/VM data, as well as LAN segments and IP subnets. The concept of a global virtual switch provides identical network connectivity across all active members within a cluster.

Members of a z/VM SSI cluster are managed, serviced, and administered as one system. Resources used by both CP and virtual machines are shared among all members. These resources include:

- User directory
- Minidisks
- Spool files
- Network device MAC addresses

The combination of enhanced functions in the SSI cluster provides the foundation that enables live guest relocation, which is the capability for moving a running Linux guest from one system to another within the cluster. Live guest relocation provides continuity for virtual server workloads over planned z/VM and machine outages. Verification that needed resources and machine features are available on the destination system prior to the relocation is provided. This verification may also be performed on request to assess a guest's eligibility for relocation. In an SSI cluster comprising different machine models, the architecture level presented to each guest is tailored to the set of machine features common to the member systems within the guest's specified relocation domain.

For more information about the z/VM SSI environment, setting up SSI clusters, and live guest relocation, see *z/VM: CP Planning and Administration*.

Performance Toolkit for VM

Performance Toolkit for VM, which is derived from the FCON/ESA program (5788-LGA), assists operators and systems programmers or analysts in the following areas:

- Operation of the system operator console in full screen mode
- Support for managing multiple VM systems
- Post processing of VM history files
- Performance monitoring
- Serving data through a web server for viewing with web browsers
- PC-based graphics
- TCP/IP performance reporting

In addition to analyzing VM performance data, the Performance Toolkit processes Linux performance data obtained from the Resource Management Facility (RMF™) Linux performance gatherer, **rmfpms**. The **rmfpms** application is available from IBM RMF PM with Support for Linux Enterprise Server (www.ibm.com/systems/z/os/zos/features/rmf/product/rmfhtmls/pmweb/).

The Linux performance data obtained from RMF can be viewed and printed in a manner similar to the presentation of VM data.

For more information, see *z/VM: Performance Toolkit Guide*.

RACF Security Server for z/VM

RACF Security Server for z/VM is a security tool that works together with existing functions in the z/VM base system to provide improved data security for an installation. RACF protects information by controlling access to it. RACF also controls what you can do on the operating system and protects your resources. It provides this security by identifying and verifying users, authorizing users to access protected resources, and recording and reporting access attempts.

To help each installation meet its unique security needs and objectives, RACF provides:

- Protection of installation-defined resources
- Flexible control of access to protected resources
- The ability to store information for other products
- A choice of centralized or decentralized control profiles
- Transparency to end users
- Exits for installation-written routines

Your organization can define individuals and groups who use the system that RACF protects. A security administrator uses RACF to define a profile for each individual that identifies that person's user ID, password, and other information. A group is a collection of individuals who have common needs and requirements. For example, a whole department may be defined as one group. Your organization can also define what authorities you have, or what authorities a group you belong to has. RACF controls what you can do on the system. Some individuals have a great degree of authority, while others have little authority. The degree of authority you are given is based on what you need to do your job.

In addition to defining user and group authorities, RACF protects resources. You can protect system resources and user resources. System resources include system minidisks, system SFS files and directories, certain VM events, and terminals. User resources include user minidisks and user SFS files and directories.

RACF stores all this information about users, groups, and resources in profiles. A profile is a record of RACF information that has been defined by the security administrator. There are user, group, and resource profiles. Using the information in its profiles, RACF authorizes access to certain resources. RACF applies user attributes, group authorities, and resource authorities to control use of the system. The security administrator or someone in authority in your organization controls the information in your user profile, in group profiles, and in resource profiles. You, as an end user, control the information in profiles describing your own resources, such as your own minidisks. You can protect your data by setting up resource profiles. You can set up an access list in your resource profile to control who has read-access and who has write-access to your data.

In addition to uniquely identifying and authorizing users, RACF can record what users do on the system. It keeps track of what happens on the system so that an organization can monitor who is logged on to the system at any given time. RACF reports if persons have attempted to perform unauthorized actions. For example, RACF can record when someone who does not have the proper authority tries to use or change your data. The security administrator can monitor these activities and generate reports.

For more information, see *z/VM: RACF Security Server General User's Guide*.

RSCS Networking for z/VM

Remote Spooling Communications Subsystem (RSCS) Networking for z/VM is a networking program that enables users on a z/VM system to send messages, files, commands, and jobs to other users within a network. RSCS connects nodes (systems, devices, and workstations) using links. These links allow data, consisting mainly of CP spool files, to be transferred between the nodes.

Running under the GCS component of z/VM, RSCS uses the spooling facilities of z/VM to store and retrieve data. z/VM handles data transfer within its system by means of spooling. RSCS extends the basic spooling capabilities of z/VM, handling data transfer between the z/VM system and outside sources. Data is stored on a spool after RSCS receives it and until RSCS can forward it to its destination. RSCS uses communications equipment to transfer data between the local z/VM system and other systems or remote locations.

A node in an RSCS network is either a system node or a station node. A station node can originate and receive information. It can be a computer, a workstation, or a printer. A system node, however, must be a computer. Besides originating and receiving information, system nodes can also relay information between two other nodes.

RSCS can communicate with system nodes that are running under the control of network job entry (NJE) compatible subsystems, such as:

- JES2 or JES3
- RSCS
- VSE/POWER
- AS/400® Communications Utilities
- Products that provide NJE functions for Linux or AIX®

For more information about NJE, see *Network Job Entry: Formats and Protocols*.

RSCS can communicate with station nodes that are:

- ASCII printers or plotters
- Computers running under the control of a system that can provide a multileaving protocol
- IBM 3270 Information Display System Printers
- Line printer router (LPR) daemons and clients in a TCP/IP network
- Unsolicited File Transfer (UFT) daemons and clients in a TCP/IP network
- Workstations running under the control of remote job entry (RJE)

Each link in an RSCS network is associated with a programming routine, called a driver, that manages the transmission and reception of files, messages, and commands over the link. The way that a driver manages the data is called a protocol. All file transmission between networking nodes uses NJE protocol, 3270 printers use 3270 data streams, workstations use RJE protocol, and ASCII printers use data streams appropriate to that printer. Systems Network Architecture (SNA) provides one set of protocols that governs communications on links. The method that RSCS uses for sending data to a node varies, depending on the type of connection used to establish the link. RSCS can support non-SNA (such as binary synchronous communication or channel-to-channel), SNA, and TCP/IP connections.

For more information, see *z/VM: RSCS Networking Planning and Configuration*.

Chapter 7. z/VM library guide

This section provides a guide to the z/VM product information. It contains the following topics:

- “Where to get z/VM information”
- “Abstracts of the publications for the z/VM base product” on page 72
- “Abstracts of the publications for the z/VM optional features” on page 85

Where to get z/VM information

The z/VM product information includes primarily z/VM publications, but also includes some z/OS publications and other IBM publications containing information that applies to z/VM. The z/VM product information is available from several sources, as described in this section.

z/VM Information Center

IBM publishes the documentation for each z/VM release in an information center using the IBM Eclipse Help System framework. The z/VM information center includes documentation for the z/VM base product and optional features, the prerequisite IBM products, and some optional IBM products and facilities. The content of the information center is identical to that in the traditional IBM BookManager and Adobe PDF formats; however, the presentation might differ to some degree.

Advantages of information centers are:

- Content is indexed by Google and other internet search engines to help locate information more easily.
- Custom searches can be created that include only the information you need for a particular task or job role.

The z/VM V6R2 Information Center is available from the following sources:

- z/VM V6R2 Information Center (publib.boulder.ibm.com/infocenter/zvm/v6r2/)
- *z/VM V6R2 Information Center, SK5T-7098-02*. This DVD includes the IBM Eclipse Help System framework, which allows you to install the information center on a workstation or intranet. The DVD can be ordered (for a fee) from IBM Publications Center (www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss).

z/VM Internet Library

IBM: z/VM Internet Library (www.ibm.com/vm/library/) is the Library page of the z/VM product website and provides links to various online z/VM information resources, such as:

- z/VM information centers
- z/VM BookManager-format bookshelves (displayed as HTML), offering standard BookManager searching capability
- z/VM PDF libraries
- z/VM program directories
- z/VM data areas, control blocks, and monitor records
- IBM Redbooks® publications
- White papers, consultant reports, and performance reports

Library Guide

- Data sheets and brochures
- Reference guides
- Journals, bulletins, and newsletters

z/VM Collection DVD

The *IBM Online Library: z/VM Collection*, SK5T-7054, includes BookManager bookshelves and PDF libraries for z/VM and current IBM licensed programs that run on z/VM.

Notes:

1. Only unlicensed publications are included.
2. Some programs might not have PDF libraries; some publications might not have PDF versions.

The DVD also includes the following programs:

IBM Softcopy Reader

Allows you to view and search BookManager files and organize BookManager files and bookshelves.

IBM Softcopy Librarian

Allows you to manage BookManager files in a repository and download BookManager files from the Internet.

Softcopy Receiver Tool

Allows you to transfer BookManager files from the DVD to the workstation or host.

IBM Publications Center

You can obtain z/VM publications from IBM Publications Center (www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss). The IBM Publications Center provides worldwide downloading and ordering facilities for a broad range of IBM publications in many languages.

Documentation supplied with z/VM

The following documentation is shipped automatically with z/VM, at no additional cost:

- One copy of the *IBM Online Library: z/VM Collection DVD*
- One printed copy of each of the following z/VM publications:
 - *z/VM: License Information*
 - *z/VM: Installation Guide*
- A complete set of z/VM HELP files (included in the z/VM system image)

Abstracts of the publications for the z/VM base product

The following sections contain abstracts of the publications that provide documentation for the z/VM base product, grouped in topic (task) categories.

- “System overview” on page 73
- “Installation, migration, and service” on page 73
- “Planning and administration” on page 74
- “Customization and tailoring” on page 77
- “Operation and use” on page 77
- “Application programming” on page 79

- “Diagnosis” on page 84

Some of the publications have new editions:

- z/VM publications that have been updated to support the general availability (GA) of z/VM V6.2 are identified with the designation **V6R2**.

Note: z/VM publications with prior release designations support z/VM V6.2 but do not contain new or changed information.

- Updated editions of non-z/VM publications (compared with the previous z/VM release) are identified with the designation **[UPD]**.

Notes:

1. The updates in a non-z/VM publication might not be z/VM related.
2. The indicated edition of a non-z/VM publication might not be the latest edition. Functions described in newer editions might not be supported in this z/VM release.

System overview

These publications provide an overview of z/VM.

z/VM V6R2 General Information, GC24-6193-02

z/VM: General Information provides the following information about z/VM:

- Introduction to the base product and optional features
- How z/VM can help you
- What is new or changed in the current release
- Hardware and software requirements
- Guide to the z/VM documentation
- IBM servers, guest operating systems, and devices supported by current z/VM releases

z/VM V6R2 Glossary, GC24-6195-02

z/VM: Glossary defines the terms used in z/VM publications.

z/VM V6R2 License Information, GC24-6200-04

z/VM: License Information provides information on the warranted functions of z/VM, the specified operating environment, and the supplemental terms.

Installation, migration, and service

These publications provide information to help you install and maintain a z/VM system.

z/VM V6R2 Installation Guide, GC24-6246-00

z/VM: Installation Guide contains step-by-step procedures for installing a z/VM system. The procedures feature an automated installation process using a panel interface to load a prebuilt z/VM System Image packaged in DASD Dump/Restore (DDR) format on magnetic tape or packaged in binary format on DVD. Installation worksheets are included that you can use to plan your installation.

z/VM V6R2 Migration Guide, GC24-6201-02

z/VM: Migration Guide provides a history of significant changes and enhancements since VM/ESA V2.1. This information is provided to help system support personnel migrate a system to the new z/VM release. It is also intended for application programmers who use external interfaces in their programs. This document contains the following information:

- An introduction to migration

- Descriptions of system changes and enhancements that you should be aware of before migrating
- Identification of specific external interfaces that have changed, with changes classified as upwardly compatible or incompatible
- Guidance for some migration tasks you might need to perform, such as converting from system definition macros to system configuration files, migrating your system directory, or migrating a pre-z/VM V6.2 Cross System Extensions (CSE) complex to z/VM V6.2

z/VM V6R2 Service Guide, GC24-6247-00

z/VM: Service Guide contains step-by-step procedures for installing preventive and corrective service to the z/VM components and the preinstalled z/VM facilities and features. This document should be used in conjunction with *z/VM: VMSES/E Introduction and Reference*.

z/VM V6R2 VMSES/E Introduction and Reference, GC24-6243-01

z/VM: VMSES/E Introduction and Reference provides an overview of the Virtual Machine Serviceability Enhancements Staged/Extended (VMSES/E) component of z/VM. It describes how to use VMSES/E tools to install, migrate, build, service, and delete program products from a more general, less z/VM-specific, point of view than *z/VM: Installation Guide* and *z/VM: Service Guide*. This document discusses the software inventory and how you can use it to manage the products on your system. Reference information is provided on the product parameter file, software inventory tables, VMSES/E commands, and other related commands.

Planning and administration

These publications provide information to help you define and manage a z/VM system. Planning is an iterative task in that many of the decisions are made before installation, are continually evaluated after installation, and are revised as appropriate. The administration task involves defining the characteristics of data processing resources to z/VM. The resources can be data files, databases, programs, users, and so forth.

z/VM V6R2 CMS File Pool Planning, Administration, and Operation, SC24-6167-01

z/VM: CMS File Pool Planning, Administration, and Operation provides information on planning for, administering, and operating CMS file pools. It provides information about using file pools as repositories for CMS Shared File System (SFS) and OpenExtensions Byte File System (BFS) data. It also provides information about using file pool server machines for Coordinated Resource Recovery (CRR) and FIFO functions. The document includes a reference section that contains descriptions of file pool startup parameters and file pool administration and server commands.

z/VM V6R2 CMS Planning and Administration, SC24-6171-01

z/VM: CMS Planning and Administration provides information on the Conversational Monitor System (CMS) component of z/VM. It includes information on the following topics:

- Tailoring CMS
- Administering CMS Pipelines
- Planning for VSE simulation and VSAM support
- Setting up and using the Programmable Operator Facility
- Managing the CMS batch facility

z/VM V6R2 Connectivity, SC24-6174-02

z/VM: Connectivity provides an overview of the z/VM facilities that enable logical connections between systems. This document presents basic connectivity concepts and planning considerations. It describes how to plan and set up real networks (using OSA-Express and HiperSockets) and virtual networks (using guest LANs and virtual switches). It also provides information about using server and requester virtual machines, the Transparent Services Access Facility (TSAF), APPC/VM VTAM Support (AVS), and the Inter-System Facility for Communications (ISFC).

z/VM V6R2 CP Planning and Administration, SC24-6178-02

z/VM: CP Planning and Administration provides information about the Control Program (CP) component of z/VM. It includes the following topic categories:

- System planning and administration

This section includes information on using the Stand-Alone Program Loader (SAPL), creating and updating the system configuration file, configuring z/VM for an ensemble, and other topics.

- User planning and administration

This section includes information on creating and updating a user directory and redefining command privilege classes.

- Storage planning and administration

This section includes information on managing real storage, allocating DASD space, sharing DASD, defining and managing SCSI FCP disks, and other topics.

- Single system image cluster planning and administration

This section includes information on setting up a z/VM single system image (SSI) cluster, preparing for guest relocations in an SSI cluster, and other topics.

z/VM V6R2 Getting Started with Linux on System z, SC24-6194-01

z/VM: Getting Started with Linux on System z describes how to configure and use z/VM functions and facilities for Linux servers running on the System z platform. This document is designed to help system administrators who have limited knowledge of z/VM, but want to deploy Linux servers on z/VM. The document provides requirements and guidelines to implement during z/VM installation, but primarily assumes that z/VM is installed and you are ready to deploy Linux servers in z/VM virtual machines. Topics covered include:

- z/VM basics
- Planning for Linux virtual servers
- Changing the z/VM system configuration
- Configuring the Directory Maintenance Facility
- Configuring TCP/IP for z/VM
- Creating and cloning Linux virtual machines
- Setting up basic system automation
- Performing runtime tasks
- Monitoring performance and capacity
- Servicing z/VM
- Preparing for live guest relocation
- Using FTP to install Linux from the hardware management console

z/VM V6R2 Group Control System, SC24-6196-01

z/VM: Group Control System provides information about the Group Control System (GCS) component of z/VM. GCS is a virtual machine supervisor. It bands many virtual machines together in a group and supervises their operations. The specific function of GCS is to support a native VM Systems Network Architecture (SNA) network. This document is intended for anyone writing programs that run under GCS. It provides planning information and describes the purpose and use of GCS commands and macroinstructions.

z/VM V6R2 I/O Configuration, SC24-6198-02

z/VM: I/O Configuration describes how to plan the I/O configuration (channel paths, control units, and I/O devices) on your z/VM system. Two methods are described. Part 1 describes how to use Hardware Configuration Definition (HCD) and Hardware Configuration Manager (HCM) to manage the hardware and software I/O configuration. HCM is a Windows-based program that provides a graphical user interface. Part 2 describes how to use z/VM dynamic I/O configuration, which is a command interface for modifying the hardware I/O configuration.

z/VM V6R2 Running Guest Operating Systems, SC24-6228-02

z/VM: Running Guest Operating Systems contains information to help you plan for and run guest operating systems under the supervision of z/VM. It discusses the fundamentals of guest support in z/VM and provides specific information on how to run z/OS, VSE, or z/VM in a virtual machine.

For information about running Linux in a virtual machine, see *z/VM: Getting Started with Linux on System z*.

z/VM V6R2 Saved Segments Planning and Administration, SC24-6229-01

z/VM: Saved Segments Planning and Administration provides information about using saved segments on your z/VM system. It includes information on the following topics:

- Planning and defining CP saved segments
- Planning and defining CMS logical saved segments
- Using VMSES/E to define, build, and manage saved segments

z/VM V6R2 Secure Configuration Guide, SC24-6230-01

z/VM: Secure Configuration Guide describes the steps necessary to configure your z/VM installation to conform with the requirements of the Common Criteria.

z/VM V6R2 TCP/IP LDAP Administration Guide, SC24-6236-01

z/VM: TCP/IP LDAP Administration Guide contains information to assist administrators in using the z/VM TCP/IP Lightweight Directory Access Protocol (LDAP) server to maintain directory information. The LDAP server provides user authentication, authorization, and auditing capabilities.

z/VM V6R2 TCP/IP Planning and Customization, SC24-6238-02

z/VM: TCP/IP Planning and Customization provides information to help you plan and set up TCP/IP networks on your z/VM system. It describes how to define and configure the virtual machines, servers, and applications available in TCP/IP for z/VM. It also describes how to customize and tune TCP/IP for your specific needs.

z/OS and z/VM: Hardware Configuration Manager User's Guide, SC33-7989-13 [UPD]

z/OS and z/VM: Hardware Configuration Manager User's Guide describes how to use the Hardware Configuration Manager (HCM) graphical user interface to create and maintain hardware configuration information.

zEnterprise System, System z10, System z9 and eServer zSeries: Open Systems Adapter-Express Customer's Guide and Reference, SA22-7935-13 [UPD]

zEnterprise System, System z10, System z9 and eServer zSeries: Open Systems Adapter-Express Customer's Guide and Reference provides step-by-step instructions for setting up and using the Open Systems Adapter Facility (OSA/SF) to configure and manage OSA-Express server features.

System z9 and eServer zSeries 890 and 990: Open Systems Adapter-Express Integrated Console Controller User's Guide, SA22-7990-01

System z9 and eServer zSeries 890 and 990: Open Systems Adapter-Express Integrated Console Controller User's Guide provides instructions for setting up and using the Open Systems Adapter-Express Integrated Console Controller (OSA-ICC), a function of the OSA-Express 1000BASE-T Ethernet feature that integrates 3270 emulation for console session connections into System z9 and zSeries® 890 and 990 servers.

System z: Open Systems Adapter-Express Integrated Console Controller 3215 Support, SA23-2247-00

System z: Open Systems Adapter-Express Integrated Console Controller 3215 Support provides instructions for setting up and using the Open Systems Adapter-Express Integrated Console Controller (OSA-ICC) for the 3215 data stream. This publication is intended to be used with *System z9 and eServer zSeries 890 and 990: Open Systems Adapter-Express Integrated Console Controller User's Guide*.

System z10: Open Systems Adapter-Express3 Integrated Console Controller Dual-Port User's Guide, SA23-2266-01 [UPD]

System z10: Open Systems Adapter-Express3 Integrated Console Controller Dual-Port User's Guide describes the configuration process for the Open Systems Adapter-Express3 Integrated Console Controller Dual-Port feature. This publication, which is intended to be used with *System z9 and eServer zSeries 890 and 990: Open Systems Adapter-Express Integrated Console Controller User's Guide*, is primarily for system programmers who are responsible for setting up system consoles for System z operating systems.

Customization and tailoring

These publications provide information about extending or enhancing a z/VM system.

z/VM V6R2 CP Exit Customization, SC24-6176-01

z/VM: CP Exit Customization describes how to customize a z/VM system using CP exit points, both IBM-defined and customer-written. Topics include:

- Creating, controlling, and calling CP exit points
- Creating, controlling, and overriding CP commands
- Creating, controlling, and overriding DIAGNOSE codes
- Creating, controlling, and using local message repositories
- Dynamically loading into and unloading from CP storage
- Using dynamic CP exits

z/VM V6R2 Performance, SC24-6208-02

z/VM: Performance contains information about the planning, managing, measuring, and tuning considerations needed to obtain optimum z/VM system performance. It provides an overview of z/VM system characteristics and discusses performance methodology, planning measures, monitoring facility tools, and tuning actions.

Operation and use

These publications provide information about operating a z/VM system and using z/VM commands and other functions.

z/VM V6R2 CMS Commands and Utilities Reference, SC24-6166-02

z/VM: CMS Commands and Utilities Reference provides detailed reference information on all general-use CMS commands, z/VM HELP Facility format words, and CMS utilities (functions intended primarily for system programmers).

z/VM V6R1 CMS Pipelines Reference, SC24-6169-00

z/VM: CMS Pipelines Reference provides reference information on the CMS PIPE command and CMS Pipelines stages, subcommands, and assembler macros.

z/VM V6R1 CMS Pipelines User's Guide, SC24-6170-00

z/VM: CMS Pipelines User's Guide describes the general concepts of CMS Pipelines and how to use CMS Pipelines, including:

- Writing a pipeline
- Using filters, host command interfaces, and device drivers
- Writing a stage command
- Writing a multistream pipeline
- Using pipelines in exec procedures and XEDIT macros
- Storing user-written stage commands in a filter package
- Migrating from CMS Pipelines 1.1.6

z/VM V6R1 CMS Primer, SC24-6172-00

z/VM: CMS Primer introduces the Conversational Monitor System (CMS) component of z/VM and provides information on basic CMS tasks. It presents, through examples, a portion of the functions and commands available with CMS, with the primary emphasis on:

- Logging on
- Editing and working with files
- Using the Shared File System (SFS)
- Communicating with other users

z/VM V6R2 CMS User's Guide, SC24-6173-01

z/VM: CMS User's Guide describes how to use various facilities provided by CMS. Topics include:

- Managing your file system
- Creating and maintaining an online Help facility
- Using windowing commands and full-screen CMS
- Modifying and running execs and programs

This document is intended for users with a general understanding of CMS, who want to use CMS at a more advanced level. Users with no CMS knowledge should first read *z/VM: CMS Primer*.

z/VM V6R2 CP Commands and Utilities Reference, SC24-6175-02

z/VM: CP Commands and Utilities Reference provides detailed reference information on Control Program (CP) commands and system utilities for users of every privilege class. System utilities perform CP functions but operate only in the CMS environment.

z/VM V6R2 System Operation, SC24-6233-02

z/VM: System Operation explains how to operate a z/VM system. It provides information about system operator tasks, such as:

- Bringing up, running, and shutting down the system
- Controlling local devices
- Responding to errors
- Collecting information about system operation

- Performing data management services
- Running utility programs

z/VM V6R2 TCP/IP User's Guide, SC24-6240-02

z/VM: TCP/IP User's Guide is intended for end users and describes how to use TCP/IP after it has been installed and customized on a network. The document explains how to use the applications available in TCP/IP, including:

- Transferring files
- Sending electronic mail
- Logging on to a foreign host
- Using LDAP operation utilities
- Monitoring the TCP/IP network
- Using Network File System commands
- Using remote printing
- Managing TCP/IP network resources with SNMP
- Using the Domain Name System

z/VM V6R1 Virtual Machine Operation, SC24-6241-00

z/VM: Virtual Machine Operation contains information about operating a z/VM virtual machine. It might be particularly useful for the user who wants to set up a virtual machine in which to run a guest operating system. The document includes topics such as:

- Setting up, starting, and operating a virtual machine
- Communicating with the guest operating system and with the z/VM Control Program
- Managing the storage, processor, and I/O resources of a virtual machine
- Using z/VM facilities to test programs running in a virtual machine

z/VM V6R1 XEDIT Commands and Macros Reference, SC24-6244-00

z/VM: XEDIT Commands and Macros Reference provides detailed reference information on the XEDIT command, subcommands, and macros. Users should first read *z/VM: XEDIT User's Guide*.

z/VM V6R1 XEDIT User's Guide, SC24-6245-00

z/VM: XEDIT User's Guide provides a working knowledge of the z/VM system editor, XEDIT. XEDIT provides a wide range of functions for text processing and programming development. Because it is both a full-screen and a line-mode editor, XEDIT can be used on display and on typewriter terminals.

CMS/TSO Pipelines: Author's Edition, SL26-0018-06

CMS/TSO Pipelines: Author's Edition provides information on CMS/TSO Pipelines from its author. (CMS/TSO Pipelines evolved from the CMS Pipelines PRPQ.) In addition to a description of CMS/TSO Pipelines terms and concepts, this document includes reference information on the PIPE command and Pipelines stages, commands, and messages. This document contains additional information not found in *z/VM: CMS Pipelines Reference* and also provides information intended for the advanced Pipelines user.

Application programming

These publications provide information about creating application programs to do specific functions on z/VM. The publications describe the primary application programming interfaces (APIs) provided by z/VM.

z/VM V6R2 CMS Application Development Guide, SC24-6162-01

z/VM: CMS Application Development Guide provides information about developing application programs in CMS. The development process includes planning, designing, writing, compiling, debugging, executing, and updating. This document also describes the following CMS services and how you can use them to develop an application:

- Shared File System (SFS)
- OpenExtensions
- Common Programming Interface (CPI) Communications
- Coordinated Resource Recovery (CRR)
- VM Data Space support
- CMS libraries
- CMS Batch Facility
- Parsing facility
- Message repositories

z/VM V6R2 CMS Application Development Guide for Assembler, SC24-6163-01

z/VM: CMS Application Development Guide for Assembler provides information to help assembler language programmers:

- Understand CMS programming interfaces and virtual machine architecture
- Use CMS services to build, load, run, and package assembler programs
- Develop OS/MVS and VSE applications under CMS
- Use Access Methods Services and VSAM under CMS and CMS/DOS

z/VM V6R1 CMS Application Multitasking, SC24-6164-00

z/VM: CMS Application Multitasking describes how you can use CMS multitasking to develop and run multitasking application programs written in C/C++, assembler, or REXX. The document provides introductory and tutorial information as well as detailed reference material.

z/VM V6R2 CMS Callable Services Reference, SC24-6165-01

z/VM: CMS Callable Services Reference describes the basic set of CMS callable services library (CSL) routines. These CSL routines primarily perform functions related to CMS file system management and CMS file pool administration.

z/VM V6R2 CMS Macros and Functions Reference, SC24-6168-01

z/VM: CMS Macros and Functions Reference describes the preferred CMS macroinstructions and functions. It also describes the CMS macros and functions supported only for compatibility (because they support only 24-bit addressing).

z/VM V6R2 CP Programming Services, SC24-6179-02

z/VM: CP Programming Services describes the application programming services and facilities available in CP, including:

- DIAGNOSE instructions
- Inter-User-Communications Vehicle (IUCV)
- Advanced Program-to-Program Communications for VM (APPC/VM)
- CP System Services
- ESA/XC address-space management macros
- Symptom record reporting
- Access Control Interface (ACI) for an external security manager

z/VM V6R1 CPI Communications User's Guide, SC24-6180-00

z/VM: CPI Communications User's Guide provides step-by-step instructions for using Systems Application Architecture (SAA) Common Programming Interface (CPI) Communications to write communications programs to run in the CMS

environment. Sample programs written in REXX/VM show how to use SAA CPI Communications calls and the CMS extensions to CPI Communications.

z/VM V6R1 Enterprise Systems Architecture/Extended Configuration Principles of Operation, SC24-6192-00

z/VM: Enterprise Systems Architecture/Extended Configuration Principles of Operation provides a detailed description of the Enterprise System Architecture/Extended Configuration (ESA/XC) virtual machine architecture. It describes how ESA/XC operates as compared to Enterprise Systems Architecture/390 (ESA/390), upon which ESA/XC is based. It is intended as a reference for programmers who write or debug programs that run in ESA/XC virtual machines.

z/VM V6R2 Language Environment User's Guide, SC24-6199-01

z/VM: Language Environment User's Guide provides information for using Language Environment for z/VM, which is based on z/OS Language Environment. It provides information unique to the z/VM platform and is intended to be used in conjunction with the z/OS Language Environment documentation (included in the z/VM library), such as:

- *z/OS: Language Environment Concepts Guide*
- *z/OS: Language Environment Programming Guide*
- *z/OS: Language Environment Programming Reference*

z/VM V6R1 OpenExtensions Advanced Application Programming Tools, SC24-6202-00

z/VM: OpenExtensions Advanced Application Programming Tools provides advanced information for using the OpenExtensions utilities **lex**, **yacc**, and **make**.

z/VM V6R2 OpenExtensions Callable Services Reference, SC24-6203-01

z/VM: OpenExtensions Callable Services Reference describes the callable services that provide interfaces between z/VM and the functions specified in the IEEE POSIX.1 standard. These services are used by the POSIX calls in the C/C++ runtime library. This document also describes callable services that are not related to the standards.

z/VM V6R1 OpenExtensions Commands Reference, SC24-6204-00

z/VM: OpenExtensions Commands Reference describes the OpenExtensions commands and utilities. It also describes the CMS OPENVM commands, which are used to obtain OpenExtensions services.

z/VM V6R1 OpenExtensions POSIX Conformance Document, GC24-6205-00

z/VM: OpenExtensions POSIX Conformance Document describes the z/VM implementation of those areas of the Institute of Electrical and Electronics Engineers (IEEE) Portable Operating System Interface for Computer Environments (POSIX) standards that were declared to be optional or implementation-defined. This document describes the z/VM implementation for both POSIX.1 and POSIX.2. The implementation of these standards in z/VM is known as OpenExtensions and is included in CMS.

z/VM V6R1 OpenExtensions User's Guide, SC24-6206-00

z/VM: OpenExtensions User's Guide describes the OpenExtensions Byte File System (BFS) and provides information for using the OpenExtensions shell commands. The shell provides a UNIX-like interactive user environment.

z/VM V6R2 Program Management Binder for CMS, SC24-6211-02

z/VM: Program Management Binder for CMS describes how to use the Program Management binder for CMS (CMS binder). The CMS binder is based on the z/OS MVS Program Management Binder, and this document describes the differences in usage and behavior between the two programs. This document is intended to be used in conjunction with the following z/OS MVS Program Management documents (included in the z/VM library):

- *z/OS MVS Program Management: User's Guide and Reference*
- *z/OS MVS Program Management: Advanced Facilities*

z/VM V6R1 Reusable Server Kernel Programmer's Guide and Reference, SC24-6220-00

z/VM: Reusable Server Kernel Programmer's Guide and Reference describes how you can use the reusable server kernel supplied with z/VM to develop and run server programs in the CMS environment. The document covers advanced material in server construction and is not intended for beginning programmers.

z/VM V6R2 REXX/VM Reference, SC24-6221-01

z/VM: REXX/VM Reference provides reference information about REXX instructions and functions and their use. It also provides information about the REXX Sockets API.

z/VM V6R1 REXX/VM User's Guide, SC24-6222-00

z/VM: REXX/VM User's Guide provides step-by-step instructions for using the REXX interpretive command and macrolanguage on z/VM. It is intended for users with some knowledge of z/VM, editors, and terminals, but previous programming experience is not needed.

z/VM V6R2 Systems Management Application Programming, SC24-6234-02

z/VM: Systems Management Application Programming describes the socket-based application programming interface (API) that z/VM provides for performing system management functions for virtual systems (guests) in a z/VM environment.

z/VM V6R2 TCP/IP Programmer's Reference, SC24-6239-02

z/VM: TCP/IP Programmer's Reference, which is intended for users and programmers who are familiar with z/VM and its CP and CMS components, contains information about the following application programming interfaces (APIs):

- z/VM C sockets
- Pascal
- Virtual Machine Communication Facility (VMCF)
- Inter-User Communication Vehicle (IUCV) sockets
- Remote Procedure Calls (RPCs)
- X Window System
- Simple Network Management Protocol (SNMP) agent distributed program interface
- Simple Mail Transfer Protocol (SMTP) virtual machine interfaces
- Telnet exits
- File Transfer Protocol (FTP) exit
- IMAP authentication exit
- Remote authorization and auditing through LDAP
- Building an LDAP server plug-in

Common Programming Interface Communications Reference, SC26-4399-09

Common Programming Interface Communications Reference describes Systems Application Architecture (SAA) Common Programming Interface (CPI) Communications in detail, including scenarios and individual routines. It is intended for anyone writing an application program that communicates with another program using the APPC protocol. The communications occur within a single TSAF collection, across many TSAF collections, or between a TSAF collection and a SNA network.

Common Programming Interface Resource Recovery Reference, SC31-6821-01

Common Programming Interface Resource Recovery Reference describes the System Application Architecture (SAA) Common Programming Interface resource recovery interface in detail, including scenarios and individual routines. It is intended for programmers who want to write applications that use the SAA resource recovery interface. The SAA resource recovery interface lets programs coordinate exchanges of data and updates to databases and other resources. This coordination ensures that either all changes become permanent or all are undone.

z/OS: IBM Tivoli Directory Server Plug-in Reference for z/OS, SA76-0148-00

z/OS: IBM Tivoli Directory Server Plug-in Reference for z/OS describes the application services routines that can be used to create plug-ins to extend the capabilities of the z/VM TCP/IP LDAP server. For information about creating LDAP server plug-ins on z/VM, see *z/VM: TCP/IP Programmer's Reference*.

z/OS: Language Environment Concepts Guide, SA22-7567-11 [UPD]

z/OS: Language Environment Concepts Guide provides information on evaluating and planning for Language Environment.

z/OS: Language Environment Debugging Guide, GA22-7560-11 [UPD]

z/OS: Language Environment Debugging Guide provides assistance in detecting and locating programming errors that might occur during run time under Language Environment. The document can help you establish a debugging process to analyze data and narrow the scope and location of where an error might have occurred. You can read about how to prepare a routine for debugging, how to classify errors, and how to use the debugging facilities that Language Environment provides.

z/OS: Language Environment Programming Guide, SA22-7561-11 [UPD]

z/OS: Language Environment Programming Guide contains information about linking, running, and using services within Language Environment. It also contains the Language Environment program management model and provides language-specific and operating system-specific information, where applicable.

z/OS: Language Environment Programming Reference, SA22-7562-12 [UPD]

z/OS: Language Environment Programming Reference provides a detailed description of each Language Environment runtime option and callable service, as well as information on how to use them. It also provides programming examples that illustrate how each callable service can be used in routines written in Language Environment-conforming high-level languages (HLLs) and assembler language.

z/OS: Language Environment Run-Time Messages, SA22-7566-11 [UPD]

z/OS: Language Environment Run-Time Messages contains runtime messages and return codes for Language Environment and supported high-level languages.

z/OS: Language Environment Writing Interlanguage Communication Applications, SA22-7563-08 [UPD]

z/OS: Language Environment Writing Interlanguage Communication Applications contains information on creating and running interlanguage communication (ILC) applications under Language Environment. ILC applications are applications built of two or more high-level languages (HLLs) and frequently assembler. ILC applications run outside the realm of a single language's environment, which creates special conditions, such as how the languages' data maps across load module boundaries, how conditions are handled, or how data can be called and received by each language.

z/OS MVS Program Management: Advanced Facilities, SA22-7644-12 [UPD]

z/OS MVS Program Management: Advanced Facilities contains the programming interfaces provided by the program management component of z/OS. It describes the binder application programming interface, macros, user exits, and buffer formats.

z/OS MVS Program Management: User's Guide and Reference, SA22-7643-10 [UPD]

z/OS MVS Program Management: User's Guide and Reference describes the end user interfaces provided by the program management component of z/OS.

Diagnosis

These publications provide information to help identify, describe, report, and correct a programming problem in a z/VM system.

z/VM V6R2 CMS and REXX/VM Messages and Codes, GC24-6161-01

z/VM: CMS and REXX/VM Messages and Codes describes the messages and codes generated by the CMS and REXX/VM components of z/VM. It explains the conditions that cause each message or code, describes the resulting system action, and suggests the proper user, operator, or programmer response.

z/VM V6R2 CP Messages and Codes, GC24-6177-02

z/VM: CP Messages and Codes describes the messages and codes generated by the CP component of z/VM. It explains the conditions that cause each message or code, describes the resulting system action, and suggests the proper user, operator, or programmer response.

z/VM V6R2 Diagnosis Guide, GC24-6187-01

z/VM: Diagnosis Guide provides diagnostic guidance information to help you identify, report, solve, and collect information about problems that might occur in z/VM.

z/VM V6R1 Dump Viewing Facility, GC24-6191-00

z/VM: Dump Viewing Facility describes how to use the Dump Viewing Facility to interactively diagnose system problems. This facility allows you to display, format, and print data interactively from virtual machine dumps, as well as display and format recorded trace data.

Note: For information on analyzing CP dumps, see *z/VM: VM Dump Tool*.

z/VM V6R2 Other Components Messages and Codes, GC24-6207-01

z/VM: Other Components Messages and Codes describes the messages and codes generated by the AVS, Dump Viewing Facility, GCS, TSAF, and VMSES/E components of z/VM. It explains the conditions that cause each message or code, describes the resulting system action, and suggests the proper user, operator, or programmer response.

Note: Messages issued by the Language Environment component of z/VM are documented in *z/OS: Language Environment Run-Time Messages*.

z/VM V6R2 TCP/IP Diagnosis Guide, GC24-6235-02

z/VM: TCP/IP Diagnosis Guide is intended for system programmers who want to diagnose and report problems that might occur in TCP/IP networks.

z/VM V6R2 TCP/IP Messages and Codes, GC24-6237-02

z/VM: TCP/IP Messages and Codes is intended to help system programmers diagnose TCP/IP problems. The document lists TCP/IP messages and codes by category and identifies system actions and user responses.

z/VM V6R2 VM Dump Tool, GC24-6242-01

z/VM: VM Dump Tool describes how to use the VM Dump Tool, which assists in analyzing dump data from a dump file created by the DUMpload utility. The tool can process CP standalone dumps, CP ABEND dumps, and virtual machine dumps of a CP system. This document describes the VM Dump Tool subcommands and macros that allow you to display, locate, and format dump data interactively.

z/OS and z/VM: Hardware Configuration Definition Messages, SC33-7986-13 [UPD]

z/OS and z/VM: Hardware Configuration Definition Messages provides explanations of the messages issued by Hardware Configuration Definition (HCD).

Abstracts of the publications for the z/VM optional features

The following sections contain abstracts of the publications that provide documentation for the z/VM optional features, grouped by feature.

- “DFSMS/VM” on page 86
- “Directory Maintenance Facility for z/VM” on page 86
- “Performance Toolkit for VM” on page 87
- “RACF Security Server for z/VM” on page 87
- “RSCS Networking for z/VM” on page 88

Note: There are no IBM z/VM Single System Image Feature publications. Information about z/VM SSI clusters is included in the z/VM base product publications. See “z/VM V6R2 CP Planning and Administration, SC24-6178-02” on page 75.

Some of the publications have new editions. z/VM publications that have been updated to support the general availability (GA) of z/VM V6.2 are identified with the designation **V6R2**.

Note: z/VM publications with prior release designations support z/VM V6.2 but do not contain new or changed information.

DFSMS/VM

z/VM V6R1 DFSMS/VM Customization, SC24-6181-00

z/VM: DFSMS/VM Customization provides information on customizing DFSMS/VM to meet the needs of your installation. Topics include:

- Customizing the DFSMS/VM control file
- Using installation-wide exits
- Customizing the Interactive Storage Management Facility (ISMF)
- Performance and tuning

z/VM V6R1 DFSMS/VM Diagnosis Guide, GC24-6182-00

z/VM: DFSMS/VM Diagnosis Guide provides information about diagnosing and reporting DFSMS/VM errors.

z/VM V6R1 DFSMS/VM Messages and Codes, GC24-6183-00

z/VM: DFSMS/VM Messages and Codes contains explanations and suggested actions for messages and codes issued by DFSMS/VM.

z/VM V6R1 DFSMS/VM Planning Guide, SC24-6184-00

z/VM: DFSMS/VM Planning Guide provides information on planning for the installation and use of DFSMS/VM. It can help you make decisions about:

- Using the DFSMS/VM storage management functions in your system
- Authorizing storage administrators
- Setting up the server machines
- Generating accounting records
- Backing up secondary storage

z/VM V6R2 DFSMS/VM Removable Media Services, SC24-6185-01

z/VM: DFSMS/VM Removable Media Services provides information about planning for, using, and customizing the DFSMS/VM Removable Media Services (RMS) subsystem. RMS provides functions that support the IBM 3494 Tape Library Dataserver and IBM 3495 Tape Library Dataserver.

z/VM V6R1 DFSMS/VM Storage Administration, SC24-6186-00

z/VM: DFSMS/VM Storage Administration provides conceptual information on System Managed Storage (SMS) and guides you through the DFSMS/VM storage management and minidisk management operations. It also contains a DFSMS/VM command reference, an ISMF reference, exit information, a DFSMS/VM sample environment, and other useful reference information.

Directory Maintenance Facility for z/VM

z/VM V6R2 Directory Maintenance Facility Commands Reference, SC24-6188-02

z/VM: Directory Maintenance Facility Commands Reference provides detailed reference information on DirMaint commands. Most of the DirMaint commands have the same name and format as the z/VM directory statements they support. This document also describes the additional utilities that DirMaint provides for minidisk management and command security.

z/VM V6R2 Directory Maintenance Facility Messages, GC24-6189-01

z/VM: Directory Maintenance Facility Messages provides the explanations, failing component names, and suggested corrective actions for the messages issued by DirMaint.

z/VM V6R2 Directory Maintenance Facility Tailoring and Administration Guide, SC24-6190-02

z/VM: Directory Maintenance Facility Tailoring and Administration Guide is the primary source of information about configuring and using DirMaint. It provides information about using DirMaint programming interfaces (for example, user exits), tailoring DirMaint files to accommodate the installation's needs, and administering user IDs and other facilities.

Performance Toolkit for VM

z/VM V6R2 Performance Toolkit Guide, SC24-6209-02

z/VM: Performance Toolkit Guide describes how the Performance Toolkit for VM can assist system operators and systems programmers or analysts in the following areas:

- **System operator console operation in full screen mode.** Features are provided to facilitate the operation of VM systems, thereby improving operator efficiency and productivity.
- **Performance monitoring on VM.** An enhanced real-time performance monitor allows systems programmers to monitor system performance and to analyze bottlenecks.

z/VM V6R2 Performance Toolkit Reference, SC24-6210-02

z/VM: Performance Toolkit Reference describes the subcommands and screens provided with the Performance Toolkit for VM. It also describes the Performance Toolkit messages.

RACF Security Server for z/VM

z/VM V6R2 RACF Security Server Auditor's Guide, SC24-6212-01

z/VM: RACF Security Server Auditor's Guide describes the role of the RACF auditor and explains the auditing tools that RACF provides. Reports on system and resource use can provide the auditor with information about the basic system-security environment of an installation.

z/VM V6R2 RACF Security Server Command Language Reference, SC24-6213-01

z/VM: RACF Security Server Command Language Reference describes the syntax and functions of RACF commands.

z/VM V6R1 RACF Security Server Diagnosis Guide, GC24-6214-00

z/VM: RACF Security Server Diagnosis Guide explains how to diagnose problems that might occur in RACF.

z/VM V6R1 RACF Security Server General User's Guide, SC24-6215-00

z/VM: RACF Security Server General User's Guide teaches general users how to use RACF to protect their own minidisks, SFS files, SFS directories, and other general resources. It contains an introduction to RACF and sections that guide the user through basic security tasks.

z/VM V6R2 RACF Security Server Macros and Interfaces, SC24-6216-01

z/VM: RACF Security Server Macros and Interfaces describes the syntax and functions of RACF customization macros. It also contains information about other interfaces such as RACF SMF records and subcodes of CP DIAGNOSE code

X'A0'. This document does not describe the RACROUTE macro or the independent RACF system macros, which are described in *z/VM: Security Server RACROUTE Macro Reference*.

z/VM V6R2 RACF Security Server Messages and Codes, GC24-6217-01

z/VM: RACF Security Server Messages and Codes describes RACF messages, abend codes, manager return codes, and utility return codes.

z/VM V6R2 RACF Security Server Security Administrator's Guide, SC24-6218-01

z/VM: RACF Security Server Security Administrator's Guide provides information to help the RACF security administrator to:

- Plan how to use RACF, which includes deciding which resources are protected and which users and groups are known to RACF
- Perform daily administration tasks, such as giving users access to the system (assigning user IDs and passwords) and giving users access to system resources or functions
- Coordinate with other administrators, such as the tape librarian

z/VM V6R2 RACF Security Server System Programmer's Guide, SC24-6219-01

z/VM: RACF Security Server System Programmer's Guide provides information for system programmers or installation personnel responsible for:

- Maintaining RACF databases
- Writing, testing, and installing RACF exits
- Modifying RACF to satisfy an installation's particular needs

z/VM V6R2 Security Server RACROUTE Macro Reference, SC24-6231-01

z/VM: Security Server RACROUTE Macro Reference describes the full-function RACROUTE external security interface for z/VM, the requests that can be invoked by it, their respective syntax, and related information. This document is intended for programmers who are writing applications that need to invoke RACF (or another external security manager) from z/VM, or programmers who are writing an external security manager for z/VM.

RSCS Networking for z/VM

z/VM V6R1 RSCS Networking Diagnosis, GC24-6223-00

z/VM: RSCS Networking Diagnosis describes the functions of RSCS and RSCS diagnostic aids and facilities. This document is intended for people who are responsible for isolating and diagnosing any problems that might occur in RSCS operation.

z/VM V6R1 RSCS Networking Exit Customization, SC24-6224-00

z/VM: RSCS Networking Exit Customization describes the RSCS exit facilities. This document describes how you can customize RSCS processing to meet the functional requirements of your facility.

z/VM V6R2 RSCS Networking Messages and Codes, GC24-6225-01

z/VM: RSCS Networking Messages and Codes provides explanations of the messages and abend codes for RSCS and RSCS facilities.

z/VM V6R2 RSCS Networking Operation and Use, SC24-6226-01

z/VM: RSCS Networking Operation and Use explains how to operate RSCS and how to use it to enable your z/VM system to communicate with other systems. This document describes how to send data across the network, how to define and manage links, and how to automate your network. It explains how to operate and use the RSCS server, RSCS domain name server, RSCS dynamic authorization, and the RSCS Data Interchange Manager. This document also describes the syntax and functions of the RSCS commands and link operational parameters and the commands for the other RSCS facilities.

z/VM V6R2 RSCS Networking Planning and Configuration, SC24-6227-01

z/VM: RSCS Networking Planning and Configuration provides a high-level description of RSCS and provides information to help you plan for your RSCS network. This document describes planning tasks for RSCS configuration, operation, administration, and diagnosis. It also describes the syntax and functions of the configuration file statements for the various types of RSCS servers.

Appendix A. IBM servers supported by z/VM

This section provides information about the IBM servers supported by current releases of z/VM. For guest considerations, see Appendix B, “IBM operating systems supported as guests of z/VM,” on page 95.

The following general notes apply to this matrix:

- The Processor Resource/Systems Manager™ (PR/SM) hardware feature supports the creation of multiple logical partitions (LPARs) on a single central processor complex (CPC), dividing server resources across the LPARs. Each LPAR supports an independent operating system. z/VM supports three LPAR types (modes):
 - ESA/390
 - LINUX only
 - z/VM (available only on the IBM System z10 and later servers)
- z/VM supports guest use of IBM mainframe specialty processors:
 - IBM Integrated Facility for Linux (IFL)
 - IBM Internal Coupling Facility (ICF)
 - IBM System z Application Assist Processor (zAAP)
 - IBM System z Integrated Information Processor (zIIP)
- A LINUX only LPAR, whether the allocated processors are IFL processors or general purpose processors, will not support any of the IBM traditional operating systems (such as z/OS, z/TPF, or z/VSE). Only Linux on System z or z/VM can run in a LINUX only LPAR.
- An IFL processor can be defined only in a LINUX only or z/VM LPAR.
- z/VM may run on IFL processors only if, on the IFL processors, z/VM is being used exclusively to run (1) Linux or OpenSolaris workloads and, (2) if required, z/VM applications in support of those Linux or OpenSolaris workloads.
- z/VM can exploit the ESCON Multiple Image Facility (EMIF) when running in an LPAR.
- If an APAR is required for a z/VM release to support a function on a particular server, that APAR might also be required for the same z/VM release to support the same function on a newer server.
- To use the integrated 3270 console task on the Hardware Management Console as a system operator console for z/VM, the server must be running at the current Licensed Internal Code (LIC) level with the Hardware Management Console at V1.8.0 or later.
- Before installing z/VM, refer to the Preventive Service Planning (PSP) bucket for your System z server for the minimum MCL level and any required updates:

PSP Bucket	System z Server
2817DEVICE	IBM zEnterprise 196
2818DEVICE	IBM zEnterprise 114
2097DEVICE	IBM System z10 Enterprise Class
2098DEVICE	IBM System z10 Business Class
2094DEVICE	IBM System z9 Enterprise Class
2096DEVICE	IBM System z9 Business Class
2084DEVICE	IBM eServer™ zSeries 990
2086DEVICE	IBM eServer zSeries 890
2064DEVICE	IBM eServer zSeries 900
2066DEVICE	IBM eServer zSeries 800

Server Support

The PSP buckets can be found at IBM Resource Link (www.ibm.com/servers/resourcelink/). A Resource Link user ID and password are required.

- The inclusion of a server in this matrix does not imply that the server is available in all countries.
- For information about IBM servers and z/VM releases not listed in this matrix, see the previous editions of this document.

Table 2. Server support

Server	Mode	z/VM Support
zEnterprise 196 (z196)	ESA/390	V5.4 (1,2,3,4,5,6,7), V6.1 (1,2,3,4,5,6,7), V6.2 (2)
	LINUX only	V5.4 (1,2,3,4,5,6,7), V6.1 (1,2,3,4,5,6,7), V6.2 (2)
	z/VM	V5.4 (1,2,3,4,5,6,7), V6.1 (1,2,3,4,5,6,7), V6.2 (2)
zEnterprise 114 (z114)	ESA/390	V5.4 (1,2,3,4,5,6,7,8,9,10,11), V6.1 (1,2,3,4,5,6,7,8,9,10,11), V6.2 (2,8)
	LINUX only	V5.4 (1,2,3,4,5,6,7,8,9,10,11), V6.1 (1,2,3,4,5,6,7,8,9,10,11), V6.2 (2,8)
	z/VM	V5.4 (1,2,3,4,5,6,7,8,9,10,11), V6.1 (1,2,3,4,5,6,7,8,9,10,11), V6.2 (2,8)
<p>General z/VM support for zEnterprise:</p> <ul style="list-style-type: none"> • Up to 60 LPARs. • Up to 256 GB of real storage (memory) per LPAR (z/VM image) and more than 1 TB of total virtual memory in use by guests. (The actual amount of usable real and virtual memory is dependent on the amount of real memory in the LPAR; the server model, firmware level, and configuration; and the number of guests and their workload characteristics.) • Up to 32 real processors (up to 64 virtual processors) per z/VM image. • Guest use of the configurable Crypto Express3 Accelerator. <p>Specific support notes for zEnterprise:</p> <p>(1) Support for guest exploitation of selected features requires APARs VM64798, VM64879, VM64881, and VM64947.</p> <p>(2) EREP support requires APAR VM64807.</p> <p>(3) CMS IOCP support requires APAR VM64799.</p> <p>(4) HCD support requires APAR VM64672.</p> <p>(5) HCM support requires APAR VM64747.</p> <p>(6) HLASM support requires APAR PK97799.</p> <p>(7) Support for protected key CPACF requires APAR VM64793.</p> <p>(8) EREP support for the z114 requires APAR VM64928.</p> <p>(9) CMS IOCP support for the z114 requires APAR VM64940.</p> <p>(10) HCD support for the z114 requires APAR VM64856.</p> <p>(11) HCM support for the z114 requires APAR VM64880.</p>		
System z10 Enterprise Class (z10 EC)	ESA/390	V5.4 (2,3,4,5), V6.1 (5), V6.2
	LINUX only	V5.4 (2,3,4,5), V6.1 (5), V6.2
	z/VM	V5.4 (2,3,4,5), V6.1 (5), V6.2
System z10 Business Class (z10 BC)	ESA/390	V5.4 (1,2,3,4,5), V6.1 (5), V6.2
	LINUX only	V5.4 (1,2,3,4,5), V6.1 (5), V6.2
	z/VM	V5.4 (1,2,3,4,5), V6.1 (5), V6.2

Table 2. Server support (continued)

Server	Mode	z/VM Support
General z/VM support for System z10:		
<ul style="list-style-type: none"> Up to 60 LPARs. Up to 256 GB of real storage (memory) per LPAR (z/VM image) and more than 1 TB of total virtual memory in use by guests. (The actual amount of usable real and virtual memory is dependent on the amount of real memory in the LPAR; the server model, firmware level, and configuration; and the number of guests and their workload characteristics.) Up to 32 real processors (up to 64 virtual processors) per z/VM image. Guest use of the configurable Crypto Express2 Accelerator and Crypto Express3 Accelerator. 		
Specific support notes for System z10:		
(1)	EREP support for the z10 BC requires APAR VM64475.	
(2)	CMS IOCP support for the z10 BC and the enhanced z10 EC requires APAR VM64474.	
(3)	HCD/HCM support for the z10 BC and the enhanced z10 EC requires APAR VM64410.	
(4)	OSA/SF support for the z10 BC and the enhanced z10 EC requires APAR OA26286.	
(5)	Support for protected key CPACF requires APAR VM64793.	
System z9 Enterprise Class (z9 EC)	ESA/390	V5.4
	LINUX only	V5.4
System z9 Business Class (z9 BC)	ESA/390	V5.4
	LINUX only	V5.4
General z/VM support for System z9:		
<ul style="list-style-type: none"> z/VM supports guest use of the configurable Crypto Express2 Accelerator. 		
zSeries 990 (z990)	ESA/390	V5.4
	LINUX only	V5.4
zSeries 890 (z890)	ESA/390	V5.4
	LINUX only	V5.4
zSeries 900 (z900)	Basic	V5.4
	ESA/390	V5.4
	LINUX only	V5.4
zSeries 800 (z800)	Basic	V5.4
	ESA/390	V5.4
	LINUX only	V5.4
General z/VM support for zSeries:		
<ul style="list-style-type: none"> z/VM supports guest use of the PCI Cryptographic Accelerator, PCI Cryptographic Coprocessor, PCIX Cryptographic Coprocessor, and Crypto Express2 features, where available. z/VM supports guest use of the CMOS Cryptographic Processor Feature (CCF) for the z900 and z800, except in a LINUX only LPAR. 		

Appendix B. IBM operating systems supported as guests of z/VM

This section provides information about the IBM operating systems supported as guests of current z/VM releases.

In general, an operating system is supported as a guest of z/VM only where support has been announced for that operating system to run on the server. For example, if an operating system is not supported to run in a LINUX only LPAR, that operating system is not supported as a guest of z/VM running in a LINUX only LPAR.

This section is not intended to completely cover the complex issues involved in hardware support. For server considerations, see Appendix A, "IBM servers supported by z/VM," on page 91.

The following general notes apply to the matrix:

- All listed guests are supported in ESA virtual machines. XA virtual machines are supported for compatibility and are functionally equivalent to ESA virtual machines.
- Guests may be uniprocessor or multiprocessor.
- For information about product releases not listed, see the previous editions of this document.

Table 3. Guest support

Guest	z/VM Host	Support Notes
Linux on System z	V6.2 V6.1 V5.4	
z/OS V1.10 or later	V6.2 V6.1 V5.4	(1,2) (1,2) (1,2)
Notes for z/OS:		
(1)	Exploitation of hardware function by z/OS might require specific levels of z/VM. For details, see <i>z/OS: Planning for Installation</i> , GA22-7504.	
(2)	z/VM does not support external coupling facilities or real coupling links.	
z/TPF V1.1	V6.2 V6.1 V5.4	(1) (1) (1)
Notes for z/TPF:		
(1)	z/TPF is supported as a guest of z/VM in environments where TPF-specific server and DASD control unit RPQs are not required. Use of the Emulation Program for 37xx communication controllers is not supported.	
z/VSE V4.2 or later	V6.2 V6.1 V5.4	
z/VM V6.2	V6.2 V6.1 V5.4	(2,3) (2,3) (1,2,3)
z/VM V6.1	V6.2 V6.1 V5.4	(2,3,4) (2,3) (1,2,3)

Guest Support

Table 3. Guest support (continued)

Guest	z/VM Host	Support Notes
z/VM V5.4	V6.2	(2,3,4)
	V6.1	(2,3)
	V5.4	(2,3)
Notes for z/VM:		
(1)	Supported only when the z/VM V5.4 host is running on a System z10 or later.	
(2)	Running z/VM in a virtual machine (z/VM as a guest of z/VM, also known as “second-level”) is functionally supported but is intended only for testing purposes for the second-level z/VM system and its guests (called “third-level” guests). Performance of the third-level guests will not compare favorably to what would be achieved running those guest workloads second-level.	
(3)	z/VM Guest Coupling Simulation will not work in the z/VM guest.	
(4)	APAR VM64980 should be installed on the z/VM V6.1 or V5.4 guest system. (This APAR is included in z/VM V6.2.) The fix addresses a rare problem that can occur on a second-level system running on top of V6.2 when a guest of the second-level system is using floating-point instructions. The problem occurs only if the guest executes certain floating-point instructions for which the guest has been told by STLFE facilities that bits are not available to it.	

Appendix C. IBM devices supported by z/VM

This section provides information about the IBM devices supported by current releases of z/VM.

The device support matrix is intended to provide a quick reference for support of various devices on z/VM. You should also check the hardware device support list for your server to be sure the hardware supports a particular device. Other devices not found in the device support matrix might be supported; consult Hardware Announcements and device documentation for this information.

Device support might not be in the base z/VM release, but might have been delivered through the service stream by the PTF for an APAR.

Devices are fully supported by z/VM unless otherwise indicated. z/VM supports devices at any of the following three levels of function:

1. **Full support:** z/VM uses the device itself and makes it available to guests. z/VM will usually manage the device on behalf of all users.
2. **Dedicated to a guest or virtual machine:** z/VM does not use the device itself but enables guests or applications in virtual machines to use the device. There are usually no restrictions on the use of the device by a guest or application, except that the device may not be dedicated to more than one virtual machine at a time.
3. **Dedicated to a guest or virtual machine as an unsupported device:** z/VM does not use the device itself but enables guests or applications in virtual machines to use the device. The guest or application is completely responsible for the management and use of the device. Usage restrictions might apply; the major common usage restriction for this type of support is that the device cannot be used as the guest's IPL device.

The device support matrix might show that support for a device was not provided after a particular z/VM release. The following clarifications on the meaning of the programming support withdrawal might be helpful:

- Where programming support for a *real* device has been withdrawn as of some z/VM release, the corresponding *virtual* device remains supported in the same or later z/VM releases.
- There are some device types which can be defined as, or use many of the device attributes of, the real devices for which programming support is being terminated. z/VM support for these currently supported devices remains unchanged.

Additional device restrictions might apply. There might be exceptions to supported devices and releases depending on your specific configuration. See *z/VM: CP Planning and Administration* for your release and applicable Hardware Announcements.

It is sometimes possible to define a device differently depending on whether the definition is done using the system configuration file or dynamic I/O configuration support. For instance, a 3745 Communications Controller would be defined as a 3705 Communications Controller using the system configuration file but as a 3745 Communications Controller using dynamic I/O configuration. For information about the system configuration file, see *z/VM: CP Planning and Administration*. For information about dynamic I/O configuration, see *z/VM: I/O Configuration*.

Device support

The following general notes apply to the matrix:

- z/VM does not support 7-track tapes.
- For information about IBM devices and z/VM releases not listed in this matrix, see the previous editions of this document.
- If a z/VM-related problem occurs with an OEM device that is equivalent to a supported IBM device listed in this section, IBM will fix the problem if it can be recreated with the corresponding IBM device. Customers should always contact the OEM manufacturer first about any problem encountered with an OEM device to find out if it is a known OEM problem.

Direct access storage devices (DASD)

Table 4. Device support – DASD

Device	z/VM Support	Support Notes
RAMAC Array DASD: 9391 Model A10 Rack with 9392 Model B13 Drawers	V5.4, V6.1, V6.2	Must be defined as a 3390 attached to a 3990 Model 3 or Model 6 Storage Control. Supported for installation of z/VM only when emulating a 3390 Model 3.
RAMAC Virtual Array Subsystem: 9393 Model 002, T42, or T82	V5.4, V6.1, V6.2	Emulates a 3990 Model 3 Storage Control with 3390 Model 1, 2, and 3 DASD or 3380 Model J and K DASD. Supported for installation of z/VM only when emulating a 3390 Model 3.
RAMAC Array Subsystem: 9394 Model 001, 002, or 003 Cluster Array Controller with 9395 Model B13 or B23 Drawers	V5.4, V6.1, V6.2	Emulates a 3990 Model 2 Storage Control with 3390 Model 3 DASD or 3380 Model K DASD. Supported for installation of z/VM only when emulating a 3390 Model 3.

Table 4. Device support – DASD (continued)

Device	z/VM Support	Support Notes
TotalStorage Enterprise Storage Server [®] (2105) [ESCON or FICON attached]	V5.4, V6.1, V6.2	<p>z/VM automatically virtualizes the controller mode that a guest supports.</p> <p>Models 750 and 800 supported in 2105 native controller mode.</p> <p>Supported in emulation mode as a 3990 Model 3 or 6 Storage Control with the following DASD device types:</p> <ul style="list-style-type: none"> • 3390 Model 2, 3, and 9 DASD (including large Model 9s known as Model 27 and Model 54) • 3390 Model 2 and 3 DASD in 3380 track compatibility mode <p>Only 3390 Model 3 or 9 is supported for installation of z/VM. Mixed 3390 models are not supported for installation of z/VM.</p> <p>Performance-oriented track level commands supported for guest use only.</p> <p>Parallel Access Volumes (PAVs) supported as minidisks for guests that exploit PAV architecture.</p> <p>DFSMS/VM FL221 support for 3390 Model 9 disks larger than 10017 cylinders requires APAR VM63004.</p> <p>For guest use of FlashCopy, the device must be dedicated to the guest or defined as a fullpack minidisk.</p> <p>PPRC supported for guest use only. Native use of PPRC requires ICKDSF V1.17.</p> <p>For ICKDSF support information, see the ICKDSF Program Directory.</p>
TotalStorage Enterprise Storage Server (2105) [FCP attached]	V5.4, V6.1, V6.2	<p>For native z/VM use, the ESS volumes appear as 9336-20 DASD.</p> <p>For guest use, the ESS volumes can appear either as 9336-20 DASD or as SCSI LUNs directly accessed through FCP subchannels.</p>

Device support

Table 4. Device support – DASD (continued)

Device	z/VM Support	Support Notes
System Storage DS8000 series (2107) [ESCON or FICON attached]	V5.4, V6.1, V6.2	<p>z/VM automatically virtualizes the controller mode that a guest supports.</p> <p>Supported in 2107 native controller mode.</p> <p>Supported in emulation mode as a 3990 Model 3 or 6 Storage Control with the following DASD device types:</p> <ul style="list-style-type: none"> • 3390 Model 2, 3, and 9 DASD (including large Model 9s known as Model 27 and Model 54) • 3390 Model 2 and 3 DASD in 3380 track compatibility mode <p>Only 3390 Model 3 or 9 is supported for installation of z/VM. Mixed 3390 models are not supported for installation of z/VM.</p> <p>Dynamic volume expansion is supported.</p> <p>Extended address 3390-A devices are supported. z/VM V5.4 and V6.1 require APARs VM64709 and VM64711. A 3390-A volume with more than 65,520 cylinders must be dedicated to a guest or defined as a fullpack minidisk.</p> <p>Performance-oriented track level commands supported for guest use only.</p> <p>Parallel Access Volumes (PAVs) supported as minidisks for guests that exploit PAV architecture.</p> <p>On FICON-attached DS8000, HyperPAV volumes supported as minidisks for guests that exploit this PAV architecture.</p> <p>(continued)</p>
DS8000 [ESCON or FICON] (continued)		<p>DFSMS/VM FL221 support for 3390 Model 9 disks larger than 10017 cylinders requires APAR VM63004.</p> <p>For guest use of FlashCopy, the device must be dedicated to the guest or defined as a fullpack minidisk. On z/VM V5.4, support for the FlashCopy/SE feature requires APAR VM64449.</p> <p>PPRC supported for guest use only. Native use of PPRC requires ICKDSF V1.17.</p> <p>For ICKDSF support information, see the ICKDSF Program Directory.</p>

Table 4. Device support – DASD (continued)

Device	z/VM Support	Support Notes
System Storage DS8000 series (2107) [FCP attached]	V5.4, V6.1, V6.2	<p>For native z/VM use, the DS8000 volumes appear as 9336-20 DASD.</p> <p>For guest use, the DS8000 volumes can appear either as 9336-20 DASD or as SCSI LUNs directly accessed through FCP subchannels.</p> <p>DFSMS/VM FL221 support for SCSI LUNs larger than 381 GB requires APAR VM63664.</p>
System Storage DS6000 series (1750) [FICON attached]	V5.4, V6.1, V6.2	<p>z/VM automatically virtualizes the controller mode that a guest supports.</p> <p>Supported in 1750 native controller mode.</p> <p>Supported in emulation mode as a 3390 Model 3 or 6 Storage Control with the following DASD device types:</p> <ul style="list-style-type: none"> • 3390 Model 2, 3, and 9 DASD (including large Model 9s known as Model 27 and Model 54) • 3390 Model 2 and 3 DASD in 3380 track compatibility mode <p>Only 3390 Model 3 or 9 is supported for installation of z/VM. Mixed 3390 models are not supported for installation of z/VM.</p> <p>Performance-oriented track level commands supported for guest use only.</p> <p>Parallel Access Volumes (PAVs) supported as minidisks for guests that exploit PAV architecture.</p> <p>(continued)</p>
DS6000 [ESCON or FICON] (continued)		<p>DFSMS/VM FL221 support for 3390 Model 9 disks larger than 10017 cylinders requires APAR VM63004.</p> <p>For guest use of FlashCopy, the device must be dedicated to the guest or defined as a fullpack minidisk.</p> <p>PPRC supported for guest use only. Native use of PPRC requires ICKDSF V1.17.</p> <p>For ICKDSF support information, see the ICKDSF Program Directory.</p>

Device support

Table 4. Device support – DASD (continued)

Device	z/VM Support	Support Notes
System Storage DS6000 series (1750) [FCP attached]	V5.4, V6.1, V6.2	<p>For native z/VM use, the DS6000 volumes appear as 9336-20 DASD.</p> <p>For guest use, the DS6000 volumes can appear either as 9336-20 DASD or as SCSI LUNs directly accessed through FCP subchannels.</p> <p>DFSMS/VM FL221 support for SCSI LUNs larger than 381 GB requires APAR VM63664.</p>
XIV Storage System	V5.4, V6.1, V6.2	<p>Direct attachment for system use on z/VM V5.4 and V6.1 requires APAR VM64708.</p> <p>HCD support on z/VM V5.4 and V6.1 requires APAR VM64672.</p> <p>ICKDSF support on z/VM V5.4 and V6.1 requires APAR PM11220.</p>

DASD control units and storage controls

Table 5. Device support – DASD control units and storage controls

Device	z/VM Releases Providing Support	Support Notes
3990 Model 6	V5.4, V6.1, V6.2	<p>All functions of former 3990 models supported plus increased cache, System Clock Support, and the Control Unit Initiated Reconfiguration feature (CUIR).</p> <p>Concurrent Copy and Extended Remote Copy functions are supported for guest use only.</p>

Tape units and tape libraries

Table 6. Device support – tape units and tape libraries

Device	z/VM Releases Providing Support	Support Notes
3480 Tape Subsystem	V5.4, V6.1, V6.2	Not supported for installation of z/VM.
3490 Tape Subsystem	V5.4, V6.1, V6.2	Not supported for installation of z/VM.
3490E (Enhanced Capability Model) Tape Subsystem	V5.4, V6.1, V6.2	Not supported for installation of z/VM.
TotalStorage Enterprise Automated Tape Library (3494)	V5.4, V6.1, V6.2	<p>Native support is provided in conjunction with DFSMS/VM.</p> <p>Not supported for installation of z/VM.</p>

Table 6. Device support – tape units and tape libraries (continued)

Device	z/VM Releases Providing Support	Support Notes
TotalStorage Virtual Tape Server (3494 VTS)	V5.4, V6.1, V6.2	Native support is provided in conjunction with DFSMS/VM. Not supported for installation of z/VM. Peer-to-Peer VTS and VTS Import/Export are supported for guests.
3495 Tape Library Dataserver	V5.4, V6.1, V6.2	Native support is provided in conjunction with DFSMS/VM. Not supported for installation of z/VM.
TotalStorage Enterprise Tape System (3590)	V5.4, V6.1, V6.2	DFSMS/VM FL221 support for the 3590 Model H requires APAR VM63097.
System Storage TS1120 Tape Drive (3592)	V5.4, V6.1, V6.2	z/VM supports drive-based data encryption, including rekey support, with the 3592 Model E05. On z/VM V5.4, support for the 3592 Model E06 requires APAR VM64459. On z/VM V5.4 and V6.1, support for the 3592 Model E07 requires APAR VM64979. DFSMS/VM FL221 requires: <ul style="list-style-type: none"> • APAR VM63353 for basic device support. • APAR VM63460 to support Write Once Read Many (WORM) media. • APAR VM64458 to support the 3592 Model E06. • APAR VM65005 to support the 3592 Model E07.

Tape control units

Table 7. Device support – tape control units

Device	z/VM Releases Providing Support	Support Notes
TotalStorage Enterprise Tape Controller 3590 Model A60	V5.4, V6.1, V6.2	

Printers

Table 8. Device support – printers

Device	z/VM Releases Providing Support	Support Notes
3203 Model 5	V5.4, V6.1, V6.2	
3262	V5.4, V6.1, V6.2	
3268 Model 2, 2C	V5.4, V6.1, V6.2	Must be defined as a 3287.
3287 Model 1, 1C, 2, 2C, 4	V5.4, V6.1, V6.2	

Device support

Table 8. Device support – printers (continued)

Device	z/VM Releases Providing Support	Support Notes
3289 Model 1, 3, 4, 8	V5.4, V6.1, V6.2	
3800 Model 1	V5.4, V6.1, V6.2	
3800 Model 3, 6, 8	V5.4, V6.1, V6.2	Full support in Model 1 compatibility mode or using Advanced Function Printing (AFP) programs.
3812	V5.4, V6.1, V6.2	Full support through the RSCS feature, using AFP programs; downloading of fonts is not supported. Loading of fonts requires Print Services Facility/VM (PSF/VM) licensed program (program number 5684-141).
3816 Model 01D, 01S	V5.4, V6.1, V6.2	Full support through the RSCS feature, using AFP programs; downloading of fonts is not supported. Loading of fonts requires PSF/VM licensed program.
3820	V5.4, V6.1, V6.2	Full support through ACF/VTAM V4.2, using AFP programs.
3825	V5.4, V6.1, V6.2	Full support using AFP programs.
3827	V5.4, V6.1, V6.2	Full support using AFP programs.
3835	V5.4, V6.1, V6.2	Full support using AFP programs.
3900	V5.4, V6.1, V6.2	Must be defined as a 3800.
4245 Model 1	V5.4, V6.1, V6.2	Supported as a 4245 or in 3262 compatibility mode.
4245 Model 12, 20	V5.4, V6.1, V6.2	
4248 Model 1, 2	V5.4, V6.1, V6.2	Supported as a 4248 or in 3211 compatibility mode.
6262 Model 14, 22	V5.4, V6.1, V6.2	Must be defined as a 4248 Model 1.

Card readers and card punches

Table 9. Device support – card readers and card punches

Device	z/VM Releases Providing Support	Support Notes
3505 Model B1, B2	V5.4, V6.1, V6.2	
3525 Model P1, P2, P3	V5.4, V6.1, V6.2	

Terminals, displays, and consoles

Most terminals are supported as a virtual machine console (in 3215 emulation mode or 3270 mode). 3270-family displays can be defined generically (for example, as 3270s) with dynamic computation of screen size based on information returned from the device.

z/VM also supports the integrated 3270 console and integrated ASCII console functions of the Hardware Management Console.

Display printers

Most display printers are supported using Advanced Function Printing (AFP) programs through RSCS.

Display control units

Table 10. Device support – display control units

Device	z/VM Releases Providing Support	Support Notes
3174	V5.4, V6.1, V6.2	
3272 Model 2	V5.4, V6.1, V6.2	
3274	V5.4, V6.1, V6.2	
3276 Display/Control Unit	V5.4, V6.1, V6.2	Must be dedicated to a virtual machine as a supported device.

Communications controllers

Table 11. Device support – communications controllers

Device	z/VM Releases Providing Support	Support Notes
3745	V5.4, V6.1, V6.2	Must be dedicated to a virtual machine as a supported device. Must be defined as a 3705, unless you are using dynamic system configuration.

Switches

Table 12. Device support – switches

Device	z/VM Releases Providing Support	Support Notes
2032 FICON Director Model 1	V5.4, V6.1, V6.2	
9032 ESCON Director Model 2	V5.4, V6.1, V6.2	
9032 ESCON Director Model 3	V5.4, V6.1, V6.2	
9032 ESCON Director Model 5	V5.4, V6.1, V6.2	
9033 ESCON Director Model 1	V5.4, V6.1, V6.2	
9033 ESCON Director Model 4	V5.4, V6.1, V6.2	

Note: Control units and I/O devices with parallel channels can be attached to ESCON channels through the ESCON Converter Model 1. (See your IBM representative for a list of parallel devices that the ESCON Converter Model 1 supports.) Conversely, the ESCON Converter Model 2 allows ESCON devices to be attached to servers that have parallel channels.

For more information about ESCON architecture, see *Introducing Enterprise Systems Connection*, GA23-0383.

Server adapters and miscellaneous devices

Table 13. Device support – server adapters and miscellaneous devices

Adapter/Device	z/VM Releases Providing Support	Support Notes
Channel-to-Channel Adapter (CTCA)	V5.4, V6.1, V6.2	Must be dedicated to a virtual machine as a supported device.

Device support

Table 13. Device support – server adapters and miscellaneous devices (continued)

Adapter/Device	z/VM Releases Providing Support	Support Notes
Crypto Express2	V5.4, V6.1, V6.2	<p>Must be configured for shared or dedicated access by a virtual machine.</p> <p>Configurable Crypto Express2 feature (Coprocesor or Accelerator) supported on z10 and z9 servers.</p>
Crypto Express3	V5.4, V6.1, V6.2	<p>Must be configured for shared or dedicated access by a virtual machine.</p> <p>Configurable Crypto Express3 feature (Coprocesor or Accelerator) supported on z10 and later servers.</p> <p>z/VM V5.4, and V6.1 require APAR VM64656.</p>
ESCON CTCA	V5.4, V6.1, V6.2	Extended mode operation only.
FICON CTCA	V5.4, V6.1, V6.2	Must be dedicated to a guest as a supported device.
FICON Express2	V5.4, V6.1, V6.2	Must be dedicated to a guest as a supported device.
FICON Express4	V5.4, V6.1, V6.2	Must be dedicated to a guest as a supported device.
FICON Express8, Express8S	V5.4, V6.1, V6.2	Must be dedicated to a guest as a supported device.
HiperSockets	V5.4, V6.1, V6.2	Must be dedicated to a virtual machine as a supported device.
OSA-2	V5.4	<p>Must be dedicated to a virtual machine as a supported device.</p> <p>OSA-2 is recognized as TYPE=OSA.</p> <p>OSA/SF is required to customize the OSA modes. OSA/SF can be running either in an LPAR in the same system complex or as a z/VM guest. OSA/SF for z/VM can be run as a CMS application.</p>

Table 13. Device support – server adapters and miscellaneous devices (continued)

Adapter/Device	z/VM Releases Providing Support	Support Notes
OSA-Express	V5.4	<p>Must be dedicated to a virtual machine as a supported device.</p> <p>OSA-Express is recognized as:</p> <ul style="list-style-type: none"> • TYPE=OSE for the Fast Ethernet and 155 ATM modes • TYPE=OSD for the Gigabit Ethernet mode • TYPE=OSC for the 1000BASE-T Ethernet feature with the Integrated Console Controller (ICC) function <p>OSA/SF is required to customize the OSA-Express modes. OSA/SF can be running either in an LPAR in the same system complex or as a z/VM guest. OSA/SF for z/VM can be run as a CMS application.</p>
OSA-Express2	V5.4, V6.1, V6.2	<p>Must be dedicated to a virtual machine as a supported device.</p> <p>OSA-Express2 is recognized as:</p> <ul style="list-style-type: none"> • TYPE=OSD for the Gigabit Ethernet and 10 Gigabit Ethernet modes. z/VM supports IEEE 802.3ad link aggregation. • TYPE=OSC for the 1000BASE-T Ethernet feature with the Integrated Console Controller (ICC) function. • TYPE=OSN for the Gigabit Ethernet and 1000BASE-T Ethernet features with the Open Systems Adapter for the Network Control Program (OSA NCP) function. <p>z/VM supports isolation of QDIO data connection communications on a shared OSA-Express2 port on a z10 or z9 server in support of multitier security zones. z/VM V5.4 requires APARs VM64463 and PK67610.</p> <p>OSA/SF is required to customize the OSA-Express2 modes. OSA/SF can be running either in an LPAR in the same system complex or as a z/VM guest. OSA/SF for z/VM can be run as a CMS application.</p>

Device support

Table 13. Device support – server adapters and miscellaneous devices (continued)

Adapter/Device	z/VM Releases Providing Support	Support Notes
OSA-Express3	V5.4, V6.1, V6.2	<p>Must be dedicated to a virtual machine as a supported device.</p> <p>OSA-Express3 is recognized as:</p> <ul style="list-style-type: none"> • TYPE=OSD for the Gigabit Ethernet, 10 Gigabit Ethernet, and 1000BASE-T modes. z/VM supports IEEE 802.3ad link aggregation. • TYPE=OSC for the 1000BASE-T Ethernet feature with the Integrated Console Controller (ICC) function. • TYPE=OSN for the Gigabit Ethernet and 1000BASE-T Ethernet features with the Open Systems Adapter for the Network Control Program (OSA NCP) function. • TYPE=OSM for the 1000BASE-T Ethernet intranode management network (INMN) in the IBM zEnterprise System. z/VM V5.4 requires APAR VM64794 for configuring via dynamic I/O (but the device cannot be varied online). z/VM V6.1 requires APAR VM64822. • TYPE=OSX for the 10 Gigabit Ethernet intraensemble data network (IEDN) in the IBM zEnterprise System. z/VM V5.4 requires APAR VM64794 for configuring via dynamic I/O (but the device cannot be varied online). z/VM V6.1 requires APAR VM64822. <p>z/VM supports four ports per feature on the Gigabit Ethernet and 1000BASE-T features on the z196, z114, and z10 EC, and two ports per feature on the z114 and z10 BC.</p> <p>z/VM supports isolation of QDIO data connection communications on a shared OSA-Express3 port on a zEnterprise or z10 server in support of multitier security zones. z/VM V5.4 requires APARs VM64463 and PK67610.</p> <p>OSA/SF is required to customize the OSA-Express3 modes. OSA/SF can be running either in an LPAR in the same system complex or as a z/VM guest. OSA/SF for z/VM can be run as a CMS application.</p>

Table 13. Device support – server adapters and miscellaneous devices (continued)

Adapter/Device	z/VM Releases Providing Support	Support Notes
OSA-Express4S	V5.4, V6.1, V6.2	<p>Must be dedicated to a virtual machine as a supported device.</p> <p>OSA-Express4S is recognized as:</p> <ul style="list-style-type: none"> • TYPE=OSD for the Gigabit Ethernet and 10 Gigabit Ethernet modes. z/VM supports IEEE 802.3ad link aggregation. • TYPE=OSX for the 10 Gigabit Ethernet intraensemble data network (IEDN) in the IBM zEnterprise System. z/VM V5.4 requires APAR VM64794 for configuring via dynamic I/O (but the device cannot be varied online). z/VM V6.1 requires APAR VM64822. <p>z/VM supports two ports on the Gigabit Ethernet feature on the z196 and z114.</p> <p>z/VM supports isolation of QDIO data connection communications on a shared OSA-Express4S port in support of multitier security zones. z/VM V5.4 requires APARs VM64463 and PK67610.</p> <p>OSA/SF with APAR OA37060 is required to customize the OSA-Express4S modes. OSA/SF can be running either in an LPAR in the same system complex or as a z/VM guest. OSA/SF for z/VM can be run as a CMS application.</p>
PCI Cryptographic Accelerator	V5.4	Must be configured for shared or dedicated access by a virtual machine.
PCI Cryptographic Coprocessor	V5.4	Must be configured for shared or dedicated access by a virtual machine.
PCIX Cryptographic Coprocessor	V5.4	Must be configured for shared or dedicated access by a virtual machine.
SAN Volume Controller 2145	V5.4, V6.1, V6.2	
3088	V5.4, V6.1, V6.2	Must be dedicated to a virtual machine as a supported device.
3088 Model 61 PSCA Card (common link access to workstations)	V5.4, V6.1, V6.2	
3174 Model 12L, 22L	V5.4, V6.1, V6.2	
3737 Remote Channel-to-Channel Unit Model 2	V5.4, V6.1, V6.2	Must be dedicated to a virtual machine as a supported device. Must be defined as a CTCA.
4753 Network Security Processor	V5.4, V6.1, V6.2	Must be dedicated to a guest as a supported device.

Device support

Table 13. Device support – server adapters and miscellaneous devices (continued)

Adapter/Device	z/VM Releases Providing Support	Support Notes
7171 Device Attachment Control Unit	V5.4, V6.1, V6.2	The 7171 ASCII Device Attachment Control Unit (DACU) allows the attachment of ASCII terminals locally or (through Start-Stop communication line) remotely. The ASCII terminals appear to the host as 3277s, 3278s, and 3279s, with the DACU itself appearing as a 3274 control unit.
9034 ESCON Converter Model 1	V5.4, V6.1, V6.2	
9035 ESCON Converter Model 2	V5.4, V6.1, V6.2	

Notices

This information was developed for products and services offered in the U.S.A.

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
Armonk, NY 10504-1785
U.S.A.

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.
1623-14, Shimotsurama, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

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 states 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 supply 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:

Site Counsel
IBM Corporation
2455 South Road
Poughkeepsie, NY 12601-5400
U.S.A.

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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

All 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 the names and addresses used by an actual business enterprise 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.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at IBM copyright and trademark information - United States (www.ibm.com/legal/us/en/copytrade.shtml).

Adobe is either a registered trademark or a trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Windows is a trademark of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java is a trademark or registered trademark of Oracle and/or its affiliates.

Index

Numerics

- 370 Accommodation Facility 45
- 370 virtual machines not supported 45
- 370-XA architecture compatibility 46

A

- abstracts of z/VM books 72
- accessibility 7
- ACF/VTAM 62
- adapter interruptions 53
- analyzing dump data 47
- APPC/VM
 - AVS 62
 - ISFC 47
 - TSAF 64
- APPC/VM VTAM Support (AVS) 62
- application development and deployment, support for
 - CMS support 56
 - enhancements 20
 - overview 6
- application programming interfaces 6
- architecture compatibility 46
- architecture level set 25
- architecture support
 - 370-XA compatibility 46
 - dynamic storage reconfiguration 48
 - ESA/370 compatibility 46
 - ESA/390 45
 - ESA/XC 46
 - ESCON 48
 - FCP 49
 - FICON 49
 - HiperSockets 49
 - IEEE Floating Point 50
 - LCSS 50
 - POSIX 57
 - QDIO 52
 - TCP/IP 59
 - z/Architecture 46
- auditability 40
- auditable events 43
- AVS 62

B

- base facilities
 - AVS 62
 - CMS 55
 - CP 45
 - Dump Viewing Facility 62
 - GCS 62
 - HCD and HCM for z/VM 62
 - Language Environment 63
 - OSA/SF 63
 - REXX/VM 64
 - TCP/IP for z/VM 59

- base facilities (*continued*)
 - TSAF 64
 - VMSES/E 64
- benefits of z/VM 1
- BFS (byte file system) 57
- books, z/VM
 - abstracts
 - base library - application programming 79
 - base library - customization and tailoring 77
 - base library - installation, migration, and service 73
 - base library - operation and use 77
 - base library - overview 73
 - base library - planning and administration 74
 - base library - system diagnosis 84
 - DFSMS/VM 86
 - DirMaint 86
 - OpenExtensions 81
 - Performance Toolkit for VM 87
 - RACF Security Server for z/VM 87
 - RSCS 88
 - basic 72
 - media 71
 - online
 - program requirements 35
 - sources 71
 - supplied with z/VM 72
- byte file system (BFS) 57

C

- cables, fiber optic 48, 49
- card readers and punches supported by z/VM 104
- changes in z/VM V6.2
 - application enablement 20
 - common service file pool 21
 - contiguous frame coalescing, enhanced 18
 - CPACF protected key 19
 - CPU-Measurement Counter Facility, host support for 13
 - Crypto Express3 19
 - Directory Maintenance Facility, enhanced 15
 - dynamic discovery of FCP disks 15
 - dynamic EDEVICE path control 13
 - EAL4+ certification for z/VM V6.1 19
 - ESM access control for real devices 18
 - ESM access control support for SET SECUSER and SET OBSERVER 19
 - extended address DASD volumes, support for 12
 - FICON Express8S 12
 - FIPS support 18
 - FTP enhancements 17
 - guest page reordering process, disabling 17
 - hard limiting of scheduled system resource, alternate method for 18
 - HyperSwap improvements 12
 - IBM z/VM Single System Image Feature 20
 - installation changes 20

- changes in z/VM V6.2 (*continued*)
 - installation process improvements 21
 - intraensemble data network (IEDN), support for 16
 - intranode management network (INMN), support for 16
 - IPv6 support 16
 - Language Environment upgrade 20
 - LDAP server upgrade 17
 - live guest relocation 10
 - memory constraint relief 17
 - MPROUTE upgrade 16
 - multiple access ports per guest 17
 - network virtualization 16
 - OSA address table support 17
 - OSA-Express4S 12
 - OSX and OSM CHPIDs 11
 - packaging changes 20
 - page release serialization, improved 18
 - Performance Toolkit, enhanced 16
 - Program Management Binder upgrade 20
 - RACF Security Server enhancements 19
 - RACF support for protected user IDs 19
 - scalability improvements 17
 - SCSIDISC utility 15
 - security enhancements 18
 - serviceability changes 20
 - shipping zMAP agents discontinued 22
 - SSL server upgrade 18
 - statement of direction 22
 - systems management APIs, enhanced 15
 - systems management improvements 13
 - technology exploitation 11
 - Unified Resource Manager, support for 14
 - VMSES/E enhancements 21
 - withdrawal of support 21
 - XIV Storage Systems, direct attachment of 13
 - XL C/C++ Compiler for z/VM, V1.3, support for 20
 - XRC time stamping 12
 - z/Architecture CMS (z/CMS) 20
 - z/VM single system image (SSI) clusters 9
 - z114 support 11
 - z196 support 11
 - zEnterprise System, support for 14
- channel subsystem
 - channel program access protection 40
- CMS (Conversational Monitor System)
 - application programming 56
 - byte file system (BFS) 57
 - ESA/390 CMS 55
 - HELP facility 59
 - multicultural support 37
 - multitasking, application 57
 - OpenExtensions 57
 - overview 55
 - Pipelines 57
 - Program Management Binder for CMS 58
 - Reusable Server Kernel 58
 - Shared File System (SFS) 55
 - systems management API 56
 - versions 55
 - Virtual Machine Resource Manager (VMRM) 58
- CMS (Conversational Monitor System) (*continued*)
 - XEDIT 59
 - z/Architecture CMS (z/CMS) 55
- Collaborative Memory Management Assist (CMMA) 47
- communications controllers supported by z/VM 105
- compatibility
 - 370-XA architecture 46
 - ESA/370 architecture 46
 - System/370 architecture 45
- components, z/VM
 - AVS 62
 - CMS 55
 - CP 45
 - Dump Viewing Facility 62
 - GCS 62
 - Language Environment 63
 - REXX/VM 64
 - TSAF 64
 - VMSES/E 64
- concurrent-sense facility support 48
- consoles supported by z/VM 104
- Control Program (CP)
 - architecture compatibility 46
 - CSE 46
 - ESA/XC architecture support 46
 - guest LAN 46
- hardware and architecture support
 - adapter interruptions 53
 - Collaborative Memory Management Assist (CMMA) 47
 - concurrent-sense facility 48
 - CPACF 48
 - CPU-Measurement Counter Facility 48
 - Crypto Express2 48
 - Crypto Express3 48
 - cryptographic facilities 48
 - dynamic storage reconfiguration 48
 - ECMDS 48
 - ESCON 48
 - expanded storage 48
 - Extended-TOD-Clock Facility 49
 - FCP 49
 - FICON channels 49
 - FICON CTCA 49
 - FlashCopy 49
 - guest ASCII console 49
 - HiperSockets 49
 - ICF 53
 - IEEE Floating Point 50
 - IFL 53
 - integrated 3270 console 50
 - integrated systems management support 50
 - LCSS 50
 - MIDAWs 51
 - Move-Page Facility 51
 - NPIV 51
 - OSA 51
 - Parallel Access Volumes (PAV) 52
 - PPRC 52
 - QDIO architecture 52
 - specialty processors 53

- Control Program (CP) *(continued)*
 - hardware and architecture support *(continued)*
 - STP 53
 - zAAP 53
 - zIIP 53
 - ISFC 47
 - multicultural support 37
 - overview 45
 - virtual machine
 - architectures 45
 - type 46
 - VM Dump Tool 47
 - z/Architecture support 46
- control units supported by z/VM
 - DASD 102
 - display 105
 - storage 102
 - tape 103
- controllers, communications, supported by z/VM 105
- Conversational Monitor System (CMS)
 - application programming 56
 - byte file system (BFS) 57
 - ESA/390 CMS 55
 - HELP facility 59
 - multicultural support 37
 - multitasking, application 57
 - OpenExtensions 57
 - overview 55
 - Pipelines 57
 - Program Management Binder for CMS 58
 - Reusable Server Kernel 58
 - Shared File System (SFS) 55
 - systems management API 56
 - versions 55
 - Virtual Machine Resource Manager (VMRM) 58
 - XEDIT 59
 - z/Architecture CMS (z/CMS) 55
- CP (Control Program)
 - architecture compatibility 46
 - CSE 46
 - ESA/XC architecture support 46
 - guest LAN 46
 - hardware and architecture support
 - adapter interruptions 53
 - Collaborative Memory Management Assist (CMMA) 47
 - concurrent-sense facility 48
 - CPACF 48
 - CPU-Measurement Counter Facility 48
 - Crypto Express2 48
 - Crypto Express3 48
 - cryptographic facilities 48
 - dynamic storage reconfiguration 48
 - ECMDS 48
 - ESCON 48
 - expanded storage 48
 - Extended-TOD-Clock Facility 49
 - FCP 49
 - FICON channels 49
 - FICON CTCA 49
 - FlashCopy 49

- CP (Control Program) *(continued)*
 - hardware and architecture support *(continued)*
 - guest ASCII console 49
 - HiperSockets 49
 - ICF 53
 - IEEE Floating Point 50
 - IFL 53
 - integrated 3270 console 50
 - integrated systems management support 50
 - LCSS 50
 - MIDAWs 51
 - Move-Page Facility 51
 - NPIV 51
 - OSA 51
 - Parallel Access Volumes (PAV) 52
 - PPRC 52
 - QDIO architecture 52
 - specialty processors 53
 - STP 53
 - zAAP 53
 - zIIP 53
 - ISFC 47
 - multicultural support 37
 - overview 45
 - virtual machine
 - architectures 45
 - type 46
 - VM Dump Tool 47
 - z/Architecture support 46
- CPU-Measurement Counter Facility 48
- cross system extensions (CSE)
 - description 46
 - hardware requirements 27
 - program requirements 34
- cryptographic facilities, support for 48
- CSE (cross system extensions)
 - description 46
 - hardware requirements 27
 - program requirements 34
- customer responsibilities for data security 40

D

- DASD (direct access storage device)
 - control units supported by z/VM 102
 - CSE requirement 27
 - HyperPAV support 52
 - Parallel Access Volumes (PAV) support 52
 - required to install z/VM 26
 - supported by z/VM 98
- data integrity
 - customer responsibilities 40
 - for guests 39
- dedicated-only device, definition of 29
- device support, types of 28
- devices supported by VM
 - card readers and punches 104
 - communications controllers 105
 - consoles 104
 - DASD 98
 - DASD control units 102

- devices supported by VM *(continued)*
 - display control units 105
 - display printers 105
 - displays 104
 - miscellaneous 105
 - printers 103
 - storage controls 102
 - switches 105
 - tape control units 103
 - tape libraries 102
 - tape units 102
 - terminals 104
- DFSMS/VM
 - books, abstract of 86
 - description 65
 - multicultural support 37
 - program requirements 31
- diagnosing system problems 62
- direct access storage device (DASD)
 - control units supported by z/VM 102
 - CSE requirement 27
 - HyperPAV support 52
 - Parallel Access Volumes (PAV) support 52
 - required to install z/VM 26
 - supported by z/VM 98
- DirMaint (Directory Maintenance Facility)
 - books, abstract of 86
 - description 65
 - multicultural support 37
 - program requirements 32
 - restrictions 36
 - security 42
 - system integrity 40
- display control units supported by z/VM 105
- display printers supported by z/VM 105
- displays supported by z/VM 104
- DOS/VSE support 56
- dump data, analyzing 47
- Dump Viewing Facility 62
- DVD
 - installing z/VM from 26
 - online books 72
 - z/VM System Image 37
- dynamic storage reconfiguration 48

E

- ECMDS (extended channel measurement data support) 48
- editor, z/VM 59
- Enterprise Systems Connection Architecture (ESCON)
 - description 48
 - devices supported 105
 - ESCON Multiple Image Facility (EMIF) 91
- ESA virtual machine 45
- ESA/370 architecture compatibility 46
- ESA/390 CMS 55
- ESA/XC architecture 46
- ESCON (Enterprise Systems Connection Architecture)
 - description 48
 - devices supported 105

- ESCON (Enterprise Systems Connection Architecture) *(continued)*
 - ESCON Multiple Image Facility (EMIF) 91
- expanded storage support 48
- extended channel measurement data support (ECMDS) 48
- Extended Configuration (ESA/XC) architecture 46
- Extended-TOD-Clock Facility support 49

F

- FCP (Fibre Channel Protocol) support 49
- features, optional
 - DFSMS/VM 65
 - DirMaint 65
 - IBM z/VM Single System Image Feature 66
 - Performance Toolkit for VM 67
 - RACF Security Server for z/VM 68
 - RSCS Networking for z/VM 69
- Federal Information Processing Standard (FIPS) support 42
- fiber optic cables 48, 49
- FICON (fibre connection) channel support 49
- FICON CTCA support 49
- FICON Express 49
- FICON Express2 49
- FICON Express4 49
- FICON Express8 49
- file pool, CMS
 - description 55
 - management using DFSMS/VM 65
 - security 41
- FIPS support 42
- first-level storage
 - minimum and maximum 26
- FlashCopy support 49
- floating point support 50

G

- GCS (Group Control System) 62
- guest ASCII console 49
- guest LAN 46
- guest operating system support
 - architecture compatibility 46
 - description 4
 - matrix 95

H

- hardware architectures and facilities, support for
 - adapter interruptions 53
 - Collaborative Memory Management Assist (CMMA) 47
 - concurrent-sense facility 48
 - CPACF 48
 - CPU-Measurement Counter Facility 48
 - Crypto Express2 48
 - Crypto Express3 48
 - cryptographic facilities 48
 - ECMDS 48

hardware architectures and facilities, support for
(continued)

- ESCON 48
- expanded storage 48
- Extended-TOD-Clock Facility 49
- FCP 49
- FICON channels 49
- FICON CTCA 49
- FlashCopy 49
- guest ASCII console 49
- HiperSockets 49
- ICF 53
- IEEE Floating Point 50
- IFL 53
- integrated 3270 console 50
- integrated systems management support 50
- LCSS 50
- MIDAWs 51
- Move-Page Facility 51
- NPIV 51
- OSA 51
- Parallel Access Volumes (PAV) 52
- PPRC 52
- QDIO architecture 52
- specialty processors 53
- STP 53
- zAAP 53
- zIIP 53

Hardware Configuration Definition and Hardware
Configuration Manager for z/VM

- description 62
- program requirements 32

Hardware Management Console 26

hardware requirements

- CSE 27
- OSA/SF 27
- Performance Toolkit 27
- server 25
- SSI clusters 27
- TCP/IP for z/VM 28
- workstation 26
- z/VM installation 25

HCD and HCM for z/VM

- description 62
- program requirements 32

HELP facility 59

High Level Assembler 30, 32

HiperSockets support 49

HyperPAV support 52

I

I/O (input/output)

- channel program access protection 40
- ESCON 48
- FICON channels 49
- FlashCopy 49

IBM Publications Center 72

ICF (Internal Coupling Facility), support for 53

IEEE Floating Point support 50

IFL (Integrated Facility for Linux), support for 53

Information Center 71

input/output (I/O)

- channel program access protection 40
- ESCON 48
- FICON channels 49
- FlashCopy 49

installation requirements 25

integrated 3270 console 50

Integrated Facility for Linux (IFL), support for 53

integrated systems management support 50

integrity

- data 39
- DirMaint 40
- system 39

Internal Coupling Facility (ICF), support for 53

IOCP program requirements 32

IPv6 support 60

ISFC (Inter-System Facility for Communications) 47

J

journaling 40

L

Language Environment

- description 63
- program requirements 32

LCSS (logical channel subsystems) 50

library, z/VM

- abstracts of books
 - base library - application programming 79
 - base library - customization and tailoring 77
 - base library - installation, migration, and service 73
 - base library - operation and use 77
 - base library - overview 73
 - base library - planning and administration 74
 - base library - system diagnosis 84
- DFSMS/VM 86
- DirMaint 86
- OpenExtensions 81
- Performance Toolkit for VM 87
- RACF Security Server for z/VM 87
- RSCS 88

- documentation supplied with z/VM 72

- sources 71

licensed programs supported on z/VM 36

limitations 35

Linux on System z

- enablement, latest z/VM enhancements for 9
- guest support for 95
- program requirements 35

logical channel subsystems (LCSS) 50

LPAR (logical partition) 91

M

main storage

- minimum and maximum 26

- memory
 - minimum and maximum 26
- MIDAW support 51
- Move-Page Facility support 51
- multicultural support 37
- multitasking, CMS application 57

N

- N_Port ID virtualization (NPIV) 51
- network interface card (NIC) 46
- networking
 - enhancements 16
- new or changed in z/VM V6.2
 - application enablement 20
 - common service file pool 21
 - contiguous frame coalescing, enhanced 18
 - CPACF protected key 19
 - CPU-Measurement Counter Facility, host support for 13
 - Crypto Express3 19
 - Directory Maintenance Facility, enhanced 15
 - dynamic discovery of FCP disks 15
 - dynamic EDEVICE path control 13
 - EAL4+ certification for z/VM V6.1 19
 - ESM access control for real devices 18
 - ESM access control support for SET SECUSER and SET OBSERVER 19
 - extended address DASD volumes, support for 12
 - FICON Express8S 12
 - FIPS support 18
 - FTP enhancements 17
 - guest page reordering process, disabling 17
 - hard limiting of scheduled system resource, alternate method for 18
 - HyperSwap improvements 12
 - IBM z/VM Single System Image Feature 20
 - installation changes 20
 - installation process improvements 21
 - intraensemble data network (IEDN), support for 16
 - intranode management network (INMN), support for 16
 - IPv6 support 16
 - Language Environment upgrade 20
 - LDAP server upgrade 17
 - live guest relocation 10
 - memory constraint relief 17
 - MROUTE upgrade 16
 - multiple access ports per guest 17
 - network virtualization 16
 - OSA address table support 17
 - OSA-Express4S 12
 - OSX and OSM CHPIDs 11
 - packaging changes 20
 - page release serialization, improved 18
 - Performance Toolkit, enhanced 16
 - Program Management Binder upgrade 20
 - RACF Security Server enhancements 19
 - RACF support for protected user IDs 19
 - scalability improvements 17
 - SCSIDISC utility 15

- new or changed in z/VM V6.2 (*continued*)
 - security enhancements 18
 - serviceability changes 20
 - shipping zMAP agents discontinued 22
 - SSL server upgrade 18
 - statement of direction 22
 - systems management APIs, enhanced 15
 - systems management improvements 13
 - technology exploitation 11
 - Unified Resource Manager, support for 14
 - VMSES/E enhancements 21
 - withdrawal of support 21
 - XIV Storage Systems, direct attachment of 13
 - XL C/C++ Compiler for z/VM, V1.3, support for 20
 - XRC time stamping 12
 - z/Architecture CMS (z/CMS) 20
 - z/VM single system image (SSI) clusters 9
 - z114 support 11
 - z196 support 11
 - zEnterprise System, support for 14

O

- object code only (OCO) 38
- online books
 - DVD 72
 - program requirements 35
- Open Systems Adapter Support Facility (OSA/SF)
 - description 63
 - hardware requirements 27
 - program requirements 33
- OpenExtensions
 - books, abstracts of 81
 - byte file system (BFS) 57
 - overview 57
 - program requirements 32
- operating systems supported as guests of z/VM 95
- Operations Manager for z/VM 5
- ordering information 37
- OS/MVS simulation 56
- OSA (Open Systems Adapter) support 51
- OSA-Express 63
- OSA-Express2 63
- OSA/SF (Open Systems Adapter Support Facility)
 - description 63
 - hardware requirements 27
 - program requirements 33

P

- packaging, product
 - changes 20
 - distribution media 37
 - System Image contents 37
- Parallel Access Volumes (PAV) support 52
- Peer-to-Peer Remote Copy (PPRC) support 52
- Performance Toolkit for VM
 - book, abstract of 87
 - description 67
 - hardware requirements 27
 - program requirements 33

- POSIX support 57
- PR/SM (Processor Resource/Systems Manager) 91
- prerequisite licensed programs 30
- printers supported by z/VM 103
- problems, diagnosing 62
- Program Management Binder for CMS 58
- program requirements
 - base facilities and optional features 31
 - CSE 34
 - DFSMS/VM 31
 - DirMaint 32
 - HCD and HCM for z/VM 32
 - installing z/VM from DVD 26
 - IOCP 32
 - Language Environment 32
 - Linux 35
 - online books 35
 - OpenExtensions 32
 - OSA/SF 33
 - other products 35
 - Performance Toolkit for VM 33
 - POSIX 35
 - prerequisite licensed programs 30
 - RACF Security Server for z/VM 33
 - RSCS Networking for z/VM 34
 - SSI clusters 34
 - tasks that require the High Level Assembler 30
 - TCP/IP for z/VM 34
- Programmable Operator 5
- programming interfaces 6
- programming languages and environments 6
- programs supported on z/VM 36
- publications, z/VM
 - abstracts
 - base library - application programming 79
 - base library - customization and tailoring 77
 - base library - installation, migration, and service 73
 - base library - operation and use 77
 - base library - overview 73
 - base library - planning and administration 74
 - base library - system diagnosis 84
 - DFSMS/VM 86
 - DirMaint 86
 - OpenExtensions 81
 - Performance Toolkit for VM 87
 - RACF Security Server for z/VM 87
 - RSCS 88
 - basic 72
 - media 71
 - online
 - program requirements 35
 - sources 71
 - supplied with z/VM 72

Q

- QDIO (queued direct I/O) architecture support 52
- queued direct I/O (QDIO) architecture support 52

R

- RACF Security Server for z/VM (RACF Security Server for z/VM) 42
 - books, abstract of 87
 - description 68
 - program requirements 33
- real storage
 - minimum and maximum 26
- Remote Spooling Communications Subsystem (RSCS) Networking for z/VM (RSCS)
 - books, abstract of 88
 - description 69
 - program requirements 34
- requirements
 - customer responsibilities for data security 40
 - hardware
 - CSE 27
 - installing z/VM 25
 - OSA/SF 27
 - Performance Toolkit 27
 - server 25
 - SSI clusters 27
 - storage 26
 - TCP/IP for z/VM 28
 - workstation 26
 - installing z/VM 25
 - program
 - base facilities and optional features 31
 - CSE 34
 - DFSMS/VM 31
 - DirMaint 32
 - HCD and HCM for z/VM 32
 - installing z/VM from DVD 26
 - IOCP 32
 - Language Environment 32
 - Linux 35
 - online books 35
 - OpenExtensions 32
 - OSA/SF 33
 - other products 35
 - Performance Toolkit for VM 33
 - POSIX 35
 - prerequisite licensed programs 30
 - RACF Security Server for z/VM 33
 - RSCS Networking for z/VM 34
 - SSI clusters 34
 - tasks that require the High Level Assembler 30
 - TCP/IP for z/VM 34
 - server 25
 - storage requirements 26
 - virtual machine size, minimum 26
- restrictions
 - DirMaint 36
 - reconfiguration of partitionable processors 35
 - virtual machine 35
- Reusable Server Kernel 58
- REXX language 64
- REXX/VM 64
- RSA-Assist support 48

RSCS (Remote Spooling Communications Subsystem (RSCS) Networking for z/VM)
books, abstract of 88
description 69
program requirements 34

S

scalability improvements 17
SCSI disks, installing z/VM on 26
Secure Sockets Layer (SSL) support 42
security
DirMaint 42
enhancements 18
file pool 41
general 40
TCP/IP for z/VM 42
Server Time Protocol (STP) 53
servers
requirements 25
supported by z/VM 91
SFS (Shared File System) 55
single system image (SSI) clusters
description 66
hardware requirements 27
program requirements 34
SNA (Systems Network Architecture)
AVS 62
GCS 62
ISFC 47
softcopy books
DVD 72
program requirements 35
software requirements
base facilities and optional features 31
CSE 34
DFSMS/VM 31
DirMaint 32
HCD and HCM for z/VM 32
installing z/VM from DVD 26
IOCP 32
Language Environment 32
Linux 35
online books 35
OpenExtensions 32
OSA/SF 33
other products 35
Performance Toolkit for VM 33
POSIX 35
prerequisite licensed programs 30
RACF Security Server for z/VM 33
RSCS Networking for z/VM 34
SSI clusters 34
tasks that require the High Level Assembler 30
TCP/IP for z/VM 34
source code availability 38
spanned HiperSockets channels 50
specialty processors, support for 53
SSI (single system image) clusters
description 66
hardware requirements 27

SSI (single system image) clusters (*continued*)
program requirements 34
SSL support 42
statement of direction 22
storage
management using DFSMS/VM 65
minimum and maximum 26
storage controls supported by z/VM 102
storage reconfiguration 48
supported device, definition of 28
switches supported by z/VM 105
System Delivery Offering (SDO) 37
system integrity
customer responsibilities for data security 40
DirMaint 40
z/VM, general statement for 39
System z Application Assist Processor (zAAP), support for 53
System z Integrated Information Processor (zIIP), support for 53
system-managed storage 65
System/370 architecture compatibility 45
systems management API 56
systems management improvements 13
Systems Network Architecture (SNA)
AVS 62
GCS 62
ISFC 47

T

tape control units supported by z/VM 103
tape libraries supported by z/VM 102
tape units supported by z/VM 102
TCP/IP for z/VM
hardware requirements 28
overview 59
program requirements 34
security 42
technology exploitation, new 11
terminals supported by z/VM 104
TSAF (Transparent Services Access Facility) 64

U

unsupported device, definition of 29

V

V=V virtual machine 46
virtual machine
architectures 45
restrictions 35
type 46
Virtual Machine Resource Manager (VMRM) 58
virtual network adapters 46
virtual switch 47
virtualization technology, extensions to 9
VM Dump Tool 47
VMRM (Virtual Machine Resource Manager) 58
VMSES/E 64

VTAM 62

X

XA virtual machine 45

XC virtual machine 45, 46

XEDIT 59

Z

z/Architecture

ESA virtual machine 46

z/Architecture CMS (z/CMS) 55

z/OS, guest support for 95

z/TPF, guest support for 95

z/VM Information Center 71

z/VM Internet Library 71

z/VM V5.4

devices supported 97

guest operating systems supported 95

servers supported 91

z/VM V6.1

devices supported 97

guest operating systems supported 95

servers supported 91

z/VM V6.2

benefits 1

devices supported 97

distribution media 37

guest operating systems supported 95

hardware requirements

CSE 27

installation 25

OSA/SF 27

Performance Toolkit 27

server 25

SSI clusters 27

storage 26

TCP/IP for z/VM 28

workstation 26

multicultural support 37

new functions, enhancements, and changes 9

ordering information 37

program requirements

base facilities and optional features 31

CSE 34

DFSMS/VM 31

DirMaint 32

HCD and HCM for z/VM 32

installing from DVD 26

IOCP 32

Language Environment 32

online books 35

OpenExtensions 32

OSA/SF 33

Performance Toolkit for VM 33

POSIX 35

prerequisite licensed programs 30

RACF Security Server for z/VM 33

RSCS Networking for z/VM 34

SSI clusters 34

z/VM V6.2 (*continued*)

program requirements (*continued*)

tasks that require the High Level Assembler 30

TCP/IP for z/VM 34

programs supported 36

servers supported 91

source code availability 38

System Image contents 37

z/VM virtual switch 47

z/VSE, guest support for 95

zAAP (System z Application Assist Processor), support for 53

zIIP (System z Integrated Information Processor), support for 53



Product Number: 5741-A07

Printed in USA

GC24-6193-02

