

z/VM
built on IBM Virtualization Technology



General Information

version 5 release 2

z/VM
built on IBM Virtualization Technology



General Information

version 5 release 2

Note:

Before using this information and the product it supports, read the information in “Notices” on page 109.

Fourth Edition (July 2005)

This edition applies to version 5, release 2, modification 0 of IBM z/VM (product number 5741-A05) and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces GC24-6095-02.

© **Copyright International Business Machines Corporation 1990, 2005. All rights reserved.**

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

Contents

About This Book	vii
Who Should Read This Book	vii
Where to Find More Information	vii
How to Send Your Comments to IBM	vii
 Chapter 1. Introducing z/VM	1
z/VM Virtualization Technology Provides Guest Support	1
z/VM Supports Application Development and Deployment	2
z/VM Provides Proven System Integrity, Security, and Reliability	3
z/VM Is Accessible by People with Disabilities	4
 Chapter 2. How z/VM Can Help You	5
 Chapter 3. What Is New or Changed in z/VM V5.2	9
Virtualization Technology and Guest Support	9
Enhanced Scalability Support for Storage-Constrained Environments	9
Improved Performance of SCSI Disk I/O	10
Support for Open Systems Adapter for Network Control Program	10
Enhanced Performance Assists for z/VM Guests	11
Enhanced Management Functions for Linux and Other Virtual Images	12
SSL Server Support for Red Hat and Additional SUSE LINUX Distributions	12
Network Virtualization and Security	13
Improved Problem Determination for Guest LANs and Virtual Switches	13
Enhanced Dynamic Routing Capabilities with a New MPRoute Server	13
Technology Exploitation	14
Support for IBM System z9 109	14
Support for IBM TotalStorage DS6000 and DS8000 Series	16
Support for IBM TotalStorage DS4000 Midrange Disk Systems	16
Systems Management	17
Simplified User Administration with the Coordination of DirMaint and RACF Changes	17
Improved Directory Management Performance	18
Application Support	18
Upgraded CMS Binder	18
Installation, Migration, and Service	18
Simplified Migration of Preinstalled Facilities, Features, and Products	18
Further Automation of the Local Modification Process	19
Withdrawal of Functions or Support	19
Vector Facility Support Removed	19
System Administration Facility Removed	19
Support for Server-Requester Programming Interface Discontinued	19
Device Support Withdrawn	19
Packaging Modifications	19
Withdrawal of CD-ROM Distribution Medium	20
Withdrawal of Tivoli Storage Manager for z/VM	20
Statement of Direction	20
 Chapter 4. Hardware and Software Requirements and Other Technical Information	21
Hardware and Storage Requirements	21
Server Requirements	21
Installation Requirements	21
Storage Requirements	22

CSE Hardware Requirements	22
OSA/SF Hardware Requirements	22
TCP/IP Hardware Requirements	23
VM Guest RSA-Assist Support for Linux Hardware Requirements	23
Device Support	24
Supported Devices	24
Unsupported Devices	24
Software (Program) Requirements	24
Prerequisite Licensed Programs	25
High Level Assembler	25
Additional Base Facilities and Optional Features	26
Shared-DASD Complex and CSE Cluster Program Requirements	26
DFSMS/VM Program Requirements	26
DirMaint Program Requirements	27
Language Environment Program Requirements	27
OpenExtensions Program Requirements	27
OSA/SF Program Requirements	28
RACF for z/VM Program Requirements	28
TCP/IP for z/VM Program Requirements	28
Online Books Program Requirements	29
Linux Program Requirements	30
Program Requirements for Other IBM Products	30
Limitations and Restrictions	30
General Restrictions for z/VM Virtual Machines	30
DirMaint Restrictions	31
Operating Systems Supported as Guests	31
Other Programs Supported on z/VM	31
National Language Support	32
Packaging and Ordering Information	32
Object Code Only and Limited Source Availability	33
Integrity and Security	34
Data Integrity for Guests	34
System Integrity Statement for z/VM	34
Security, Auditability, and Control	35
Chapter 5. z/VM Base Product	39
Control Program	39
Virtual Machine Modes	39
Architecture Compatibility	40
z/Architecture Support	40
ESA/XC Support	40
Cross System Extensions	40
Guest LANs and Virtual Switches	40
Inter-System Facility for Communications	41
VM Dump Tool	41
Support for Hardware Facilities, Features, and Architectures	41
Conversational Monitor System	45
Shared File System	46
CMS Application Programming	46
CMS Pipelines	46
CMS Application Multitasking	47
OpenExtensions	47
Program Management Binder for CMS	48
Systems Management APIs	48
Reusable Server Kernel	48
Virtual Machine Resource Manager	48

XEDIT	49
HELP Facility	49
TCP/IP for z/VM	49
Link Protocols	50
Network Protocols	50
Transport Protocols	51
Applications and Protocols	51
APPC/VM VTAM Support	52
Dump Viewing Facility	52
Group Control System	53
Hardware Configuration Definition and Hardware Configuration Manager for z/VM	53
Language Environment	54
Open Systems Adapter Support Facility	54
REXX/VM	54
Transparent Services Access Facility	55
Virtual Machine Serviceability Enhancements Staged/Extended	55
Chapter 6. z/VM Optional Features	57
Data Facility Storage Management Subsystem for VM	57
Directory Maintenance Facility	57
Performance Toolkit for VM	58
Resource Access Control Facility for z/VM	59
Chapter 7. z/VM Library Guide	61
Where to Get z/VM Documentation	61
IBM Publications Center	61
z/VM Internet Library	61
z/VM Library on Disk	61
Basic (Automatically Shipped) z/VM Documentation	62
Books in the z/VM Base Library	62
System Overview	62
Installation, Migration, and Service	63
Planning and Administration	64
Customization and Tailoring	66
Operation	67
Application Programming	67
End Use	71
System Diagnosis	73
Books for z/VM Optional Features	74
DFSMS/VM	74
DirMaint	75
Performance Toolkit for VM	75
RACF for z/VM	76
Book Editions and Media	77
z/VM Base Library	78
z/VM Optional Features	80
Appendix A. IBM Servers Supported by z/VM	83
Appendix B. IBM Operating Systems Supported as Guests of z/VM	87
Appendix C. IBM Devices Supported by z/VM	91
Direct Access Storage Devices (DASD)	92
DASD Control Units and Storage Controls	99
Tape Units and Tape Libraries	100

Tape Control Units	101
Printers	101
Card Readers and Card Punches	102
Terminals/Displays/Consoles	102
Display Printers	102
Display Control Units	103
Communications Controllers	103
Switches.	103
Miscellaneous Devices and Server Features	104
Notices	109
Trademarks.	111
Index	113

About This Book

This book provides general information about IBM® z/VM®, version 5 release 2:

- Introduction to z/VM
- Cost/performance benefits
- What is new or changed in z/VM V5.2
- Hardware and software requirements
- Base facilities and optional features
- Product documentation
- IBM servers, guest operating systems, and devices supported by current z/VM releases

Who Should Read This Book

This book is intended for anyone who wants a general overview of z/VM V5.2. It also provides information for those who need to evaluate the capabilities of z/VM V5.2 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 books listed in the z/VM library guide, which begins on page 61. The library guide includes an abstract of each book.

Note: This edition has been published for the announcement of z/VM V5.2. Sources of additional information identified in this book might not be available until the general availability of z/VM V5.2.

Links to Other Online Books

If you are viewing the Adobe Portable Document Format (PDF) version of this book, it may contain links to other books. A link to another book is based on the name of the requested PDF file. The name of the PDF file for an IBM book is unique and identifies both the book and the edition. The book links provided in this book are for the editions (PDF names) that were current when the PDF file for this book was generated. However, newer editions of some books (with different PDF names) may exist. A link from this book to another book works only when a PDF file with the requested name resides in the same directory as this book.

How to Send Your Comments to IBM

IBM welcomes your comments. You can use any of the following methods:

- Complete and mail the Readers' Comments form (if one is provided at the back of this book) or send your comments to the following address:

IBM Corporation
Department 55JA, Mail Station P384
2455 South Road
Poughkeepsie, New York 12601-5400
U.S.A.

- Send your comments by FAX:
 - United States and Canada: 1-845-432-9405
 - Other Countries: +1 845 432 9405
- Send your comments by electronic mail to one of the following addresses:
 - Internet: mhvrcfs@us.ibm.com
 - IBMLink™ (US customers only): IBMUSM10(MHVRCFS)

Be sure to include the following in your comment or note:

- Title and complete publication number of the book
- Page number, section title, or topic you are commenting on

If you would like a reply, be sure to also include your name, postal or email address, telephone number, or FAX number.

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

Chapter 1. Introducing z/VM

z/VM is an operating system for the IBM System z9™ and @server® zSeries® platforms, or equivalent servers that support IBM z/Architecture™. z/VM provides a highly flexible test and production environment for enterprises deploying the latest On Demand Business solutions. The z/VM implementation of IBM virtualization technology provides the capability to run other full-function operating systems, such as Linux® on System z9 and zSeries and z/OS®, under z/VM as “guest” systems. z/VM supports z/Architecture (64-bit) guests and IBM Enterprise Systems Architecture/390® (31-bit) guests.

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

- Control Program (CP)
- Conversational Monitor System (CMS)
- Transmission Control Protocol/Internet Protocol (TCP/IP) for z/VM
- Advanced Program-to-Program Communication/Virtual Machine (APPC/VM) Virtual Telecommunications Access Method (VTAM®) Support (AVS)
- Dump Viewing Facility
- Group Control System (GCS)
- Hardware Configuration Definition (HCD) and Hardware Configuration Manager (HCM) for z/VM
- Language Environment®
- Open Systems Adapter Support Facility (OSA/SF)
- Restructured Extended Executor/Virtual Machine (REXX/VM)
- Transparent Services Access Facility (TSAF)
- Virtual Machine Serviceability Enhancements Staged/Extended (VMSES/E)

For more information, see Chapter 5, “z/VM Base Product,” on page 39.

z/VM also offers the following optional features:

- Data Facility Storage Management Subsystem for VM (DFSMS/VM®)
- Directory Maintenance Facility (DirMaint™)
- Performance Toolkit for VM™
- Resource Access Control Facility (RACF®) for z/VM

For more information, see Chapter 6, “z/VM Optional Features,” on page 57.

z/VM Virtualization Technology Provides Guest Support

z/VM presents a unique approach to computer operating systems. It provides each user with an individual working environment known as a *virtual machine*. The virtual machine simulates the existence of a dedicated real machine, including server functions, storage, and input/output (I/O) resources.

Operating systems and application programs can run in virtual machines as guests. For example, you can run multiple Linux and z/OS images on the same z/VM system that is also supporting various applications and end users. 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 and, at the same time, process production work.** 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 may 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.
- **Test a new operating system.** The existing operating system can be used to process production work concurrently with the generation and testing of a new operating system. Experience with the new system can be obtained before it is used on a production basis, 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 guest operating systems 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 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 APIs and other facilities are available to CMS applications, such as:

- 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
- Systems management APIs
- Reusable Server Kernel
- Assembler macros and functions
- OS/MVS simulation
- DOS/VSE support

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 may 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 Provides Proven System Integrity, Security, and Reliability

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

- Integrity and security:

Introduction

- z/VM supports guest use of the cryptographic facilities provided by supported IBM servers.
- IBM will correct any integrity exposures introduced by unauthorized programs into the system.
- Kerberos authentication and Secure Sockets Layer (SSL) support are provided through TCP/IP for z/VM.
- Integrated access control and authentication services may be augmented with the addition of the RACF for z/VM feature or other external security managers.
- 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 SystemView[®] Host Management Facilities/VM can be added. Both the Programmable Operator and Host Management Facilities/VM can interact with NetView[®] on z/VM, which in turn can interact with NetView on z/OS.
 - 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.

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

Chapter 2. How z/VM Can Help You

z/VM version 5 extends the strength of the IBM mainframe environment as an ideal platform for integrating applications and data and consolidating select UNIX®, Microsoft® Windows®, and Linux workloads on a single physical server. This can allow you to run tens to hundreds of Linux servers with the potential to support new workloads, achieve productivity gains, and lower IT costs. z/VM supports Linux on the mainframe and enables Linux servers to share hardware resources and use internal high-speed communication within the System z9 and zSeries environments.

z/VM version 5 allows you to put the power of System z9 and zSeries server partitioning and z/VM virtualization technology to work for you and help realize the benefits of workload isolation and resource sharing, including:

- Reliability, availability, and serviceability of IBM mainframes
- Flexibility to create as many as 60 logical partitions (LPARs) with the System z9 109 server
- Ability to virtually partition each LPAR into hundreds or more virtual machines
- Ability to virtualize processor, communication, memory, storage, I/O, and networking resources
- Help with maximizing resources to achieve high system utilization
- Advanced dynamic resource allocation
- High-speed communications among LPARs and guests with HiperSockets™
- Advanced systems management, administration, and accounting tools

z/VM version 5 has a pricing model that:

- Helps reduce the expense of incremental growth; that is, the price per-engine for z/VM version 5 becomes less as more engines are added (purchased).
- Can make it more affordable for you to add z/VM to an existing standard engine environment.

The success of Linux on System z9 and zSeries is largely attributed to the business value that Linux-based solutions derive from IBM mainframe virtualization technology as provided by z/VM and the System z9 and zSeries platforms. z/VM technology offers you the ability to host a large number of Linux servers on a single mainframe while also providing an operational environment that is well-suited for on demand computing: highly flexible, adaptable, and efficient.

z/VM version 5 offers new levels of price/performance, functional capabilities, and hardware exploitation that are expected to increase the attractiveness of deploying Linux solutions on the mainframe. You can add capacity to IBM mainframe servers for hosting Linux-on-z/VM workloads by configuring the servers with Integrated Facility for Linux (IFL) engines.

The z/VM version 5 pricing model can make it more feasible for you to add z/VM virtualization technology to your standard-engine environment (compared with the pricing models for z/VM version 3 and version 4). This helps enable you to use z/VM to host workloads that cannot run on IFL engines, such as z/OS, VSE/ESA™, z/VSE™, TPF, or z/TPF™ (when available), to consolidate existing VM/ESA® or z/VM version 3 and version 4 workloads onto a single, larger System z9 or zSeries server.

Benefits

As the successor product to VM/ESA, z/VM version 3, and z/VM version 4, z/VM version 5 provides additional support and exploitation opportunities for the thousands of users who have built enterprise-wide automation and infrastructure enhancements on the VM platform in support of their applications, database systems, and On Demand Business solutions.

z/VM version 5 is intended to address the following situations:

- **Running more Linux server images on a single physical server.**
Considerably more images than are currently supported by the LPAR mode of operation (up to 60 on the System z9 109, up to 30 on the zSeries 990 and 890, up to 15 on the zSeries 900 and 800) can be achieved with z/VM guest support. These Linux server images can be deployed on standard engines (CPs) or IFL engines with z/VM version 5. Running multiple Linux images on an IFL-configured z/VM system may not increase the IBM software charges of your existing System z9 or zSeries environment. You can add z/VM version 5 running on IFL engines to your existing z/OS, z/OS.e, z/VM, TPF, z/TPF (when available), VSE/ESA, z/VSE, or Linux environments without necessarily increasing software costs on the standard engines (CPs).
- **Moving select Linux, Windows, and UNIX workloads to a single physical server while maintaining distinct server images and current LAN topology.**
This ability can help reduce systems management complexity. Because the number of real hardware servers and associated physical LANs is reduced, cost savings may be realized by managing large server farms deployed on virtual servers instead of using multiple hardware servers. Deploying Linux workloads on z/VM version 5 is particularly attractive if they interact with applications or data located on the same server.
- **Enhancing virtual networking.** z/VM virtual switch support provides external connectivity for guest LANs through an OSA-Express adapter without requiring a z/VM or Linux router virtual machine.
- **Consolidating operating systems on the System z9 or zSeries platform.**
z/VM V5.2 provides considerable relief for real memory-constrained VM environments. The z/VM Control Program provides added support for large real memory configurations, improving the efficiency of managing storage located above the 2 GB address line. This may enable a single z/VM instance to support a significantly larger number of guest operating systems, or guests with greater aggregate memory, than prior VM releases. This constraint relief is provided for both z/Architecture and ESA/390 guest operating systems such as Linux on System z9 and zSeries, TPF, z/TPF (when available), z/OS, VSE/ESA, or z/VSE. If you have experienced real memory constraints below the 2 GB address line with z/VM V5.1 and earlier levels of z/VM and VM/ESA, you are likely to experience improved system throughput with z/VM V5.2.
- **Migrating from VM/ESA to z/VM version 5.** This helps enable:
 - More memory to cache CMS minidisks by exploiting memory above 2 GB for minidisk cache
 - Connectivity (TCP/IP) enhancements and additional device support
 - Added security with SSL-protected TCP/IP network traffic, such as Telnet sessions and Web transactions
- **Migrating to the new version or release of a guest operating system using z/VM version 5.** This can provide added flexibility for migration, production, and testing. For example, z/VM can help you migrate from OS/390® to z/OS or z/OS.e.

- **Enhancing guest Parallel Sysplex® support in z/VM with the exploitation of z/Architecture.** This can enable addressability of larger amounts of real and virtual memory, allowing the development and testing of 64-bit Parallel Sysplex applications in a guest environment.

Benefits

Chapter 3. What Is New or Changed in z/VM V5.2

This chapter provides an overview of the new functions, enhancements, and other changes included in z/VM V5.2.

The product changes are described under the following topics:

- “Virtualization Technology and Guest Support”
- “Network Virtualization and Security” on page 13
- “Technology Exploitation” on page 14
- “Systems Management” on page 17
- “Application Support” on page 18
- “Installation, Migration, and Service” on page 18
- “Withdrawal of Functions or Support” on page 19
- “Packaging Modifications” on page 19

This chapter also includes a “Statement of Direction” on page 20.

See the *z/VM: Migration Guide* for information about the changes to external interfaces to support the z/VM V5.2 enhancements, and for information about the enhancements included in previous z/VM releases.

Virtualization Technology and Guest Support

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

Enhanced Scalability Support for Storage-Constrained Environments

Exploitation of the 64-bit addressing capability of z/Architecture is significantly improved with z/VM V5.2. z/VM V3.1 provided partial exploitation of z/Architecture to help enable the execution of 64-bit-capable operating systems and applications as guests of z/VM when running on a System z9 109 (z9-109) or zSeries server, but z/VM was still constrained by the need for many types of data to reside in real storage below 2 GB. This exploitation provides some throughput improvements when running operating systems and applications on z9-109 or zSeries servers with large real storage.

z/VM V5.2 provides major improvements when operating on z9-109 and zSeries servers with large storage configurations. Scalability is improved with the z/VM Control Program (CP) now using storage above 2 GB for a much broader set of operations. Previously, guest pages had to be moved below 2 GB for many reasons. For example, guest I/O buffers for both standard I/O and Queued Direct Input/Output (QDIO), used for networking and Fibre Channel Protocol (FCP) devices, were moved below 2 GB when an I/O operation was initiated. For QDIO, CP queue-control structures had to reside below 2 GB, as did guest queue-control structures, which were moved below 2 GB when needed. Now I/O can be done using buffers anywhere in real storage, and QDIO structures can reside above 2 GB, as can most CP control blocks. These improvements offer constraint relief for large-real-storage virtual-server environments that are storage-intensive.

The CP functions that produce hard and soft abend dumps and SNAPDUMPs have been enhanced to dump storage above 2 GB if such storage exists. Likewise, a

New or Changed

standalone dump of a z/VM system or a VMDUMP of a virtual machine in z/Architecture mode can now include storage above 2 GB.

In addition to the CP enhancements, other functions within z/VM have changed to take advantage of the removal of the 2 GB address-line constraint, including the following:

- **TCP/IP for z/VM:** Enhanced QDIO device driver to use I/O buffers above 2 GB when possible. This will help to increase system reliability by reducing contention for storage below 2 GB.
- **Block I/O (Diagnose X'250'):** Extended to allow a virtual machine to specify parameter addresses and I/O buffers with addresses above 2 GB. This will benefit operating systems that use DIAGNOSE code X'250' in a 64-bit environment. IBM is working with its Linux distribution partners to exploit this function in future Linux on System z9 and zSeries distributions or service updates.
- **VM Dump Tool:** Enhanced to support the new dump formats.

Improved Performance of SCSI Disk I/O

z/VM V5.2 provides performance improvements for Small Computer System Interface (SCSI) disk I/O, including:

- QDIO efficiency improvements
The z/VM SCSI I/O subsystem (SCSI-driver stack) will now support and exploit the QDIO architecture's ability to chain address lists, allowing more data to be moved with a single I/O request. Chaining the address lists helps to reduce fragmentation of large channel programs, decreases CP overhead, and significantly lowers the number of I/O operations driven through the SCSI stack, thereby improving the performance of SCSI I/O.
- Paging and spooling optimization
CP's paging and spooling operations on SCSI disks currently use FBA emulation to transform standard I/O requests (that is, channel programs) into SCSI commands. With this new optimization, CP's paging and spooling function will build and send SCSI commands directly to the z/VM SCSI stack, thereby improving efficiency and performance.
- Improved FBA emulation efficiency
 - CP's FBA emulation function will now read CMS I/O buffers more efficiently, reducing the number of I/O requests to the z/VM SCSI stack. This enhancement will help to significantly improve SCSI I/O performance for CMS guests.
 - FBA emulation will now handle LOCATE CCWs more efficiently, reducing the number of I/O requests to the z/VM SCSI stack. This enhancement helps to significantly improve SCSI I/O performance for Linux guests, including the length of time to IPL a Linux guest.
 - FBA emulation will now handle the FBA padding function (initiated using a CCW data count less than the associated LOCATE data count) more efficiently, reducing the number of I/O requests to the z/VM SCSI stack. This enhancement helps to significantly improve SCSI I/O performance for format functions.

Support for Open Systems Adapter for Network Control Program

The OSA-Express2 Gigabit Ethernet and 1000BASE-T Ethernet features now have the capability to provide channel connectivity from System z9 operating systems to the IBM Communication Controller for Linux (CCL) with the introduction of the Open

Systems Adapter for the Network Control Program (OSA-Express2 OSN) supporting the Channel Data Link Control (CDLC) protocol.

OSA-Express2 OSN can help eliminate the requirement to have any form of external medium (and all related hardware) for communications between the host operating system and the CCL image. Traffic between the two images (operating system and CCL) is no longer required to flow on an external Local Area Network (LAN) or ESCON® channel; traffic can flow LPAR-to-LPAR. OSA-Express2 OSN provides an efficient method of communication, and is designed to create a secure and seamless integration of the host operating system and CCL.

If you continue to need SNA solutions that require NCP functions, you can now consider CCL as a migration strategy to replace your IBM Communications Controllers (374x). The CDLC connectivity option enables TPF or z/TPF (when available) environments to exploit CCL.

z/VM V5.2 enables VTAM on z/VM and other guests to access OSA-Express CDLC channels. The Open Systems Adapter for OSN is also supported on z/VM V5.1.

OSA-Express2 OSN support is exclusive to the z9-109, OSA-Express2 Gigabit Ethernet SX, Gigabit Ethernet LX, and 1000BASE-T Ethernet features, and requires the port to be configured as CHPID type OSN, which can be done on a port-by-port basis. OSA-Express2 OSN is exclusively for internal communication, LPAR-to-LPAR.

For additional information on the OSA-Express2 OSN, refer to the System z9 109 announcement dated July 26, 2005.

IBM is working with its Linux distribution partners to exploit this function in future Linux on System z9 and zSeries distributions or service updates

Enhanced Performance Assists for z/VM Guests

z/VM V4.4 added an important virtualization technology designed to improve the performance of guest operating systems that use the QDIO architecture (such as Linux on System z9 and zSeries). Virtualization technology allows adapter interruptions to be passed directly to z/VM guests for HiperSockets, FCP, and OSA on the z9-109, z990, and z890 servers.

A complementary virtualization technology is being introduced for the z9-109, z990, and z890:

- QDIO Enhanced Buffer-State Management (QEBSM) - two new hardware instructions designed to help eliminate the overhead of hypervisor interception
- Host Page-Management Assist (HPMA) - an interface to the z/VM paging-storage management function designed to allow page frames to be assigned, locked, and unlocked without z/VM hypervisor assistance, primarily benefiting the QEBSM environment.

These new hardware assists 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.

These new virtualization technologies are available only to first-level guests of z/VM V5.2. That is, they are not available to guests of a z/VM system that is itself running as a guest of another z/VM system.

New or Changed

IBM is working with its Linux distribution partners to provide exploitation of QEBSM in future Linux on System z9 and zSeries distribution releases or service updates.

Your z990 and z890 servers will require microcode level (MCL) updates for these new assists. Support is integrated in the z9-109 licensed internal code (LIC). These performance assists are provided on the z9-109, z990, and z890 for the following channel path types: HiperSockets (CHPID type IQD), FICON® (CHPID type FCP), and all OSA features (CHPID type OSD). For additional information on these new hardware assists, refer to the System z9 109 announcement dated July 26, 2005.

Enhanced Management Functions for Linux and Other Virtual Images

z/VM virtual systems management APIs, first introduced in z/VM V4.4, are provided for platform provisioning applications, including programs developed by Independent Software Vendors (ISVs), for ease of use in creating and managing large numbers of Linux and other virtual images running on z/VM.

All of the new and enhanced APIs in z/VM V5.2 have been implemented using a new version of the z/VM systems management server, version 3. APIs provided in earlier releases of z/VM, implemented using V1 or V2 of the server, can also function with the V3 server.

z/VM V5.2 provides new APIs to support the following:

- Creating and updating of the LOADDEV directory statement for a virtual image, as well as responding to queries of a virtual image's LOADDEV settings in support of List-Directed-IPL
- Obtaining the time when a virtual image was activated
- Defining local tags by the directory manager
- Scanning (searching) the directory for a specified pattern

The following enhancements have been made to existing APIs in V5.2:

- Larger disk sizes (up to 2,147,483,640 blocks) when creating disks and defining disk space on volumes for virtual images
- Capability to create persistent virtual switch definitions
- New definition parameters for virtual switch APIs

For more information, see *z/VM: Systems Management Application Programming*.

SSL Server Support for Red Hat and Additional SUSE LINUX Distributions

The SSL server support that shipped with TCP/IP for z/VM V4.4 and V5.1 provided Red Hat Package Manager (RPM) packages for the Novell SUSE LINUX Enterprise Server SLES7 (Linux kernel 2.4.7) and SLES8 (Linux kernel 2.4.19) distributions. In z/VM V5.2, support will be provided for:

- SUSE SLES8 Service Pack 3 (31-bit)
- SUSE SLES9 Service Pack 2 (31-bit)
- SUSE SLES9 Service Pack 2 (64-bit)
- Red Hat Enterprise Linux AS V3 (31-bit)
- Red Hat Enterprise Linux AS V3 (64-bit)

Additionally, this upgraded SSL server includes:

- A variety of industry-standard encryption algorithms, including DES, triple-DES, RC2, and RC4, with keys up to 128 bits in length. Hashes are provided by SHA-1 and MD5.
- Certificate activation and removal without server restart.
This capability allows a new certificate to be added or an existing one to be deactivated while existing SSL-secured sessions are active, removing the need to shut down and restart the SSL server when certificates are added or deleted.
- Federal Information Processing Standard (FIPS 140-2) operational mode support.
The SSL server may be operated in FIPS mode, restricting connections to those that employ FIPS-approved cipher suites.

Network Virtualization and Security

This section describes enhancements to z/VM network virtualization and security.

Improved Problem Determination for Guest LANs and Virtual Switches

z/VM V5.2 virtualizes a LAN sniffer to capture network traffic on a z/VM guest LAN or virtual switch. This capability helps an administrator (or the owner of the guest virtual machine) to capture network data to resolve virtual networking problems. The use of this sniffer function is predicated on the setting of the proper authorization controls. As with other potentially sensitive functions, LAN sniffer authorization is under the control of the system administrator, not of general users or guests running on the system. Capabilities are provided to capture and process the data for both Linux and traditional environments:

- Native Linux tracing capability on a guest LAN or virtual switch

When a Linux guest is deployed, traffic can be traced, recorded, and analyzed by existing tools directly from the guest virtual machine. This Linux guest must be authorized to use this capability through CP commands. The authorized guest can then use CP commands or the Linux device driver (when available) to put the guest NIC in “Promiscuous Mode”. The collected data can be printed or written to a file using tools such as **tcpdump** (available with Linux on System z9 and zSeries) to determine the causes of the network problems.

IBM is working with its distribution partners to exploit this function in future Linux on System z9 and zSeries distributions or service updates.

- Native z/VM tracing capability on a guest LAN or virtual switch

LAN traffic can be traced, recorded, and analyzed using native z/VM facilities. This function is authorized only to users with Class C privileges. The CP TRSOURCE command has been enhanced to trace and record the data transmissions, which can be analyzed with a new IPFORMAT tool to determine the causes of network problems.

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

The optional RACF for z/VM feature for z/VM V5.2 provides the ability to control promiscuous mode authorizations for guest LANs and virtual switches.

Enhanced Dynamic Routing Capabilities with a New MPRoute Server

The multi-protocol dynamic routing server (MPRoute) implements Open Shortest Path First (OSPF) and Routing Information Protocol (RIP), providing a powerful alternative to static TCP/IP gateway definitions. When properly configured, a z/VM host running the MPRoute server can become an active OSPF or RIP (or both) network router, providing network access to z/VM virtual networks. However, the current TCP/IP for z/VM MPRoute server available with z/VM V5.1 does not support

Internet Protocol version 6 (IPv6) dynamic routing and cannot receive RIPv1 and RIPv2 on the same interface. A new MPRoute server has been adapted from z/OS V1.7 to operate on z/VM V5.2 and supports the following protocols:

- For IPv4, MPRoute implements the OSPF protocol described in RFC 1583 (OSPF version 2) and the RIP protocols described in RFC 1058 (RIP version 1) and in RFC 1723 (RIP version 2).
- For IPv6, MPRoute implements the IPv6 OSPF protocol described in RFC 2740 (OSPF for IPv6) and the IPv6 RIP protocol described in RFC 2080 (RIPng for IPv6).

The new MPRoute server removes the limit of four equal-cost paths, allowing the generation of up to 16 equal-cost routes to a destination, thus providing improved load-balancing support. The current TCP/IP stack can only handle up to four equal-cost paths to a destination, but the z/VM V5.2 TCP/IP stack can handle an unlimited number of equal-cost paths.

For IPv4, MPRoute supports the RIP version 1, RIP version 2, and OSPF routing protocols. You can send RIP version 1 or RIP version 2 (but not both at the same time) on a single interface. However, you can configure a RIP interface to receive both versions of RIP. For IPv6, MPRoute supports the IPv6 OSPF and IPv6 RIP protocols.

MPRoute is the routing server recommended for use with TCP/IP for z/VM. IBM plans to remove the Routed server from a future release of z/VM. z/VM V5.2 is planned to be the last release on which this server will be available.

By using MPRoute, greater efficiency may be achieved within an IP network and manual network routing table updates may be reduced or eliminated.

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

Technology Exploitation

This section describes z/VM support and exploitation of new hardware, hardware facilities, and architectures.

Note: Exploitation of hardware functions depends on hardware level and availability. For detailed information on hardware capabilities, requirements, and availability, refer to the appropriate hardware announcement materials.

Support for IBM System z9 109

z/VM V5.2 supports the new System z9 109 server, including the following new facilities and features.

Crypto Express2 Supported as an Accelerator Card

The Crypto Express2 feature, with two PCI-X adapters, can be defined as a Coprocessor or as an Accelerator.

- Crypto Express2 Coprocessor - for secure-key-encrypted transactions (default)
 - Designed to support highly secure cryptographic functions, use of secure-encrypted key values, and User Defined Extensions (UDX)
 - Designed for Federal Information Processing Standard (FIPS) 140-2 Level 4 certification
- Crypto Express2 Accelerator - for SSL acceleration
 - Designed to support clear-key RSA operations

- Offloads compute-intensive RSA public-key and private-key cryptographic operations employed in the SSL protocol

The configurable Crypto Express2 feature is exclusive to the z9-109 and is supported by z/VM V5.2, providing guest support for z/OS and Linux on System z9 and zSeries that includes:

- Dedicated-queue support for clear-key and secure-key cryptographic functions for z/OS guests
- Shared-queue and dedicated-queue support for clear-key cryptographic functions for Linux on System z9 and zSeries guests.

z/VM V5.1 already provides support for the Crypto Express2 Coprocessor. Support for the configurable Crypto Express2 Accelerator is also planned to be supported by z/VM V5.1 with the PTF for APAR VM63646 at the availability of the hardware feature. For additional information on the Crypto Express2 Accelerator feature, refer to the System z9 109 announcement dated July 26, 2005.

N_Port ID Virtualization

N_Port Identifier Virtualization (NPIV) for FCP channels (CHPID type FCP) is designed to allow the sharing of a single physical FCP channel among multiple operating-system images, whether in logical partitions (LPARs) or as z/VM guests in virtual machines.

z/VM exploits this capability to define multiple virtual FCP channels, each with its own unique Fibre Channel port name and Fibre Channel identifier (FC_ID). By assigning different virtual port names to different guests, the guests can use the virtual FCP channels as if they were using dedicated physical FCP channels. In particular, access controls based on the virtual port names may be applied in the fabric using standard mechanisms like zoning in the switches and Logical Unit Number (LUN) masking in the storage controllers.

NPIV offers improved FCP channel utilization and sharing among operating-system images, joining ESCON and native FICON in offering channel-sharing through virtualization. This may help to reduce your hardware requirements and may help to facilitate infrastructure simplification.

NPIV is the industry standard solution which supersedes the previously previewed FCP LUN Access Control in the zSeries announcement dated October 7, 2004.

The improved sharing and channel utilization of FCP channels among operating system images in LPARs or virtual machines allowed by NPIV is exclusive to the z9-109 and is applicable to all of the FICON features supported on the z9-109. For additional information on NPIV, refer to the System z9 109 announcement dated July 26, 2005.

z/VM V4.4, V5.1, and V5.2 enable guest use of NPIV when FCP subchannels are dedicated to a guest.

z/VM V5.1 and V5.2, which support the use of SCSI FCP devices, provide for CP use of NPIV; z/VM V5.1 requires the PTF for APAR VM63744. z/VM V5 cannot be installed from DVD to SCSI disks when NPIV is enabled, and a future enhancement is planned to provide this capability in z/VM V5.2.

IBM also intends to provide a future enhancement to z/VM V5.2 for NPIV such that guest operating systems and z/VM users can obtain virtual port names. This

New or Changed

planned enhancement is to the CP QUERY command for z/VM users and to z/Architecture virtualization for guest operating systems.

IBM is working with its distribution partners to provide this function in future Linux on System z9 and zSeries distributions or service updates.

Dynamic Addition or Deletion of a Logical Partition Name

The ability to add meaningful LPAR names to your configuration without a Power-On Reset was made available on the z890 and z990 servers on May 28, 2004. Prior to this support, extra LPARs could be defined by adding reserved names in the Input/Output Configuration Data Set (IOCDs), but predicting meaningful names in advance of the need for extra LPARs was difficult.

z/VM V5.2 is designed to exploit this hardware capability on z9-109, z890 and z990 servers by providing facilities to dynamically define and delete logical partitions using CP's Dynamic I/O command interface and z/VM HCD/HCM support.

For additional information on this hardware capability, refer to the z890 announcement or z990 announcement dated April 7, 2004.

Support for up to 60 LPARs

IBM is once again doubling the number of LPARs. You now have the ability to define up to 60 LPARs, 15 per Logical Channel Subsystem, which may provide you with even more flexibility in allocating hardware resources. As before, with the Processor Resource/Systems Manager™ (PR/SM™) and the Multiple Image Facility (MIF) you can share ESCON and FICON channels, ISC-3s, and OSA ports across LPARs.

Support of up to 60 LPARs is exclusive to the z9-109 and is supported by z/VM V4.4, V5.1, and later releases, satisfying the Statement of General Direction in the z/VM V4.4 announcement dated May 13, 2003.

Support for IBM TotalStorage DS6000 and DS8000 Series

z/VM V5.2 supports the TotalStorage® DS6000 and DS8000 series disk storage subsystems. The DS6000 series is designed to deliver enterprise-class storage capabilities in a space-efficient, modular design at a low price. The DS8000 series is designed to provide unmatched functionality, flexibility, and performance for enterprise disk storage systems at new levels of cost effectiveness. Also supported is the capability to define and operate FCP-attached SCSI disks with capacities near 1 TB (2,147,483,640 512-byte blocks) for CP volumes and 381 GB for CMS and GCS volumes.

The DS6000 and DS8000 series are also supported by z/VM V4.4 and z/VM V5.1 with the PTFs for APARs VM63535 (DS6000), VM63534 (DS8000), and VM63653 (DS6000 and DS8000). In addition, PTFs for DirMaint APAR VM63700 and DFSMS/VM APAR VM63664 are required to support the near 1 TB CP SCSI disks on z/VM 5.1.

For additional information on the DS6000, refer to the DS6000 announcement; for the DS8000, refer to the DS8000 announcement dated October 12, 2004.

Support for IBM TotalStorage DS4000 Midrange Disk Systems

z/VM V5.1 provided support for FCP-attached SCSI disk logical units (SCSI disks) when connected to a Fibre-Channel fabric via System z9 and zSeries FCP channels. This support included the native-SCSI interface for guest use, and an

FBA interface for both system and guest use. z/VM V5.2 SCSI support extends this to the lower-priced DS4000 (formerly FAStT) Disk Systems with the addition of the Redundant Disk Array Controller (RDAC) driver. The RDAC includes multipath attachment with failover support and is planned to be provided for the DS4300 (FAStT 600), DS4400 (FAStT 700), and DS4500 (FAStT 900). Multipath failover support is designed to continue accessing storage on attached devices even if an element in the path (for example, an adapter, cable, or switch) fails.

SCSI disks can be used directly by a guest operating system when an FCP subchannel is dedicated to a guest. Such a guest must provide its own SCSI device support. Linux on System z9 and zSeries is designed to be such a guest. SCSI disks can also be used as emulated 9336 Model 20 FBA devices.

z/VM CP and CMS rely almost exclusively on this emulated-FBA support for their SCSI use. Specifically, this use includes directory services, minidisks, and all other system functions and programming services that can use FBA disks. Guests that support FBA disks (such as CMS, GCS, and VSE) also can use SCSI disks through the emulated-FBA support, without requiring their own SCSI support.

z/VM's SCSI support allows a Linux server farm to be deployed on z/VM in a configuration that includes only SCSI disks. ECKD™ disks are no longer required. Installation of z/VM from DVD to a SCSI disk, IPL from a SCSI disk using the Stand-Alone Program Loader (SAPL), and VM system dumps to a SCSI disk are supported. DASD Dump/Restore (DDR) services using SCSI disks are supported when DDR is running under CMS.

Systems Management

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

Simplified User Administration with the Coordination of DirMaint and RACF Changes

DirMaint is a CMS application that helps manage your z/VM directory. Directory statements can be added, deleted, or altered using the DirMaint directory statement-like commands. DirMaint provides automated validation and extent allocation routines to reduce the chance of operator error.

RACF for z/VM can be used to enhance the security and integrity of your system by:

- Helping your installation implement its security policy
- Identifying and authenticating each user
- Controlling each user's access to sensitive data
- Logging and reporting events that are relevant to the system's security.

In z/VM V5.2, the directory management functions of DirMaint can be integrated with the security management functions of RACF. DirMaint can be configured to notify RACF whenever important changes are made to user definitions and the resources they own. This configuration reduces the administrative effort and skills needed to deploy and manage users and their resources when DirMaint and RACF are used together. By eliminating the need to manually define and manage z/VM resources in RACF, the possibility of incomplete or incorrect RACF configuration is reduced. Functions that are coordinated by this new DirMaint support include:

- User creation, deletion, and changes

New or Changed

- Password management
- POSIX segment management
- Access Control Interface (ACI) group management
- Profile creation and deletion for selected VM functions

For more information, see *z/VM: Directory Maintenance Facility Tailoring and Administration Guide*.

Improved Directory Management Performance

On z/VM systems with large user directories, changes made to the user directory using the DirMaint feature should be faster than on previous releases. This enhancement allows a change to the directory without requiring reprocessing of the entire directory, resulting in less time to make the change. Performance improvements depend on the type of directory changes being made and the size of the VM directory being changed. For performance expectations, refer to the z/VM V5.2 Performance Report, planned to be available by December 16, 2005.

Application Support

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.

Upgraded CMS Binder

The CMS binder has been upgraded to a new level based on the z/OS MVS™ Program Management Binder supplied with z/OS V1.6. For more information, see the *z/VM: Program Management Binder for CMS*.

Installation, Migration, and Service

This section describes changes to the z/VM installation, migration, and service processes.

Simplified Migration of Preinstalled Facilities, Features, and Products

This migration process allows the disks associated with the preinstalled additional facilities, optional features, and licensed products on your first-level z/VM V5.1 system to be made available to your second-level z/VM V5.2 system and then transfers the following types of files:

- Customized files
- Local modifications
- Service
- User-created files residing on selected disks

The preinstalled facilities, features, and products include:

- HCD and HCM for z/VM
- OSA/SF
- TCP/IP for z/VM
- DirMaint
- Performance Toolkit for VM
- RACF for z/VM
- EREP/VM

- ICKDSF
- RSCS

For more information, see the *z/VM: Guide for Automated Installation and Service*.

Further Automation of the Local Modification Process

z/VM V5.1 provided a new command, LOCALMOD, to make the procedure for creating and applying local modifications easier. z/VM V5.2 provides further automation of the LOCALMOD process by adding a capability to rework local modifications and provide support for local service. For more information, see the *z/VM: VMSES/E Introduction and Reference*.

Withdrawal of Functions or Support

This section identifies functions or support that have been removed from z/VM.

Vector Facility Support Removed

The Vector Facility is an S/390® hardware feature. z/VM version 5 does not support S/390 servers, so Vector Facility support and the CP virtual vector facility function have been removed from z/VM V5.2.

System Administration Facility Removed

The System Administration Facility tools that were added in z/VM V4.2 to help ease migration from the IBM S/390 Virtual Image Facility to z/VM V4.2 have been removed from z/VM V5.2. Some of the same functionality may be obtained through the z/VM Systems Management APIs, first introduced in z/VM V4.4. This satisfies the Statement of Direction in the z/VM V5.1 announcement dated April 7, 2004.

Support for Server-Requester Programming Interface Discontinued

Support for the Server-Requester Programming Interface (SRPI) has been discontinued with the availability of z/VM V5.2. This satisfies the Statement of Direction in the z/VM V5.1 announcement dated April 7, 2004.

Device Support Withdrawn

The following real I/O devices are not supported by z/VM V5.2:

- 3705 Communications Controller
- 3720 Communications Controller
- 3725 Communications Controller
- 8232 LAN Channel Station

For a list of the real devices that are supported by z/VM V5.2, see Appendix C, “IBM Devices Supported by z/VM,” on page 91.

In addition, the IBM 2741 Communications Terminal and the TWX Terminal Model 33/35 (TTY) are not supported as virtual consoles.

Packaging Modifications

This section describes changes to contents of the z/VM V5.2 product package and the z/VM V5.2 System Delivery Offering (SDO).

Withdrawal of CD-ROM Distribution Medium

z/VM V5.2 will not be available on CD-ROM. z/VM continues to be distributed on 3480/3490 and 3590 tape and on DVD.

Withdrawal of Tivoli Storage Manager for z/VM

The Tivoli® Storage Manager for VM product was removed from the z/VM V5.1 system image and has not been orderable with the z/VM SDO since September 24, 2004. With the announcement of the Tivoli Storage Manager V5.3 family of products in the Tivoli Storage Manager announcement dated December 14, 2004, support for z/VM was not migrated to Tivoli Storage Manager V5.3 nor is it supported.

When running in a Linux guest environment, IBM recommends Tivoli Storage Manager, V5.3 (5608-ISM), or Tivoli Storage Manager Extended Edition, V5.3 (5608-ISX), for key data protection activities of backup, archive, recovery, space management, and disaster planning.

Statement of Direction

- IBM intends to provide future enhancements to z/VM supporting the following System z9 109 functions:
 - System and guest exploitation of HyperSockets supporting the IPv6 protocol
 - Improved memory management between z/VM and Linux on System z9 and zSeries
 - Simplified networking administration and management of VLANs with support for GARP VLAN Registration Protocol (GVRP) using OSA-Express2
 - Capability to allow guest operating systems and z/VM users to query virtual port names when using N_Port ID Virtualization
- IBM intends to evaluate z/VM V5.2 with the RACF for z/VM optional feature for conformance to the Controlled Access Protection Profile (CAPP) and Labeled Security Protection Profile (LSPP) of the Common Criteria standard for IT security, ISO/IEC 15408, at Evaluation Assurance Level 4 (EAL4).
- IBM intends to provide IBM Director support for Linux on System z9 and zSeries, extending IBM virtualization technology leadership with the exploitation of system management and virtual server deployment functions based on the Common Information Model (CIM) standard. IBM Director is a key component of the IBM Virtualization Engine™.
- IBM intends to provide exploitation of the IBM TotalStorage Parallel Access Volume (PAV) feature for z/VM system data and guest data residing on VM minidisks in a future release of z/VM.
- IBM plans to remove the RouteD and BOOTP servers from a future release of z/VM. z/VM V5.2 is planned to be the last release in which these servers will be available.

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

Chapter 4. Hardware and Software Requirements and Other Technical Information

This chapter identifies the hardware and software requirements of the z/VM V5.2 base product and its optional features. This chapter also includes other technical information about z/VM V5.2. It contains the following sections:

- “Hardware and Storage Requirements”
- “Device Support” on page 24
- “Software (Program) Requirements” on page 24
- “Operating Systems Supported as Guests” on page 31
- “Other Programs Supported on z/VM” on page 31
- “National Language Support” on page 32
- “Packaging and Ordering Information” on page 32
- “Object Code Only and Limited Source Availability” on page 33
- “Limitations and Restrictions” on page 30
- “Integrity and Security” on page 34

Note: Some technical information may have changed since the publication of this edition. There may be a later edition of this publication. For the latest product updates, you should contact your IBM Support Center or visit the z/VM Web site at www.ibm.com/eserver/zseries/zvm/.

Hardware and Storage Requirements

This section identifies the hardware and storage required to install and operate z/VM V5.2. It includes the following topics:

- “Server Requirements”
- “Installation Requirements”
- “Storage Requirements” on page 22
- “CSE Hardware Requirements” on page 22
- “OSA/SF Hardware Requirements” on page 22
- “TCP/IP Hardware Requirements” on page 23
- “VM Guest RSA-Assist Support for Linux Hardware Requirements” on page 23

Server Requirements

z/VM V5.2 requires an IBM System z9 109 or zSeries server, or an equivalent server that supports IBM z/Architecture (64-bit). See Appendix A, “IBM Servers Supported by z/VM,” on page 83.

Specific server facilities are also required. See the architecture level set information at www.ibm.com/eserver/zseries/zvm/zvm520/architecture/.

Installation Requirements

z/VM V5.2 is available in the following formats:

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

To install z/VM V5.2:

Technical Information

- A local non-SNA 3270 terminal, such as:
 - Open Systems Adapter Integrated Communications Controller (OSA-ICC)
 - Integrated 3270 console available with the z9-109, z990, and z890 servers (provided through the IBM Hardware Management Console)
 - 2074 Console Support Controller
 - 3174 Control Unit
 - Equivalent console
- Installation from magnetic tape requires a tape drive capable of reading IBM 3480/3490 or 3590 tape cartridges.
- Installation from DVD requires the Hardware Management Console, V1.8 or later.
- Installation to extended count key data (ECKD) DASD requires IBM 3390 Model 3 or Model 9 DASD, or DASD that is defined to emulate these specific models.

Note: For performance reasons, IBM does not recommend installing z/VM to a real 3390 Model 9 device.

- Installation to FCP-attached SCSI LUNs requires IBM TotalStorage SCSI disks or equivalent.
- Installation first level from tape requires a minimum of 32 MB of real storage.
- Installation first level from DVD requires that a minimum of 512 MB of real storage 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 storage for the installation user ID on your first level system.

For specific DASD volume requirements, see the *z/VM: Guide for Automated Installation and Service*.

Storage Requirements

To generate and operate z/VM, at least 32 MB of real storage is required. The maximum amount of real storage that z/VM supports is 128 GB.

The block paging function requires more DASD space to transfer pages at maximum efficiency and speed.

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.

CSE Hardware Requirements

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

- The CSE complex consists of two, three, or four VM systems.
- The 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. OSA/SF requires its own user ID and also requires a maintenance user ID for its installation. Refer to 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 book. Depending on your OSA type, refer to one of the following:

- *eServer zSeries 900: Planning for the Open Systems Adapter-2 Feature*
- *eServer zSeries: Open Systems Adapter-Express Customer's Guide and Reference*

TCP/IP Hardware Requirements

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

- IBM 3172 Interconnect Controller with the Interconnect Controller Program (ICP), version 3 (5621-425), or equivalent
- IBM Open Systems Adapter-2 (OSA-2) (supported OSA-2 features are server-dependent)
- IBM Open Systems Adapter-Express (OSA-Express) (supported OSA-Express features are server-dependent)
- IBM Open Systems Adapter-Express2 (OSA-Express2) (supported OSA-Express2 features are server-dependent)
- 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, refer to 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 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 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 z/OS, using the IBM FICON Channel-to-Channel Adapter.

- IBM HiperSockets

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

TCP/IP for z/VM also requires a 3270-equivalent workstation for TCP/IP administration.

VM Guest RSA-Assist Support for Linux Hardware Requirements

In order for the VM Guest RSA-Assist Support for Linux to operate, corresponding function in Linux on System z9 and zSeries and one of the following are required:

- IBM Crypto Express2 Coprocessor or Accelerator
- IBM PCIX Cryptographic Coprocessor (PCIXCC)
- IBM PCI Cryptographic Accelerator (PCICA)
- IBM PCI Cryptographic Coprocessor (PCICC)

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 91. The use of such a device is fully supported by IBM through the service support available with z/VM. A supported device can be supported either for CP and guest use or for dedicated use by only a single guest.

A device supported for CP and guest use is one that CP and virtual machines can use. 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.

A device supported for dedicated use by a single guest can be logically attached to only a single 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 91, 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 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 or CP-owned device. You define an unsupported device by calling it a device type other than any of those shown in Appendix C, “IBM Devices Supported by z/VM,” on page 91, and by specifying the CLASS operand of the RDEVICE macro. The CLASS value should be the one that comes closest to accurately describing the unsupported device (for example, DASD or TAPE). See *z/VM: CP Planning and Administration* for details of the definition and customization processes for unsupported devices.

Software (Program) Requirements

This section identifies the additional licensed programs and other program materials required by z/VM V5.2. It includes the following topics:

- “Prerequisite Licensed Programs” on page 25
- “High Level Assembler” on page 25
- “Additional Base Facilities and Optional Features” on page 26
- “Shared-DASD Complex and CSE Cluster Program Requirements” on page 26

- “DFSMS/VM Program Requirements” on page 26
- “DirMaint Program Requirements” on page 27
- “Language Environment Program Requirements” on page 27
- “OpenExtensions Program Requirements” on page 27
- “OSA/SF Program Requirements” on page 28
- “RACF for z/VM Program Requirements” on page 28
- “TCP/IP for z/VM Program Requirements” on page 28
- “Online Books Program Requirements” on page 29
- “Program Requirements for Other IBM Products” on page 30
- “Linux Program Requirements” on page 30

Prerequisite Licensed Programs

z/VM V5.2 requires the following additional IBM licensed programs:

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

The PTFs for APARs VM63624 (to support large disk sizes on the DS6000 and DS8000) and VM63743 (to support the z9-109 server) have been applied to EREP/VM preinstalled with z/VM V5.2.

- 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 G14. 'G' indicates release 17 of ICKDSF and '14' indicates VMSES/E format.

The PTFs for APARs PQ84848 (to support large FBA devices), PQ95319 (to support large disk sizes on the DS6000), and PQ96706 (to support large disk sizes on the DS8000) have been applied to ICKDSF preinstalled with z/VM V5.2.

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 already be licensed for these programs or you must also 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, the IBM High Level Assembler, V1.4 (5696-234) or later, or an equivalent product, is required to:

- Create a new DMSTRT for system languages (NLS)
- Create image libraries for system printers (FCBs)
- Create GCS application segments (CONTENTS macro)
- Access major CMS application interfaces (CMSCALL)
- Access most CP application interfaces (DIAGNOSE)
- Modify the AVS tuning control module (AGWTUN)
- Use RAS tools (such as MDCHECK, FS2SFSE, AFTCHAIN, PRINTBLK, and PRINTFST)
- Use the API for data compression
- Use the CMS pipelines assembler macro interface
- Customize Language Environment or compile assembler routines used in mixed-language user applications
- Use exit routines in z/VM features, such as DirMaint, or in VM-related products, such as RSCS

Technical Information

- Run applications that exploit the IEEE Floating Point hardware facility
- Add devices that cannot be sensed (updating HCPRIO ASSEMBLE)
- Perform local modifications to modules written in assembler
- Service the CP Loader (HCPLDR)
- Assemble any CP modules
- Use CP exit routines
- Reassemble CP replacement parts for the RACF for z/VM optional feature

Additional Base Facilities and Optional Features

z/VM V5.2 supports the following additional base facilities and optional features:

- Additional base facilities:
 - HCD and HCM for z/VM, function level 520 (preinstalled)
 - OSA/SF, function level 440 (preinstalled)
 - TCP/IP for z/VM, level 520 (preinstalled)
- Optional features:
 - DFSMS/VM, function level 221 (available with the z/VM V5.2 SDO)
 - DirMaint, function level 510 (preinstalled disabled)
 - Performance Toolkit for VM, function level 520 (preinstalled disabled)
 - RACF for z/VM, V1.10 (preinstalled disabled)

Previous levels of these facilities and features are not supported by z/VM V5.2. Additional PTFs may be required to support z/VM V5.2 functions, as indicated in this chapter.

Shared-DASD Complex and CSE Cluster 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 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 510
Within the shared-DASD complex, all VM systems must be running the same DirMaint FL510 service level. Therefore, you must be licensed for the DirMaint FL510 feature on each z/VM version 5 system in the complex; you must be licensed for the DirMaint FL410 feature on each z/VM version 4 system in the complex; and you must be licensed for the IBM Directory Maintenance VM/ESA program product, V1.5 (5748-XE4), on each z/VM version 3 or VM/ESA version 2 system in the complex.

A CSE cluster consists of two, three, or four VM systems. 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

DFSMS/VM 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:
 - DirMaint, function level 510
- If an external security manager is needed for authorizations:
 - RACF for z/VM, V1.10, 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 520
- For remote operations:
 - Remote Spooling Communications Subsystem Networking (RSCS), V3.2 (5684-096) or later
- 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)
- Support for more than 10017 cylinders on an emulated 3390 Model 9 on TotalStorage DASD subsystems requires the PTF for APAR VM63004. Support for very large (near 1 TB) SCSI disks on the DS8000 and DS6000 requires the PTF for APAR VM63664. These PTFs have been applied to the DFSMS/VM optional feature available with the z/VM V5.2 SDO.

DirMaint Program Requirements

DirMaint, function level 510, has the following additional program requirements:

- To support very large (near 1 TB) SCSI disks on the DS8000 and DS6000, the PTF for APAR VM63700 is required.
- To support simplified user administration with the coordination of DirMaint and RACF changes and performance improvements for large user directories, the PTF for APAR VM63733 is required.

These PTFs have been applied to the DirMaint optional feature preinstalled with z/VM V5.2.

Language Environment Program Requirements

To customize Language Environment, or to compile assembler routines included in mixed-language user applications, you need the following licensed program:

- IBM High Level Assembler, V1.4 (5696-234) or later

OpenExtensions Program Requirements

To develop OpenExtensions applications in C, you need:

- One of the following compilers:
 - IBM C/C++ for z/VM, V1.1 (5654-A22)
 - IBM C for VM/ESA, V3.1 (5654-033)
- Language Environment:
 - If you compile your C program with C/C++ for z/VM, you must use the C/C++ runtime library included in the Language Environment component of z/VM.
 - If you compile your C program with C for VM/ESA, you can use either the C runtime library included in Language Environment 1.8 or the C/C++ runtime library included in the Language Environment component of z/VM.

To develop OpenExtensions applications in C++, you need:

- IBM C/C++ for z/VM, V1.1 (5654-A22)
- Language Environment component of z/VM

OSA/SF 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:

- To support the z990 and z890 servers, the PTF for APAR OA06666 is required.
- To support the z9-109 server, the PTF for APAR OA11650 is required.

These PTFs have been applied to the OSA/SF preinstalled with z/VM V5.2.

To support an OSA in TCP/IP Passthru mode in a z/VM V5.2 environment, OSA/SF requires TCP/IP for z/VM, level 520. OSA/SF is required in this mode only if access to one of the OSA's ports is being shared by more than one host program or if an ATM OSA-2 is being used as a LAN Emulation Client (LEC).

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

- *eServer zSeries 900: Planning for the Open Systems Adapter-2 Feature*
- *eServer zSeries: Open Systems Adapter-Express Customer's Guide and Reference*

RACF for z/VM Program Requirements

RACF for z/VM, V1.10, has the following additional program requirement:

- To support controlling promiscuous mode authorizations for guest LANs and virtual switches, the PTF for APAR VM63750 is required. This PTF has been applied to the RACF for z/VM optional feature preinstalled with z/VM V5.2.

To reassemble CP replacement parts for RACF for z/VM, you need the following licensed program:

- IBM High Level Assembler, V1.4 (5696-234) or later

For possible additional program requirements, see *Resource Access Control Facility: General Information*.

TCP/IP for z/VM Program Requirements

TCP/IP for z/VM, level 520, has the following additional program requirements:

- If the SSL server is to be run, one of the following Linux distributions must be installed and configured for exclusive use by the SSL server:
 - SUSE LINUX Enterprise Server 8 (SLES8) Service Pack 3 (31-bit)
 - SUSE LINUX Enterprise Server 9 (SLES9) (31-bit), with 2005-04-28 kernel 2.6.5 bug fix patch 18

- SUSE LINUX Enterprise Server 9 (SLES9) (64-bit), with 2005-04-28 kernel 2.6.5 bug fix patch 18
- Red Hat Enterprise Linux AS V3 (31-bit)
- Red Hat Enterprise Linux AS V3 (64-bit)

IBM is working with its Linux distribution partners to exploit this function in future Linux on System z9 and zSeries distributions or service updates.

- If a primary or secondary domain name server is to be run (not a caching-only name server), or the network database server is to be run:
 - IBM DB2[®] Server for VSE & VM, V6.1 (5648-A70), or IBM DB2 Server for VSE & VM, V7.1 (5697-F42), or later
- For SNAlink LU0 interface support:
 - ACF/VTAM for VM/ESA, version 4 (5654-010)
- For X.25 interface support:
 - X.25 NCP Packet Switching Interface (NPSI), V3.4 (5688-035) or later, for 3745
 - Corresponding levels of ACF/VTAM and ACF/NCP that support NPSI
- If an external security manager is needed:
 - RACF for z/VM, V1.10, or an equivalent product
- To develop programs in C:
 - One of the following compilers:
 - IBM C/C++ for z/VM, V1.1 (5654-A22)
 - IBM C for VM/ESA, V3.1 (5654-033)
 - Language Environment:
 - If you compile your C program with C/C++ for z/VM, you must use the C/C++ runtime library included in the Language Environment component of z/VM.
 - If you compile your C program with C for VM/ESA, you can use either the C runtime library included in Language Environment 1.8 or the C/C++ runtime library included in the Language Environment component of z/VM.
- 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 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 (3.0 or later is required; 4.0 or later is recommended). You can also use the Adobe Reader to print PDF books or sections of PDF books. The Adobe Reader is available free from the Adobe Web site at www.adobe.com.

BookManager Books

To view BookManager books loaded on the z/VM system, IBM BookManager READ/VM (5684-062), release 3 or later, is required. This program has the following prerequisite licensed programs:

- GDDM[®]/VM, version 3 (5684-168) release 1.1 or later

Technical Information

- Language Environment 1.8 or later
- Compiler for SAA REXX/370 (5695-013) or Library for SAA REXX/370 (5695-014)

To run READ/VM Public Library, 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* CD-ROM supplied with z/VM.

For possible additional program requirements, refer to the BookManager announcement documentation.

Linux Program Requirements

Linux on System z9 and zSeries is not provided with z/VM. In order to operate Linux as a guest of z/VM, Linux must be obtained from distributors.

For specific function and for the most current information on Linux distributions, see www.ibm.com/eserver/zseries/os/linux/dist.html.

Program Requirements for Other IBM Products

Other IBM licensed products might also have additional program requirements to run on z/VM V5.2. Refer to the product documentation and product Web sites. See the IBM VM-related licensed products and features page at www.ibm.com/eserver/zseries/zvm/related/.

Limitations and Restrictions

z/VM version 5 has the following limitations:

- z/VM will operate only on hardware servers that support z/Architecture (64-bit), including the z9-109, z990, z890, z900, and z880, or equivalent.
- Linux is the only guest operating system that z/VM will support on IFL engines, other than z/VM version 4 or z/VM version 5. (Linux can also run on IFL engines in LPAR mode.)
- Use of IFL engines on z800 and z900 machines requires LPAR mode. Customers who choose to add IFL engines to a z800 or z900 machine running in basic mode will also have to run their operating system on the standard engine(s) in LPAR mode. This is not a consideration for z9-109, z990, and z890 machines, which do not have basic mode.
- z/VM does not provide support for virtual=real (V=R) or virtual=fixed (V=F) guest images. Customers running V=R and V=F guests should consider this when migrating to z/VM version 5. The processor requirements to run as a V=V guest what was previously run as V=R may be significant due to the loss of various hardware assists.
- z/VM does not support the storage reconfiguration of partitionable processors, either by z/VM itself or by guests running on z/VM.

General Restrictions for z/VM Virtual Machines

z/VM virtual machines have the following general restrictions:

- 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.

- The CPUID that appears in most EREP records is not the same as the CPUID of the physical processor on which the error occurred (because there is more than one partition per physical processor in LPAR mode).

DirMaint Restrictions

The following restrictions apply to DirMaint, function level 510:

- 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 may 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 are expanded to four digits (if not already so) 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.

Operating Systems Supported as Guests

For information about the specific levels of IBM operating systems supported as guests, see Appendix B, "IBM Operating Systems Supported as Guests of z/VM," on page 87.

Other Programs Supported on z/VM

For information about other IBM licensed programs supported on z/VM, see www.ibm.com/eserver/zseries/zvm/related/.

For information about non-IBM Solution Developer products that run on z/VM, see www.ibm.com/eserver/zseries/zvm/vendor/.

National Language Support

The system default language for z/VM is mixed-case American English. However, z/VM's national language support allows you to install a language other than American English on your z/VM system. End users can selectively (through commands) receive messages and z/VM HELP information in that language. Any installed language can be set as the system default language. Each session on the same system can select a language that is different from another session, as long as the language selected has been installed on that system. For example, one user can be using German while another user is using Japanese.

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.

Note: Some parts of the listed items may not be translated.

Table 1. IBM-Translated Parts of z/VM

Language	IBM-Translated Parts of z/VM
German	<ul style="list-style-type: none"> • CP messages, HELP for CP commands and messages • CMS messages, HELP for CMS commands and messages • REXX/VM messages, HELP for REXX/VM statements and messages
Japanese (KANJI)	<ul style="list-style-type: none"> • CP messages, HELP for CP commands and messages • CMS messages, HELP for CMS commands and messages • REXX/VM messages, HELP for REXX/VM statements and messages • DFSMS/VM panels and messages • DirMaint messages

Packaging and Ordering Information

To order z/VM V5.2 and its features, as well as other licensed products that run on z/VM, order the z/VM V5.2 System Delivery Offering (SDO), program number 5741-A06. The z/VM V5.2 SDO package contains:

- z/VM V5.2 System Image
- z/VM SDO Optional Products Enabling Aid
- z/VM SDO licensed products and optional features, if ordered

For more information about the z/VM V5.2 SDO, see ibm.com/eserver/zseries/zvm/sdo/.

The z/VM V5.2 System Image is available as:

- Multivolume 3390 DASD Dump Restore (DDR) image on 3480/3490 or 3590 tape cartridges
- Binary image on DVD

The z/VM V5.2 System Image contains:

- z/VM base:
 - CP
 - CMS
 - TCP/IP for z/VM, level 520 (includes source)
 - AVS

- Dump Viewing Facility
- GCS
- HCD and HCM for z/VM, function level 520
- Language Environment
- OSA/SF, function level 440
- REXX/VM
- TSAF
- VMSES/E
- SFS and CRR file pools
- DASD Dump Restore (DDR) program
- 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 510
 - Performance Toolkit for VM, function level 520
 - RACF for z/VM, V1.10
- Selected related IBM licensed programs:
 - EREP/VM, V3.5 (5654-260)
 - ICKDSF for VM, V1.17 (5684-042) (includes ICKDSF Standalone release 17)
 - RSCS, V3.2 (5684-096)

Notes:

1. The DirMaint, Performance Toolkit for VM, and RACF for z/VM 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 so that you can enable them.
2. The DFSMS/VM optional feature is not supplied automatically with z/VM. If desired, DFSMS/VM must be ordered as a no-cost optional feature of the z/VM SDO.
3. ICKDSF, EREP/VM, and RSCS are not part of the z/VM product. They are separate IBM licensed programs that are included in the z/VM System Image for convenience. ICKDSF and EREP/VM are prerequisites of z/VM. Before using ICKDSF and EREP/VM you must be already licensed for them, or you must place orders for them to establish licenses. RSCS is installed disabled. If you wish to use RSCS, you must place an order to establish a license so that you can enable it.

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
- SSL server included with TCP/IP for z/VM
- DFSMS/VM
- Performance Toolkit for VM

CP, CMS, DirMaint, and RACF for z/VM 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 security and integrity of the system.

Data Integrity for Guests

Operating system failures that occur in virtual machines do not normally affect the z/VM operating system 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 testing and production work running in other virtual machines.

System Integrity Statement for z/VM

System integrity is an important characteristic of z/VM. This statement extends IBM's previous statements on system integrity to the z/VM environment.

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

The z/VM control program system integrity is the inability of any program running in a virtual machine not authorized by a z/VM control program mechanism under the customer's control or a guest operating system mechanism under the customer's control to:

- Circumvent or disable the control program real or auxiliary storage protection.
- Access a resource protected by RACF. Resources protected by RACF 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.

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

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.

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's commitment to the system integrity of the z/VM environment, as described in this statement, represents a further significant step 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 may 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.

Support for Cryptographic Facilities on IBM Servers

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

- CMOS Cryptographic Coprocessor is supported for use by guests.
- PCI Cryptographic Accelerator and PCI Cryptographic Coprocessor are supported for use by Linux on System z9 and zSeries guests.
- PCIX Cryptographic Coprocessor and Crypto Express2 are supported for use by Linux on System z9 and zSeries guests and z/OS guests.
- Configurable Crypto Express2 (Coprocessor or Accelerator) is supported for use by Linux on System z9 and zSeries guests 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 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:

- Kerberos authentication service
The Kerberos authentication server provides a way for authenticated users to prove their identity to other servers in a network. The authentication server uses the Kerberos database to verify that the client making the request is, in fact, the client named in the request.
- 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.
- Exits to an ESM
TCP/IP servers can be configured to use RACF for z/VM or another ESM to:
 - Authenticate and authorize access to resources managed by the server
 - Validate client user IDs and passwords

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 customer-tailorable filtering of unnecessary audit details.

Security Provided by RACF for z/VM

RACF 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

Technical Information

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.

Chapter 5. z/VM Base Product

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

- “Control Program”
- “Conversational Monitor System” on page 45
- “TCP/IP for z/VM” on page 49
- “APPC/VM VTAM Support” on page 52
- “Dump Viewing Facility” on page 52
- “Group Control System” on page 53
- “Hardware Configuration Definition and Hardware Configuration Manager for z/VM” on page 53
- “Language Environment” on page 54
- “Open Systems Adapter Support Facility” on page 54
- “REXX/VM” on page 54
- “Transparent Services Access Facility” on page 55
- “Virtual Machine Serviceability Enhancements Staged/Extended” on page 55

Control Program

The Control Program (CP) is primarily a 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 may require CMS to be running in an XA virtual machine.
- **XC** virtual machines process according to ESA/XC architecture.

Notes:

1. 370-mode (System/370™ architecture) virtual machines are not supported. However, the 370 Accommodation Facility allows many CMS applications 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); the virtual machine's guest real storage is paged into and out of host real storage. The V=R (virtual=real) and V=F (virtual=fixed) types of virtual machines, also known as preferred virtual machines, which permanently mapped guest real storage to host real storage, are not supported.

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 may 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 the *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

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 52.

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 V5.2,” on page 9. For detailed information on hardware capabilities, requirements, and availability, refer to the appropriate hardware announcements.
2. In some cases, the CP support is available only to guests and is not used by CP.

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.

Cryptographic Facilities

CP supports guest use of server cryptographic facilities:

- CMOS Cryptographic Coprocessor is supported for guests.
- PCI Cryptographic Coprocessor (PCICC) and Accelerator (PCICA) are supported for Linux on System z9 and zSeries and z/OS guests.
- PCIX Cryptographic Coprocessor (PCIXCC) is supported for Linux on System z9 and zSeries and z/OS guests.
- Crypto Express2 Coprocessor and Accelerator are supported for Linux on System z9 and zSeries and z/OS guests.

CP provides clear-key RSA support for a considerably large number of Linux on System z9 and zSeries guests, which helps enable hardware SSL acceleration on the servers.

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

CP supports the Extended Channel Measurement Data Support (ECMDS) provided in the z9-109, z990, and z890 servers. The extended I/O-measurement facilities in the server provide 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

CP supports the Fibre Channel Protocol (FCP) hardware channel, which 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 standard FICON channel can provide up to 120 MBps effective bandwidth, compared with 17 MBps for ESCON. A FICON Express channel can provide up to 170 MBps. A FICON Express2 channel is designed to deliver up to 270 MBps.

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 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.

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.

On z9-109 and z990 servers, 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 on System z9 and zSeries 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 Facility for Linux

The IBM Integrated Facility for Linux (IFL) is a hardware feature which provides additional processing capacity for Linux workloads. Implementation of this facility requires an LPAR definition, following normal LPAR activation procedures. Only certain applications can run on an LPAR IFL; see Appendix A, “IBM Servers Supported by z/VM,” on page 83.

Logical Channel Subsystem

The z9-109, z990, and z890 servers provide a logical channel subsystem (LCSS) structure, which 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.

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 (NPIV) 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 capability to define multiple virtual FCP channels, each with its own unique Fibre Channel port name and Fibre Channel identifier.

NPIV is exclusive to the z9-109 server and is applicable to all of the FICON features supported on the z9-109. NPIV 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.

- Open Systems Adapter-2 (OSA-2)
- Open Systems Adapter-Express (OSA-Express)
- Open Systems Adapter-Express2 (OSA-Express2)

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 z9-109, z990, and z890 servers
- Open Systems Adapter for the Network Control Program (OSA NCP), a function of the OSA-Express2 Gigabit Ethernet and 1000BASE-T Ethernet features on the z9-109 server that provides channel connectivity from System z9 and zSeries operating systems to the IBM Communication Controller for Linux (CCL) on System z9.

z/VM includes the Open Systems Adapter Support Facility (OSA/SF) for customizing an OSA's modes of operation. For more information, see "Open Systems Adapter Support Facility" on page 54.

Parallel Access Volumes

CP provides guest support for the Parallel Access Volumes feature on TotalStorage 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.

Peer-to-Peer Remote Copy

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

Processor Resource/Systems Manager

The Processor Resource/Systems Manager (PR/SM) feature allows server resources to be divided across multiple logical partitions (LPARs). Each LPAR contains some portion of the processor, storage, and channel path resources and, if available, some portion of the processor's expanded storage resources. When running in an LPAR, z/VM supports logical path measurement in an ESCON Multiple Image Facility (EMIF) environment.

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. On the z9-109, z990, and z890 servers, adapter interruptions are also used by OSA-Express and FCP channels.

In addition, the z9-109, z990, and z890 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 z9-109, z990, and z890 also provide 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.

Conversational Monitor System

The Conversational Monitor System (CMS) provides a high-capacity environment that supports 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 the *z/VM: CMS User's Guide*.

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.
- 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 may 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 57.

CMS Application Programming

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

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

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

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 the *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.

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

Base Product

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 the *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*.

Systems Management APIs

The systems management application programming interfaces (APIs) provide a standard, platform-independent client interface that reduces the amount of VM-specific programming skills required to manage resources for virtual systems (guests). The APIs include 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 APIs are invoked by a client through a set of Remote Procedure Calls (RPCs) that can be called from the C programming language on any platform that supports the Open Network Computing RPC standard. A set of CMS CSL routines provides the actual functionality to the server. These CSL routines are replaceable to allow customized solutions for those who wish to implement their own RPC functions or other applications. A dedicated z/VM server virtual machine manages the RPC-CSL interface. For more information, see *z/VM: Systems Management Application Programming*.

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 the *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.

For more information, 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 the *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, may provide HELP information for display through the HELP Facility. You can also write your own HELP information. For more information, see the *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. Using the TCP/IP protocol suite of TCP/IP for z/VM, you can reach open multi-vendor networking environments from your z/VM system. TCP/IP for z/VM allows z/VM systems to act as peers of other central computers in TCP/IP open networks. Applications can be shared transparently across VM, Linux, and other environments. As shown in Figure 1 on page 50, users can send messages, transfer files, share printers, and access remote resources across a broad range of systems from multiple vendors.

TCP/IP for z/VM provides the following types of functions:

- Connectivity and gateway functions, which handle the physical interfaces and routing of data.
- Server functions, which provide a service to a client (that is, send or transfer a file).
- Client functions, which request a certain service from a server anywhere in the network.

Base Product

- Network status and management functions, which detect and solve network problems.
- Application programming interfaces, which allow you to write your own client/server applications.

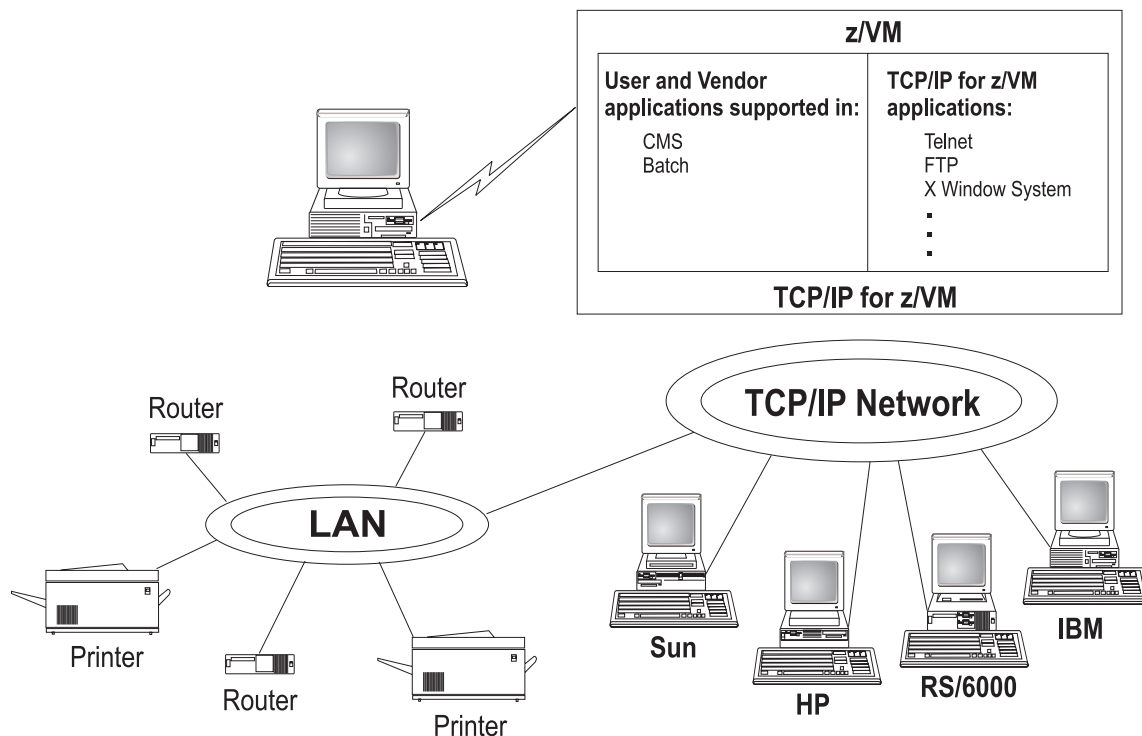


Figure 1. TCP/IP for z/VM as Part of a Multi-Vendor Network

TCP/IP consists of a layered structure of protocols and functions.

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. The link protocols include Token-Ring, Ethernet, and X.25.

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. However, a Telnet client that supports IPv6 can connect to the Telnet server in the z/VM TCP/IP stack.

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 includes the following applications and protocols:

- **Telnet Protocol** provides a standard method to interface terminal devices and terminal-oriented processes with each other.
- **File Transfer Protocol (FTP)** allows you to transfer data between local and foreign hosts or between two foreign hosts.
- **Trivial File Transfer Protocol (TFTP)** reads and writes files to and from a foreign host.
- **Internet Message Access Protocol (IMAP)** provides the processing that allows a client to access electronic mail that is kept in an IMAP Mailstore server.
- **Simple Mail Transfer Protocol (SMTP)** is an electronic mail protocol that is implemented with the CMS NOTE and SENDFILE commands.
- **Domain Name System (DNS)** is a hierarchical-naming system for naming hosts.
- **Simple Network Management Protocol (SNMP)** provides a means for managing an internet environment.
- **Kerberos Authentication System** provides additional security by allowing authorization checking at the user level rather than at the node level.
- **Remote Printing (LPR and LPD)** provides both client and server support for remote printing.
- **RouteD** uses the Routing Information Protocol (RIP) to dynamically create and maintain network routing tables.

- **MPRoute** uses Open Shortest Path First (OSPF) and RIP (either or both) to dynamically create and maintain network routing tables.
- **X Window System** supports network transparent windowing and graphics.
- **X Toolkit** is a collection of basic C language routines for developing a variety of application environments.
- **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.
- **Remote Procedure Call (RPC)** is a programming interface that calls subroutines to be executed on a foreign host.
- **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.
- **Bootstrap Protocol (BOOTP)** allows a diskless client machine to discover its own IP address, the address of a server host, and the name of a file to be loaded into memory and executed.
- **Dynamic Host Configuration Protocol (DHCP)** provides a framework for passing configuration information to hosts on a TCP/IP network.
- **Network Database System (NDB)** is used to access relational database systems in a TCP/IP-based internet environment.
- **Socket interfaces** allow you to write your own applications to supplement those supplied by TCP/IP.
- **Secure Sockets Layer (SSL)** protocol provides privacy between two communicating applications — a client and a server.

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

APPC/VM VTAM Support

APPC/VM VTAM Support (AVS) is a Virtual Telecommunications Access Method (VTAM) application that provides advanced program-to-program communication (APPC) services between VM and non-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 VM systems within the SNA network
- APPC applications residing in non-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 41.

Group Control System

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*.

Hardware Configuration Definition and Hardware Configuration Manager for z/VM

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

z/VM's original dynamic I/O configuration capabilities 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*

Open Systems Adapter Support Facility

The Open Systems Adapter-2 (OSA-2), Open Systems Adapter-Express (OSA-Express), and Open Systems Adapter Express2 (OSA-Express2) are integrated hardware features that allow System z9 and zSeries platforms 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 one of the following:

- *eServer zSeries 900: Planning for the Open Systems Adapter-2 Feature*
- *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 Restructured Extended Executor (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 the *z/VM: REXX/VM User's Guide*.

Transparent Services Access Facility

The Transparent Services Access Facility (TSAF) provides communication services within a collection of VM systems without using VTAM. TSAF runs in a CMS virtual machine. A group of up to eight VM systems that each have TSAF installed and running can form a TSAF collection. APPC/VM programs on one VM system in the TSAF collection can communicate with other APPC/VM programs on the other 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*.

Virtual Machine Serviceability Enhancements Staged/Extended

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 chapter provides an overview of the optional features of z/VM:

- “Data Facility Storage Management Subsystem for VM”
- “Directory Maintenance Facility”
- “Performance Toolkit for VM” on page 58
- “Resource Access Control Facility for z/VM” on page 59

Data Facility Storage Management Subsystem for VM

The 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-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-managed storage to DFSMS-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 provides native VM support for the IBM 3494 and 3495 Tape Library Dataservers.

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

Directory Maintenance Facility

The Directory Maintenance Facility (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 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:

Optional Features

- DirMaint operates as a CMS application. It 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 products, 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 31
- “DirMaint System Integrity” on page 35
- “DirMaint Security” on page 37
- *z/VM: Directory Maintenance Facility Tailoring and Administration Guide*

Performance Toolkit for VM

The 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.com/zseries/zos/rmf/rmfhtmls/pmweb/pmlin.htm.

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*.

Resource Access Control Facility for z/VM

Resource Access Control Facility (RACF) 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

Optional Features

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 *Resource Access Control Facility: General Information*.

Chapter 7. z/VM Library Guide

This chapter provides information about the z/VM product documentation. It contains the following sections:

- “Where to Get z/VM Documentation”
- “Books in the z/VM Base Library” on page 62
- “Books for z/VM Optional Features” on page 74
- “Book Editions and Media” on page 77

Where to Get z/VM Documentation

z/VM product documentation is available from several sources, as described in this section.

IBM Publications Center

You can obtain z/VM books from the IBM Publications Center, at www.ibm.com/shop/publications/order/. The IBM Publications Center provides worldwide downloading and ordering facilities for a broad range of IBM product documentation in many languages.

z/VM Internet Library

The z/VM Internet Library, at www.ibm.com/eserver/zseries/zvm/library/, provides links to various online z/VM information resources, such as:

- **z/VM LibraryCenter**, which provides filters for searching the z/VM bookshelf. The books in the z/VM bookshelf have been enabled for the following filters:
 - Command name
 - Message ID
- Softcopy bookshelves, which provide BookManager-format books displayed as HTML and offer standard BookManager searching capability.
- Adobe Portable Document Format (PDF) libraries
- z/VM program directories
- z/VM data areas, control blocks, and monitor records
- **LookAt**, which provides an online message lookup facility.
- IBM Redbooks™
- White papers, consultant reports, and performance reports
- Specification sheets
- Reference guides
- Marketing brochures
- Journals and bulletins
- Libraries for related IBM products

z/VM Library on Disk

The z/VM product library is available on CD-ROM and DVD:

- *IBM Online Library: z/VM Collection*, SK2T-2067
- *IBM Online Library: z/VM Collection on DVD*, SK5T-7054

These collections contain BookManager bookshelves and PDF libraries for current IBM z/VM system products and IBM licensed products that run on z/VM.

Notes:

1. Only unlicensed books are included.
2. Some products may not have PDF libraries; some books may not have PDF versions.
3. The CD-ROM is a multi-disk set; the DVD is a single disk.

These collections also include the following programs:

- **IBM Softcopy Reader** allows you to view and search BookManager files and organize BookManager books 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 CD-ROM or DVD to the workstation or host.

Basic (Automatically Shipped) z/VM Documentation

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

- One copy of the z/VM publications CD-ROM set
- One copy of the z/VM publications DVD
- One printed copy of each of the following z/VM publications:
 - *z/VM: License Information*
 - *z/VM: Guide for Automated Installation and Service*
 - *z/VM: Getting Started with Linux on System z9 and zSeries*
- A complete set of z/VM HELP files (included in the z/VM System Image)

Books in the z/VM Base Library

This section provides abstracts of the books in the z/VM base library. The base library contains books that describe the facilities included in the z/VM base product. The books are listed under the following task categories:

- “System Overview”
- “Installation, Migration, and Service” on page 63
- “Planning and Administration” on page 64
- “Customization and Tailoring” on page 66
- “Operation” on page 67
- “Application Programming” on page 67
- “End Use” on page 71
- “System Diagnosis” on page 73

Note: Within each category, z/VM branded books are listed first, followed by any non-z/VM books also included in that category.

The title of each book is followed by its publication number. For current editions and available media, see “Book Editions and Media” on page 77.

System Overview

The following books provide information to help you decide if z/VM meets the needs of your installation.

z/VM: General Information, GC24-6095

z/VM: General Information provides general information about z/VM. It contains:

- An introduction to z/VM
- Information about z/VM's cost/performance benefits
- An overview of what is new or changed in the current z/VM release
- Hardware and software requirements and other z/VM technical information
- Descriptions of the z/VM base facilities and optional features
- A guide to the z/VM product documentation
- Lists of the IBM servers, guest operating systems, and devices supported by current z/VM releases

z/VM: Glossary, GC24-6097

The *z/VM: Glossary* provides definitions of the terms used in z/VM books.

z/VM: License Information, GC24-6102

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

The following books provide information to help you install and maintain a z/VM system.

z/VM: Guide for Automated Installation and Service, GC24-6099

The *z/VM: Guide for Automated Installation and Service* 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.

This book also describes automated procedures for installing preventive and corrective service to the z/VM components and the preinstalled z/VM facilities and features. (More detailed service procedures for the z/VM components are provided in the *z/VM: Service Guide*.)

Included with this book are two additional documents that provide streamlined installation and service procedures:

- *z/VM Summary for Automated Installation and Service (Tape Installation)*
- *z/VM Summary for Automated Installation and Service (DVD Installation)*

z/VM: Migration Guide, GC24-6103

The *z/VM: Migration Guide* provides a history of significant VM product changes and enhancements since VM/ESA V2.1. This information is provided to help system support personnel migrate to the new z/VM release from a previous VM system. It is also intended for application programmers who use VM external interfaces in their programs. The book 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
- Guidance for some migration tasks you might need to perform, such as converting from system definition macros to system configuration files, and migrating your system directory

z/VM: Service Guide, GC24-6117

The *z/VM: Service Guide* contains step-by-step procedures for receiving and applying service and rebuilding serviced parts of the z/VM components. This book should be used in conjunction with the *z/VM: VMSES/E Introduction and Reference*.

z/VM: VMSES/E Introduction and Reference, GC24-6130

The *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 the *z/VM: Guide for Automated Installation and Service* and the *z/VM: Service Guide*. This book discusses the Software Inventory and how you can use it to manage the products on your system. Reference information on the product parameter file, Software Inventory tables, VMSES/E commands, and other related commands is also provided.

Planning and Administration

The following books provide information to help you make decisions about how to define and manage a z/VM system. Planning is an iterative task in that many of the decisions are made before installation, continually evaluated after installation, and 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: CMS File Pool Planning, Administration, and Operation, SC24-6074

z/VM: CMS File Pool Planning, Administration, and Operation provides guidance and reference 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 book includes a reference section that contains descriptions of file pool startup parameters and file pool administration and server commands.

z/VM: CMS Planning and Administration, SC24-6078

z/VM: CMS Planning and Administration describes how to plan for and administer the Conversational Monitor System (CMS) on your z/VM system. 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: Connectivity, SC24-6080

z/VM: Connectivity provides an overview of the z/VM facilities that enable logical connections between systems. This book 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: CP Planning and Administration, SC24-6083

z/VM: CP Planning and Administration describes how to plan for and administer the z/VM Control Program (CP). It includes information on the following topics:

- System planning and administration (including the system configuration file)
- User planning and administration (including the user directory)
- Storage planning and administration

z/VM: Getting Started with Linux on System z9 and zSeries, SC24-6096

z/VM: Getting Started with Linux on System z9 and zSeries describes how to configure and use z/VM functions and facilities for Linux servers running on the System z9 and zSeries platforms. This book is designed to help system administrators who have limited knowledge of z/VM, but want to deploy Linux servers on z/VM. The book 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 servers
- Configuring and administering a z/VM system
- Configuring TCP/IP for z/VM
- Creating and cloning Linux virtual machines
- Setting up rudimentary system automation
- Monitoring performance and capacity

z/VM: Group Control System, SC24-6098

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 book 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: I/O Configuration, SC24-6100

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 of the book describes how to use Hardware Configuration Definition and Hardware Configuration Manager for z/VM (HCD and HCM for z/VM) to manage the hardware and software I/O configuration. HCM is a Windows-based program that provides a graphical user interface. Part 2 of the book describes how to use z/VM dynamic I/O configuration, which is a command interface for modifying the hardware I/O configuration.

z/VM: Running Guest Operating Systems, SC24-6115

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 z9 and zSeries*.

z/VM: Saved Segments Planning and Administration, SC24-6116

z/VM: Saved Segments Planning and Administration describes how to plan for and administer 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: TCP/IP Planning and Customization, SC24-6125

z/VM: TCP/IP Planning and Customization is intended to help system administrators plan for TCP/IP networks on a z/VM host and customize TCP/IP to their systems.

eServer zSeries 900: Planning for the Open Systems Adapter-2 Feature, GA22-7477

eServer zSeries 900: Planning for the Open Systems Adapter-2 Feature provides information to help you plan for, install, and use the Open Systems Adapter-2 (OSA-2) hardware feature on the z900 server.

eServer zSeries: Open Systems Adapter-Express Customer's Guide and Reference, SA22-7935

The *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 features on zSeries servers.

eServer zSeries: Open Systems Adapter-Express Integrated Console Controller User's Guide, SA22-7990

The *eServer zSeries: 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 zSeries servers.

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

The *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.

Customization and Tailoring

The following books provide information about extending or enhancing a z/VM system.

z/VM: CP Exit Customization, SC24-6082

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: Performance, SC24-6109

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

The following books provide information about how to operate a z/VM system and z/VM virtual machines.

z/VM: System Operation, SC24-6121

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

- Using the Stand-Alone Program Loader
- 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: Virtual Machine Operation, SC24-6128

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

Application Programming

The following books provide information about designing, coding, compiling, running, debugging, and maintaining application programs to do specific functions on z/VM.

z/VM: CMS Application Development Guide, SC24-6069

The *z/VM: CMS Application Development Guide* provides information about developing an application program in CMS. The development process includes planning, designing, writing, compiling, debugging, executing, and updating. This book 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: CMS Application Development Guide for Assembler, SC24-6070

The *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: CMS Application Multitasking, SC24-6071

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 book provides introductory and tutorial information as well as detailed reference material.

z/VM: CMS Callable Services Reference, SC24-6072

The *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: CMS Macros and Functions Reference, SC24-6075

The *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: CP Programming Services, SC24-6084

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: CPI Communications User's Guide, SC24-6085

The *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: Enterprise Systems Architecture/Extended Configuration Principles of Operation, SC24-6094

The *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: Language Environment User's Guide, SC24-6101

The *z/VM: Language Environment User's Guide* provides information for using Language Environment on z/VM. It describes the differences between using Language Environment on z/OS and z/VM and points to the z/OS Language Environment documentation for more information.

z/VM: OpenExtensions Advanced Application Programming Tools, SC24-6104

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

z/VM: OpenExtensions Callable Services Reference, SC24-6105

The *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 book also describes callable services that are not related to the standards.

z/VM: OpenExtensions Commands Reference, SC24-6106

The *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: OpenExtensions POSIX Conformance Document, GC24-6107

The *z/VM: OpenExtensions POSIX Conformance Document* describes the z/VM implementation of those areas of the IEEE POSIX.1 and POSIX.2 standards that were declared to be optional or implementation-defined.

z/VM: OpenExtensions User's Guide, SC24-6108

The *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: Program Management Binder for CMS, SC24-6110

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 book describes the differences in usage and behavior between the two programs. This book should be used in conjunction with:

- *z/OS MVS Program Management: Advanced Facilities*, SA22-7644
- *z/OS MVS Program Management: User's Guide and Reference*, SA22-7643

z/VM: Reusable Server Kernel Programmer's Guide and Reference, SC24-6112

The *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 book covers advanced material in server construction and is not intended for beginning programmers.

z/VM: REXX/VM Reference, SC24-6113

The *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: REXX/VM User's Guide, SC24-6114

The *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: Systems Management Application Programming, SC24-6122

z/VM: Systems Management Application Programming describes application programming interfaces (APIs) for performing systems management functions for virtual systems (guests) in a z/VM environment. Two sets of APIs are described:

- Remote procedure calls (RPCs), which can be called from C applications on any platform that supports the RPC standard.
- Callable services library (CSL) routines, which are invoked by the RPCs. These CSL routines can also be called directly from CMS applications, and they are replaceable to allow for customization.

z/VM: TCP/IP Programmer's Reference, SC24-6126

The *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 program interfaces (APIs):

- C sockets
- IUCV sockets
- Pascal
- Virtual Machine Communication Facility (VMCF)
- Remote Procedure Calls (RPCs)
- X Window System
- Kerberos Authentication System
- Simple Network Management Protocol (SNMP) agent distributed program interface
- Network Computing System (NCS)
- CMS command interface to the name server
- Simple Mail Transfer Protocol (SMTP)

Common Programming Interface Communications Reference, SC26-4399

The *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

The *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: Language Environment Concepts Guide, SA22-7567

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

z/OS: Language Environment Debugging Guide, GA22-7560

The *z/OS: Language Environment Debugging Guide* provides assistance in detecting and locating programming errors that may occur during runtime under Language Environment. The book 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

The *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

The *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

The *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 ILC Applications, SA22-7563

z/OS: Language Environment Writing ILC 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

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

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

End Use

The following books describe how to perform various user tasks on z/VM.

z/VM: CMS Commands and Utilities Reference, SC24-6073

The *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: CMS Pipelines Reference, SC24-6076

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

z/VM: CMS Pipelines User's Guide, SC24-6077

The *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 multi-stream 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: CMS Primer, SC24-6137

The *z/VM: CMS Primer* provides information on basic CMS tasks and commands. It presents, through examples, only 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: CMS User's Guide, SC24-6079

The *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 book 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 read the *z/VM: CMS Primer* before reading this book.

z/VM: CP Commands and Utilities Reference, SC24-6081

The *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: Quick Reference, SC24-6111

The *z/VM: Quick Reference* contains an alphabetical listing of z/VM commands, utilities, and statements. It shows the syntax of each function, indicates the environment and intended user, and identifies the reference book where you can find detailed information. This book also contains a summary of reference information about communicating with z/VM and interpreting its responses.

z/VM: TCP/IP User's Guide, SC24-6127

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

- Transferring files
- Sending electronic mail
- Logging on to a foreign host
- Monitoring the network
- Authenticating network users
- Remote printing
- Managing network resources
- Using the Domain Name System

z/VM: XEDIT Commands and Macros Reference, SC24-6131

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

z/VM: XEDIT User's Guide, SC24-6132

The *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

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 book includes reference information on the PIPE command and Pipelines stages, commands, and messages. This book contains additional information not found in the *z/VM: CMS Pipelines Reference* and also provides information intended for the advanced Pipelines user.

System Diagnosis

The following books provide information to help identify, describe, report, and correct a programming problem in a z/VM system.

z/VM: CMS and REXX/VM Messages and Codes, GC24-6118

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: CP Messages and Codes, GC24-6119

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: Diagnosis Guide, GC24-6092

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

z/VM: Dump Viewing Facility, GC24-6093

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: Other Components Messages and Codes, GC24-6120

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: TCP/IP Diagnosis Guide, GC24-6123

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

z/VM: TCP/IP Messages and Codes, GC24-6124

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

z/VM: VM Dump Tool, GC24-6129

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 book 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

z/OS and z/VM: Hardware Configuration Definition Messages provides explanations of the messages issued by Hardware Configuration Definition (HCD).

Books for z/VM Optional Features

This section provides abstracts of the books that describe the z/VM optional features. The title of each book is followed by its publication number. For current editions and available media, see “Book Editions and Media” on page 77.

DFSMS/VM

z/VM: DFSMS/VM Customization, SC24-6086

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: DFSMS/VM Diagnosis Guide, GC24-6087

The *z/VM: DFSMS/VM Diagnosis Guide* provides information about diagnosing and reporting DFSMS/VM errors.

z/VM: DFSMS/VM Messages and Codes, GC24-6088

z/VM: DFSMS/VM Messages and Codes contains explanations and suggested actions for messages and codes issued by DFSMS/VM.

z/VM: DFSMS/VM Planning Guide, SC24-6089

The *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: DFSMS/VM Removable Media Services, SC24-6090

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: DFSMS/VM Storage Administration, SC24-6091

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.

DirMaint**z/VM: Directory Maintenance Facility Commands Reference, SC24-6133**

The *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 book also describes the additional utilities that DirMaint provides for minidisk management and command security.

z/VM: Directory Maintenance Facility Messages, GC24-6134

z/VM: Directory Maintenance Facility Messages provides the explanations, failing component names, and suggested corrective actions for the messages issued by DirMaint.

z/VM: Directory Maintenance Facility Tailoring and Administration Guide, SC24-6135

The *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: Performance Toolkit, SC24-6136**

z/VM: Performance Toolkit 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 have been provided to facilitate the operation of VM systems, thereby improving operator efficiency and productivity.
- **Performance monitoring on VM.** An enhanced realtime performance monitor allows systems programmers to monitor system performance and to analyze bottlenecks. Features included in the monitor, and the manner in which data is displayed, have both been designed to improve the systems programmer's productivity when analyzing the system, and to allow even a new user to work efficiently with the tool. The program can help systems programmers to make more efficient use of system resources, increase system productivity and improve end-user satisfaction.

RACF for z/VM

External Security Interface (RACROUTE) Macro Reference for MVS and VM, GC28-1366

The *External Security Interface (RACROUTE) Macro Reference for MVS and VM* describes the interface that MVS™ and VM provide to an external security manager (ESM). This information is intended for programmers who are writing applications that need to invoke RACF (or another ESM) from MVS or VM, or for programmers who are writing an ESM for MVS or VM.

Resource Access Control Facility: Auditor's Guide, SC28-1342

The *Resource Access Control Facility: 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.

Resource Access Control Facility: Command Language Reference, SC28-0733

The *Resource Access Control Facility: Command Language Reference* describes the syntax and functions of RACF commands.

Resource Access Control Facility: Diagnosis Guide, GY28-1016

The *Resource Access Control Facility: Diagnosis Guide* explains how to diagnose problems that may occur in RACF.

Resource Access Control Facility: General Information, GC28-0722

Resource Access Control Facility: General Information contains an overview of RACF and highlights the latest enhancements.

Resource Access Control Facility: General User's Guide, SC28-1341

The *Resource Access Control Facility: 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.

Resource Access Control Facility: Macros and Interfaces, SC28-1345

Resource Access Control Facility: 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 book does not document the RACROUTE macro or the independent RACF system macros, which are documented in the *External Security Interface (RACROUTE) Macro Reference for MVS and VM*.

Resource Access Control Facility: Messages and Codes, SC38-1014

Resource Access Control Facility: Messages and Codes describes RACF messages, abend codes, manager return codes, and utility return codes.

Resource Access Control Facility: Migration and Planning, GC23-3054

Resource Access Control Facility: Migration and Planning provides information to guide installations through the migration process from RACF for VM, 1.9.2, to RACF for VM, 1.10.0. This book is not intended for customers installing RACF for the first time.

Resource Access Control Facility: Security Administrator's Guide, SC28-1340

The *Resource Access Control Facility: 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

Resource Access Control Facility: System Programmer's Guide, SC28-1343

The *Resource Access Control Facility: 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

Book Editions and Media

This section indicates the current edition and available media for each book in the z/VM library. Table 2 on page 78 lists the books for the z/VM base facilities. Table 3 on page 80 lists the books for the z/VM optional features.

The following notations are used in the tables:

- NEW** indicates that the book is a **new addition** to the z/VM library.
- ✓ indicates that the book has been updated with a **new edition** for this z/VM release. Updates may be major or minor. If the book is not z/VM branded, the updates might not be z/VM related. Refer to the book's Summary of Changes.
- *** indicates that a printed copy of the book is shipped automatically with z/VM. It can also be ordered from the IBM Publications Center for a fee.
- SC** indicates that the book is provided in softcopy (PDF and BookManager) versions. Softcopy editions are included on the z/VM publications CD-ROM and DVD. They are also available from the IBM Publications Center and the z/VM Internet Library.
- HC** indicates that a hard copy (printed) version of the book is available, but it must be ordered from the IBM Publications Center for a fee.

Notes:

1. The current editions of the z/VM publications CD-ROM and DVD are:
 - *IBM Online Library: z/VM Collection*, SK2T-2067-23
 - *IBM Online Library: z/VM Collection on DVD*, SK5T-7054-00
2. BookManager files are primarily intended for online viewing and searching. They are not intended for printing, although some programs may provide that

Library Guide

function. If you attempt to print BookManager files, you may encounter unexpected results or error conditions. Use the PDF version of a book for printing.

z/VM Base Library

The following table shows the current edition and available media for each book included in the z/VM base library.

Note: z/VM branded books are listed first, followed by other branded books also included in the z/VM base library.

Table 2. Current Edition and Available Media for Books in the z/VM Base Library

Title	Edition	Updated	Media
z/VM: CMS and REXX/VM Messages and Codes	GC24-6118-01	✓	SC
z/VM: CMS Application Development Guide	SC24-6069-01	✓	SC
z/VM: CMS Application Development Guide for Assembler	SC24-6070-00		SC
z/VM: CMS Application Multitasking	SC24-6071-00		SC
z/VM: CMS Callable Services Reference	SC24-6072-01	✓	SC
z/VM: CMS Commands and Utilities Reference	SC24-6073-01	✓	SC
z/VM: CMS File Pool Planning, Administration, and Operation	SC24-6074-01	✓	SC
z/VM: CMS Macros and Functions Reference	SC24-6075-00		SC
z/VM: CMS Pipelines Reference	SC24-6076-01	✓	SC
z/VM: CMS Pipelines User's Guide	SC24-6077-01	✓	SC
z/VM: CMS Planning and Administration	SC24-6078-01	✓	SC
z/VM: CMS Primer	SC24-6137-00		SC HC
z/VM: CMS User's Guide	SC24-6079-00		SC
z/VM: Connectivity	SC24-6080-02	✓	SC
z/VM: CP Commands and Utilities Reference	SC24-6081-02	✓	SC
z/VM: CP Exit Customization	SC24-6082-01	✓	SC
z/VM: CP Messages and Codes	GC24-6119-02	✓	SC
z/VM: CP Planning and Administration	SC24-6083-02	✓	SC
z/VM: CP Programming Services	SC24-6084-01	✓	SC
z/VM: CPI Communications User's Guide	SC24-6085-00		SC
z/VM: Diagnosis Guide	GC24-6092-01	✓	SC
z/VM: Dump Viewing Facility	GC24-6093-01	✓	SC
z/VM: Enterprise Systems Architecture/Extended Configuration Principles of Operation	SC24-6094-01	✓	SC
z/VM: General Information	GC24-6095-03	✓	SC ¹
z/VM: Getting Started with Linux on System z9 and zSeries	SC24-6096-01	✓	* SC
z/VM: Glossary	GC24-6097-01	✓	SC
z/VM: Group Control System	SC24-6098-01	✓	SC
z/VM: Guide for Automated Installation and Service	GC24-6099-02	✓	* SC ^{1, 2}
z/VM: I/O Configuration	SC24-6100-01	✓	SC
z/VM: Language Environment User's Guide	SC24-6101-01	✓	SC
z/VM: License Information	GC24-6102-02	✓	* SC ²

Table 2. Current Edition and Available Media for Books in the z/VM Base Library (continued)

Title	Edition	Updated	Media
z/VM: Migration Guide	GC24-6103-02	✓	SC
z/VM: OpenExtensions Advanced Application Programming Tools	SC24-6104-01	✓	SC
z/VM: OpenExtensions Callable Services Reference	SC24-6105-01	✓	SC
z/VM: OpenExtensions Commands Reference	SC24-6106-01	✓	SC
z/VM: OpenExtensions POSIX Conformance Document	GC24-6107-00		SC
z/VM: OpenExtensions User's Guide	SC24-6108-01	✓	SC
z/VM: Other Components Messages and Codes	GC24-6120-01	✓	SC
z/VM: Performance	SC24-6109-01	✓	SC
z/VM: Program Management Binder for CMS	SC24-6110-01	✓	SC
z/VM: Quick Reference	SC24-6111-01	✓	SC
z/VM: Reusable Server Kernel Programmer's Guide and Reference	SC24-6112-01	✓	SC
z/VM: REXX/VM Reference	SC24-6113-00		SC
z/VM: REXX/VM User's Guide	SC24-6114-01	✓	SC
z/VM: Running Guest Operating Systems	SC24-6115-01	✓	SC
z/VM: Saved Segments Planning and Administration	SC24-6116-01	✓	SC
z/VM: Service Guide	GC24-6117-01	✓	SC
z/VM: System Operation	SC24-6121-01	✓	SC
z/VM: Systems Management Application Programming	SC24-6122-01	✓	SC
z/VM: TCP/IP Diagnosis Guide	GC24-6123-01	✓	SC
z/VM: TCP/IP Messages and Codes	GC24-6124-01	✓	SC
z/VM: TCP/IP Planning and Customization	SC24-6125-01	✓	SC
z/VM: TCP/IP Programmer's Reference	SC24-6126-00		SC
z/VM: TCP/IP User's Guide	SC24-6127-01	✓	SC
z/VM: Virtual Machine Operation	SC24-6128-01	✓	SC
z/VM: VM Dump Tool	GC24-6129-01	✓	SC
z/VM: VMSES/E Introduction and Reference	GC24-6130-01	✓	SC
z/VM: XEDIT Commands and Macros Reference	SC24-6131-00		SC
z/VM: XEDIT User's Guide	SC24-6132-01	✓	SC
CMS/TSO Pipelines Author's Edition	SL26-0018-03	✓	SC
Common Programming Interface Communications Reference	SC26-4399-09 ³		SC ⁴ HC
Common Programming Interface Resource Recovery Reference	SC31-6821-01		SC ⁴ HC
eServer zSeries 900: Planning for the Open Systems Adapter-2 Feature	GA22-7477-00		SC
eServer zSeries: Open Systems Adapter-Express Customer's Guide and Reference	SA22-7935-04	✓	SC
eServer zSeries: Open Systems Adapter-Express Integrated Console Controller User's Guide	SA22-7990-01	✓	SC
z/OS: Language Environment Concepts Guide	SA22-7567-06	✓	SC
z/OS: Language Environment Debugging Guide	GA22-7560-06	✓	SC

Library Guide

Table 2. Current Edition and Available Media for Books in the z/VM Base Library (continued)

Title	Edition	Updated	Media
z/OS: Language Environment Programming Guide	SA22-7561-06	✓	SC
z/OS: Language Environment Programming Reference	SA22-7562-06	✓	SC
z/OS: Language Environment Run-Time Messages	SA22-7566-06	✓	SC
z/OS: Language Environment Writing ILC Applications	SA22-7563-04	✓	SC
z/OS and z/VM: Hardware Configuration Definition Messages	SC33-7986-06	✓	SC
z/OS and z/VM: Hardware Configuration Manager User's Guide	SC33-7989-06	✓	SC
z/OS MVS Program Management: Advanced Facilities	SA22-7644-04	NEW	SC
z/OS MVS Program Management: User's Guide and Reference	SA22-7643-03	NEW	SC
Notes on the z/VM Base Library: 1. Not included on the <i>z/VM Collection</i> or <i>z/VM Collection on DVD</i> . 2. Not available in a BookManager version. 3. The latest BookManager version of the <i>Common Programming Interface Communications Reference</i> is SC26-4399-08. 4. Not available in a PDF version.			

z/VM Optional Features

The following table shows the current edition and available media for each book provided for the z/VM optional features.

Table 3. Current Edition and Available Media for Books for z/VM Optional Features

Feature / Book Title	Number	Updated	Media
DFSMS/VM			
z/VM: DFSMS/VM Customization	SC24-6086-01	✓	SC
z/VM: DFSMS/VM Diagnosis Guide	GC24-6087-00		SC
z/VM: DFSMS/VM Messages and Codes	GC24-6088-00		SC
z/VM: DFSMS/VM Planning Guide	SC24-6089-00		SC
z/VM: DFSMS/VM Removable Media Services	SC24-6090-01	✓	SC
z/VM: DFSMS/VM Storage Administration	SC24-6091-01	✓	SC
DirMaint			
z/VM: Directory Maintenance Facility Commands Reference	SC24-6133-01	✓	SC
z/VM: Directory Maintenance Facility Messages	GC24-6134-01	✓	SC
z/VM: Directory Maintenance Facility Tailoring and Administration Guide	SC24-6135-01	✓	SC
Performance Toolkit for VM			
z/VM: Performance Toolkit	SC24-6136-02	✓	SC
RACF for z/VM			
External Security Interface (RACROUTE) Macro Reference for MVS and VM	GC28-1366-10	✓	SC
Resource Access Control Facility: Auditor's Guide	SC28-1342-14	✓	SC
Resource Access Control Facility: Command Language Reference	SC28-0733-19	✓	SC
Resource Access Control Facility: Diagnosis Guide	GY28-1016-09	✓	SC

Table 3. Current Edition and Available Media for Books for z/VM Optional Features (continued)

Feature / Book Title	Number	Updated	Media
Resource Access Control Facility: General Information	GC28-0722-20	✓	SC
Resource Access Control Facility: General User's Guide	SC28-1341-11	✓	SC
Resource Access Control Facility: Macros and Interfaces	SC28-1345-10	✓	SC
Resource Access Control Facility: Messages and Codes	SC38-1014-19	✓	SC
Resource Access Control Facility: Migration and Planning	GC23-3054-06	✓	SC
Resource Access Control Facility: Security Administrator's Guide	SC28-1340-15	✓	SC
Resource Access Control Facility: System Programmer's Guide	SC28-1343-13	✓	SC

Appendix A. IBM Servers Supported by z/VM

This appendix provides information about the IBM servers supported by current z/VM operating systems. It addresses z/VM operating systems running first level on the hardware or in a logical partition. For guest considerations, see Appendix B, “IBM Operating Systems Supported as Guests of z/VM,” on page 87.

The following terms are used in the server support matrix:

Term	Meaning
PR/SM	Processor Resource/Systems Manager. A hardware feature that supports LPAR mode and MPG capability.
LPAR	Logically Partitioned or Logical Partition. On a server with PR/SM operating in LPAR mode, server resources are divided across multiple logical partitions. The number of LPARs depends on the server model.
IFL	IBM Integrated Facility for Linux. A hardware feature that provides additional processing capacity for Linux workloads in an LPAR.
MPG	Multiple Preferred Guests. On a server with PR/SM operating in basic mode, a z/VM host operating system can support up to six preferred (high performance) guests.

Note: z/VM version 5 does not support preferred guests.

The following general notes apply to this matrix:

- z/VM can exploit the ESCON Multiple Image Facility (EMIF) when running in an LPAR.
- The inclusion of a server in this matrix does not imply the server is available in all countries.
- 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.
- For specific hardware service levels required, contact your IBM representative.
- For information about older servers and earlier z/VM releases, see the previous editions of this book.

Table 4. z/VM Server Support Matrix

Server	Mode	z/VM Support
S/390 Multiprise® 3000 Enterprise Server	Basic	V4.4 (2)
	LPAR	V4.4 (2)
	LPAR IFL (1)	V4.4 (2)

Server Support

Table 4. z/VM Server Support Matrix (continued)

Server	Mode	z/VM Support
Notes on Multiprise 3000:		
(1)	Only the following run in LPAR IFL mode on this server: Linux on S/390, z/VM V4.4, and CMS and standalone utilities supplied with those z/VM releases. For ICKDSF support information, see the ICKDSF Program Directory.	
(2)	This server does not support using the integrated 3270 console task on the Hardware Management Console as a system operator console for z/VM V4.4.	
S/390 Parallel Enterprise Server™ Generation 5 (9672 G5)	Basic	V4.4 (2,3)
	LPAR	V4.4 (2,3)
	LPAR IFL (1)	V4.4 (3)
S/390 Parallel Enterprise Server Generation 6 (9672 G6)	Basic	V4.4 (2,3)
	LPAR	V4.4 (2,3)
	LPAR IFL (1)	V4.4 (3)
Notes on G5 and G6:		
(1)	Only the following run in LPAR IFL mode on this server: Linux on S/390, z/VM V4.4, and CMS and standalone utilities supplied with z/VM V4.4. For ICKDSF support information, see the ICKDSF Program Directory.	
(2)	The CMOS Cryptographic Processor is supported for use by guests.	
(3)	The PCI Cryptographic Accelerator and PCI Cryptographic Coprocessor are supported for use by Linux on S/390 guests. APAR VM62933 is required for DirMaint support.	
@server zSeries 800 (z800)	Basic	V4.4 (2,3), V5.1 (2,3), V5.2 (2,3)
	LPAR	V4.4 (2,3), V5.1 (2,3), V5.2 (2,3)
	LPAR IFL (1)	V4.4 (3), V5.1 (3), V5.2 (3)
@server zSeries 900 (z900)	Basic	V4.4 (2,3), V5.1 (2,3), V5.2 (2,3)
	LPAR	V4.4 (2,3), V5.1 (2,3), V5.2 (2,3)
	LPAR IFL (1)	V4.4 (3), V5.1 (3), V5.2 (3)
@server zSeries 890 (z890)	LPAR	V4.4 (2,3,4), V5.1 (2,3,4,6), V5.2 (2,3,5,6)
	LPAR IFL (1)	V4.4 (3,4), V5.1 (3,4,6), V5.2 (3,5,6)
@server zSeries 990 (z990)	LPAR	V4.4 (2,3,4), V5.1 (2,3,4,6), V5.2 (2,3,5,6)
	LPAR IFL (1)	V4.4 (3,4), V5.1 (3,4,6), V5.2 (3,5,6)
Notes on zSeries:		
(1)	Only the following run in LPAR IFL mode on this server: Linux on System z9 and zSeries, z/VM V4.4, z/VM V5.1 or later, and CMS and standalone utilities supplied with those z/VM releases. For ICKDSF support information, see the ICKDSF Program Directory.	
(2)	The CMOS Cryptographic Processor is supported for use by guests.	
(3)	The PCI Cryptographic Accelerator and PCI Cryptographic Coprocessor are supported for use by Linux on System z9 and zSeries and z/OS guests. On z/VM V4.4, APAR VM62933 is required for DirMaint support.	
(4)	HCD/HCM function level 440 support for the z890 and z990 requires APAR VM63610.	
(5)	OSA/SF function level 440 support for the z890 and z990 requires APAR OA06666.	
(6)	The PCIX Cryptographic Coprocessor (Crypto Express) and Crypto Express2 Coprocessor are supported for use by Linux on System z9 and zSeries and z/OS guests.	
System z9 109 (z9-109)	LPAR	V4.4 (2,3,4,6,7,8), V5.1 (2,3,4,5,6,7,8,9), V5.2 (3,4,5,6,7)
	LPAR IFL (1)	V4.4 (2,3,6,7,8), V5.1 (2,3,5,6,7,8,9), V5.2 (3,5,6,7)

Table 4. z/VM Server Support Matrix (continued)

Server	Mode	z/VM Support
Notes on z9-109:		
(1)	Only the following run in LPAR IFL mode on this server: Linux on System z9 and zSeries, z/VM V4.4, z/VM V5.1 or later, and CMS and standalone utilities supplied with those z/VM releases. For ICKDSF support information, see the ICKDSF Program Directory.	
(2)	HCD/HCM function level 440 support for the z9-109 requires APAR VM63721.	
(3)	OSA/SF function level 440 support for the z9-109 requires APAR OA11650. The PTF for this APAR has been applied to OSA/SF preinstalled with z/VM V5.2.	
(4)	The CMOS Cryptographic Processor is supported for use by guests.	
(5)	The configurable Crypto Express2 Accelerator is supported for use by Linux on System z9 and zSeries and z/OS guests. z/VM V5.1 requires APAR VM63646.	
(6)	N_Port ID Virtualization (NPIV) is supported for guest use when FCP subchannels are dedicated to a guest. CP use of NPIV is supported on z/VM V5.1 and later; CP use on z/VM V5.1 requires APAR VM63744.	
(7)	EREP support for the z9-109 requires APAR VM63743. The PTF for this APAR has been applied to EREP preinstalled with z/VM V5.2.	
(8)	CMS IOCP support for the z9-109 requires APAR VM63740.	
(9)	CP support for OSA-Express2 OSN on the z9-109 requires APAR VM63722.	

Appendix B. IBM Operating Systems Supported as Guests of z/VM

This appendix provides information about the IBM operating systems supported as guests of current z/VM releases.

In general, a version or release of an operating system is supported as a guest of z/VM only on servers for which support has been announced for that operating system version or release to run native, in an LPAR, or as a guest of z/VM. For example, if an operating system is not supported to run in an LPAR IFL, that operating system is not supported as a guest of z/VM running in an LPAR IFL.

This appendix 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 83.

Running z/VM in an LPAR offers production-level performance for V=V workloads, including a mix of CMS and full-function operating systems. An operating system running as a V=V guest of a z/VM image running in an LPAR will experience performance comparable to running as a guest of a z/VM system running natively. Several hardware assists available to z/VM running native are not available to z/VM in an LPAR.

Notes:

1. When z/VM V4.4 is run in an LPAR, V=F guests are not supported and I/O passthrough is not available to the V=R guest. For more details, consult the *PR/SM Planning Guide*, GA22-7123, and documentation for the applicable server.
2. z/VM version 5 does support V=R or V=F guests.

Running z/VM in a virtual machine (that is, z/VM as a guest of z/VM, also known as “second-level” z/VM) 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, and IBM will not accept performance PMRs against such environments.

The following symbols and abbreviations are used in the “Support” column of the matrix to indicate the level of support:

Symbol	Meaning
—	Guest is not supported.
V	Guest is supported in a V=V virtual machine.
R	Either of the following: <ol style="list-style-type: none">1. Guest is supported in a V=R virtual machine.2. Guest is supported in a V=F virtual machine if the server has the PR/SM feature installed.
U	Guest may be UP (uniprocessor).
M	Guest may be MP (multiprocessor).

The following general notes apply to the matrix:

Guest Support

- All listed guests are supported in ESA virtual machines. XA virtual machines are supported for compatibility and are functionally equivalent to ESA virtual machines.
- For information about product releases not listed, see the previous editions of this book.

Table 5. z/VM Guest Support Matrix

Guest	Host	Support
Linux on S/390	z/VM V4.4 z/VM V5.1 z/VM V5.2	VR UM V UM V UM
Linux on System z9 and zSeries	z/VM V4.4 z/VM V5.1 z/VM V5.2	VR UM (1) V UM V UM
Notes for Linux on System z9 and zSeries Guests:		
(1) The z/VM host must be running on a z9-109 or zSeries server and the 64-bit CP image must be loaded.		
TPF V4 (1)	z/VM V4.4 z/VM V5.1 z/VM V5.2	VR UM V UM V UM
Notes for TPF Guests:		
(1) TPF is supported as a guest of z/VM in environments where TPF-specific server and DASD control unit RPQs are not required. If in a V=V virtual machine, use of the Emulation Program for 37xx communication controllers is not supported.		
z/TPF V1 (1)	z/VM V4.4 z/VM V5.1 z/VM V5.2	VR UM V UM V UM
Notes for z/TPF Guests:		
(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. If in a V=V virtual machine, use of the Emulation Program for 37xx communication controllers is not supported.		
VSE/ESA™ V2	z/VM V4.4 z/VM V5.1 z/VM V5.2	VR UM V UM V UM
z/VSE V3 (1)	z/VM V4.4 z/VM V5.1 z/VM V5.2	VR UM V UM V UM
Notes for z/VSE Guests:		
(1) z/VSE can execute in 31-bit mode only. It does not implement z/Architecture, and specifically does not implement 64-bit mode capabilities. z/VSE is designed to exploit selected features of System z9 and zSeries hardware.		
z/OS V1 (1)	z/VM V4.4 z/VM V5.1 z/VM V5.2	VR UM (2,3) V UM (3) V UM (3)
Notes for z/OS Guests:		
(1) Exploitation of hardware function by z/OS may require specific levels of z/VM. For details, see <i>z/OS and z/OS.e: Planning for Installation</i> , GA22-7504.		
(2) On an S/390 Multiprise 3000 or S/390 G5 or G6, the z/OS guest runs only in 31-bit mode. On a z9-109 or zSeries server, the z/OS guest must run in 64-bit mode on a 64-bit CP image.		
(3) z/VM does not support external coupling facilities or real coupling links.		

Table 5. z/VM Guest Support Matrix (continued)

Guest	Host	Support
z/OS.e V1 (1)	z/VM V4.4 z/VM V5.1 z/VM V5.2	VR UM (2,4) V UM (3,4) V UM (3,4)
Notes for z/OS.e Guests: (1) Exploitation of hardware function by z/OS.e may require specific levels of z/VM. For details, see <i>z/OS and z/OS.e: Planning for Installation</i> , GA22-7504. (2) The z/VM host must be running the 64-bit CP image on a z890 or z800. (3) The z/VM host must be running on a z890 or z800. (4) z/VM does not support external coupling facilities or real coupling links.		
z/VM V4.4	z/VM V4.4 z/VM V5.1 z/VM V5.2	VR UM (1,2,3) V UM (2) V UM (2)
z/VM V5.1	z/VM V4.4 z/VM V5.1 z/VM V5.2	VR UM (1,2,4) V UM (2) V UM (2)
z/VM V5.2	z/VM V4.4 z/VM V5.1 z/VM V5.2	VR UM (1,2,4) V UM (2) V UM (2)
Notes for z/VM Guests: (1) V=R recovery is not supported for a V=R guest. (2) z/VM Guest Coupling Simulation will not work in the z/VM guest. (3) The z/VM guest can run in 64-bit mode only when the z/VM host is running a 64-bit CP image on a z9-109 or zSeries server. (4) The z/VM host must be running a 64-bit CP image on a z9-109 or zSeries server.		

Appendix C. IBM Devices Supported by z/VM

This appendix provides information about the IBM devices supported by current z/VM operating systems.

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 may be supported; consult device announcements and publications for this information.

Device support might not be in the base release of any system shown, but may have been delivered through the service stream by 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 guest operating systems or applications in virtual machines to use the device. There are usually no restrictions on the use of the device by a virtual machine operating system 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 guest operating systems or applications in virtual machines to use the device. The guest or virtual machine application is completely responsible for the management and use of the device. Usage restrictions may 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 may 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 may 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 may apply. There may be exceptions to supported devices and releases depending on your specific configuration. See the *z/VM: CP Planning and Administration* book for your release and hardware announcements for additional information.

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 specific details about the system configuration file, see *z/VM: CP Planning and Administration*. For specific details 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 older devices and earlier z/VM releases, see the previous editions of this book.

Direct Access Storage Devices (DASD)

Table 6. Device Support List—DASD

Device	z/VM Support	Support Notes
3370 Model A01, B01, A02, B02	V4.4	Not supported for installation of z/VM.
3375 Model A1, B1, D1	V4.4	Not supported for installation of z/VM.
3380 Model A04, AA4, B04, AD4, BD4, AE4, BE4, AJ4, BJ4, AK4, BK4	V4.4	Not supported for installation of z/VM.
3380 Model CJ2	V4.4	Not supported for installation of z/VM.
3390 Model A14, A18, B14, B18, B1C, A24, A28, B24, B28, B2C	V4.4, V5.1, V5.2	Supported in 3390 mode (native mode) with VSE/VSAM 2.1.0 or later. Mixed 3390 models are not supported for installation of z/VM. Not supported for installation of z/VM V5.1 or later.
3390 Model A34, A38, B34, B38, B3C	V4.4, V5.1, V5.2	Supported in 3390 mode (native mode) with VSE/VSAM 2.1.0 or later. Mixed 3390 models are not supported for installation of z/VM.
3390 Model A94, A98, B94, B98, B9C	V4.4, V5.1, V5.2	Not supported for installation of z/VM V4.4. Although installation of z/VM V5.1 or later to this device is supported, it is not recommended. Intended as a mass-storage device. Should not be used for system data or for data for applications that require high performance DASD. 3380 Track Compatibility Mode is not supported. Minidisks used with VSE/VSAM or CMS/VSAM are limited to 65,536 tracks (4369 cylinders).
3995 Model 131, 132 Optical Library Dataserver	V4.4	Must be dedicated to a guest as a 3088. Not supported for installation of z/VM.
3995 Model 151, 111 Optical Library Dataserver	V4.4	Must be defined as a 3390. Not supported for installation of z/VM. Intended as a mass-storage device. Should not be used for system data or for data for applications that require high performance DASD.

Table 6. Device Support List—DASD (continued)

Device	z/VM Support	Support Notes
3995 Model 153, 113 Optical Library Dataserver	V4.4	<p>Must be defined as a 3390.</p> <p>Not supported for installation of z/VM.</p> <p>Intended as a mass-storage device. Should not be used for system data or for data for applications that require high performance DASD.</p>
9332 Model 400, 402, 600, 602	V4.4	Not supported for installation of z/VM.
9335 Model B01	V4.4	Not supported for installation of z/VM.
9336 Model 010, 020, 025	V4.4	Not supported for installation of z/VM.
9340 Direct Access Storage Subsystem: 9345 Model B12, B22	V4.4	Not supported for installation of z/VM.
RAMAC® Array DASD: 9391 Model A10 Rack with 9392 Model B13 Drawers	V4.4, V5.1, V5.2	<p>Must be defined as a 3390. Attaches to 3990 Model 3 or Model 6 Storage Control.</p> <p>Only 3390 Model 3 is supported for installation of z/VM V5.1 or later.</p>
RAMAC Virtual Array Subsystem: 9393 Model 002, T42, or T82	V4.4, V5.1, V5.2	<p>Emulates a 3990 Model 3 Storage Control with 3390 Model 1, 2, and 3 DASD or 3380 Model J and K DASD.</p> <p>Supported for installation of z/VM only when emulating a 3390.</p> <p>Mixed 3390 models are not supported for installation of z/VM.</p> <p>Only 3390 Model 3 is supported for installation of z/VM V5.1 or later.</p>
RAMAC Array Subsystem: 9394 Model 001, 002, or 003 Cluster Array Controller with 9395 Model B13 or B23 Drawers	V4.4, V5.1, V5.2	<p>Emulates a 3990 Model 2 Storage Control with 3390 Model 3 DASD or 3380 Model K DASD.</p> <p>Supported for installation of z/VM only when emulating a 3390.</p> <p>Mixed 3390 models are not supported for installation of z/VM.</p> <p>Only 3390 Model 3 is supported for installation of z/VM V5.1 or later.</p>

Device Support

Table 6. Device Support List—DASD (continued)

Device	z/VM Support	Support Notes
Internal Disk (for Multiprise 3000)	V4.4	<p>This device is a feature of the S/390 Multiprise 3000 Enterprise Server.</p> <p>Emulates a 3990 Model 2 Storage Control with 3390 Model 1, 2, 3, and 9 DASD or 3380 Model E, J, and K DASD.</p> <p>Supported for installation of z/VM only when emulating a 3390 Model 1, 2, or 3.</p> <p>Mixed 3390 models are not supported for installation of z/VM.</p> <p>The performance-oriented track level commands are supported for guest use only.</p> <p>IOCP requires APAR VM62312.</p> <p>For ICKDSF support information, see the ICKDSF Program Directory.</p>

Table 6. Device Support List—DASD (continued)

Device	z/VM Support	Support Notes
TotalStorage Enterprise Storage Server [®] (2105) [ESCON or FICON attached]	V4.4, V5.1, V5.2	<p>Supported only in emulation mode. Must be defined as a 3390 Model 3 or 6 Storage Control with either of the following:</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>Mixed 3390 models are not supported for installation of z/VM.</p> <p>3390 Model 9 is not supported for installation of z/VM V4.4.</p> <p>Only 3390 Model 3 or 9 is supported for installation of z/VM V5.1 or later.</p> <p>The performance-oriented track level commands are supported for guest use only.</p> <p>To use the Parallel Access Volumes feature, the device must be dedicated to a guest.</p> <p>To support more than 10017 cylinders on a 3390 Model 9:</p> <ul style="list-style-type: none"> • DFSMS/VM Function Level 221 requires APAR VM63004. • DirMaint Function Level 410 requires APAR VM62907. <p>For guest use of FlashCopy, the device must be dedicated to the guest or defined as a fullpack minidisk.</p> <p>PPRC is supported for guest use only. Native use of PPRC requires ICKDSF V1R17.</p> <p>For ICKDSF support information, see the ICKDSF Program Directory.</p>
TotalStorage Enterprise Storage Server (2105) [FCP attached]	V5.1, V5.2	<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 6. Device Support List—DASD (continued)

Device	z/VM Support	Support Notes
TotalStorage DS8000 (2107) [ESCON or FICON attached]	V4.4, V5.1, V5.2	<p>APARs VM63534 and VM63653 are required on z/VM V4.4 or V5.1.</p> <p>Supported only in emulation mode. Must be defined as a 3390 Model 3 or 6 Storage Control with either of the following:</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>Mixed 3390 models are not supported for installation of z/VM.</p> <p>3390 Model 9 is not supported for installation of z/VM V4.4.</p> <p>Only 3390 Model 3 or 9 is supported for installation of z/VM V5.1 or later.</p> <p>The performance-oriented track level commands are supported for guest use only.</p> <p>To use the Parallel Access Volumes feature, the device must be dedicated to a guest.</p> <p>HCD and HCM support on z/VM V4.4 or V5.1 requires APAR VM63450.</p> <p>DFSMS/VM FL221 support for 3390 Model 9 disks larger than 10017 cylinders requires APAR VM63004.</p> <p>DirMaint FL410 support for 3390 Model 9 disks with 10018-32760 cylinders requires APAR VM62907.</p> <p>DirMaint FL410 support for 3390 Model 9 disks larger than 32760 cylinders requires APAR VM63656.</p> <p>DirMaint FL510 support for 3390 Model 9 disks larger than 32760 cylinders requires APAR VM63700.</p> <p>For guest use of FlashCopy, the device must be dedicated to the guest or defined as a fullpack minidisk.</p> <p>PPRC is supported for guest use only. Native use of PPRC requires ICKDSF V1R17.</p> <p>For ICKDSF support information, see the ICKDSF Program Directory.</p>

Table 6. Device Support List—DASD (continued)

Device	z/VM Support	Support Notes
TotalStorage DS8000 (2107) [FCP attached]	V5.1, V5.2	<p>APARs VM63534 and VM63653 are required on z/VM V5.1.</p> <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>HCD and HCM support on z/VM V5.1 requires APAR VM63450.</p> <p>DFSMS/VM FL221 support for SCSI LUNs larger than 381 GB requires APAR VM63664.</p> <p>DirMaint FL510 support for SCSI LUNs larger than 381 GB requires APAR VM63700.</p>

Device Support

Table 6. Device Support List—DASD (continued)

Device	z/VM Support	Support Notes
TotalStorage DS6000 (1750) [FICON attached]	V4.4, V5.1, V5.2	<p>APARs VM63535 and VM63653 are required on z/VM V4.4 or V5.1.</p> <p>Supported only in emulation mode. Must be defined as a 3390 Model 3 or 6 Storage Control with either of the following:</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>Mixed 3390 models are not supported for installation of z/VM.</p> <p>3390 Model 9 is not supported for installation of z/VM V4.4.</p> <p>Only 3390 Model 3 or 9 is supported for installation of z/VM V5.1 or later.</p> <p>The performance-oriented track level commands are supported for guest use only.</p> <p>To use the Parallel Access Volumes feature, the device must be dedicated to a guest.</p> <p>HCD and HCM support on z/VM V4.4 or V5.1 requires APAR VM63450.</p> <p>DFSMS/VM FL221 support for 3390 Model 9 disks larger than 10017 cylinders requires APAR VM63004.</p> <p>DirMaint FL410 support for 3390 Model 9 disks with 10018-32760 cylinders requires APAR VM62907.</p> <p>DirMaint FL410 support for 3390 Model 9 disks larger than 32760 cylinders requires APAR VM63656.</p> <p>DirMaint FL510 support for 3390 Model 9 disks larger than 32760 cylinders requires APAR VM63700.</p> <p>For guest use of FlashCopy, the device must be dedicated to the guest or defined as a fullpack minidisk.</p> <p>PPRC is supported for guest use only. Native use of PPRC requires ICKDSF V1R17.</p> <p>For ICKDSF support information, see the ICKDSF Program Directory.</p>

Table 6. Device Support List—DASD (continued)

Device	z/VM Support	Support Notes
TotalStorage DS6000 (1750) [FCP attached]	V5.1, V5.2	<p>APARs VM63535 and VM63653 are required on z/VM V5.1.</p> <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>HCD and HCM support on z/VM V5.1 requires APAR VM63450.</p> <p>DFSMS/VM FL221 support for SCSI LUNs larger than 381 GB requires APAR VM63664.</p> <p>DirMaint FL510 support for SCSI LUNs larger than 381 GB requires APAR VM63700.</p>
TotalStorage DS4000	V5.2	<p>For native z/VM use, the DS4000 volumes appear as 9336-20 DASD.</p> <p>For guest use, the DS4000 volumes can appear either as 9336-20 DASD or as SCSI LUNs directly accessed through FCP subchannels.</p>

DASD Control Units and Storage Controls

Table 7. Device Support List—DASD Control Units and Storage Controls

Device	z/VM Releases Providing Support	Support Notes
3830 Model 2	V4.4	
3880 Model 1, 3, 4	V4.4	
3880 Model E21, G21, H21, J21, D23, E23, G23, H23, J23	V4.4	
3990 Model 1, 2	V4.4, V5.1, V5.2	
3990 Model 3	V4.4, V5.1, V5.2	Concurrent Copy function is supported for guest use only.
3990 Model 6	V4.4, V5.1, V5.2	<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>
9221 DASD/Tape Subsystem Control	V4.4	
9335 Model A01	V4.4	

Device Support

Table 7. Device Support List—DASD Control Units and Storage Controls (continued)

Device	z/VM Releases Providing Support	Support Notes
9340 Direct Access Storage Subsystem: 9341 Model A02 9343 Model C02, C04	V4.4	
9340 Direct Access Storage Subsystem: 9343 Model D04	V4.4	ESCON-attached only.
9340 Direct Access Storage Subsystem: 9343 Model CC2, CC4	V4.4	Models CC2 and CC4 are cache models (basic caching only). The cache function is not visible to or controllable by software. No additional z/VM support is required.
9340 Direct Access Storage Subsystem: 9343 Model DC4	V4.4	ESCON-attached only. Model DC4 is a cache model (basic caching only). The cache function is not visible to or controllable by software. No additional z/VM support is required.

Tape Units and Tape Libraries

Table 8. Device Support List—Tape Units and Tape Libraries

Device	z/VM Releases Providing Support	Support Notes
2440	V4.4	Must be defined as a 3420 Model 4. Not supported for installation of z/VM.
3420 Model 3, 4, 5, 6, 7, 8	V4.4	Not supported for installation of z/VM.
3422 Tape Unit and Control	V4.4	Not supported for installation of z/VM.
3424 Tape Subsystem	V4.4	For Brazil only. Not supported for installation of z/VM.
3430	V4.4	Not supported for installation of z/VM.
3480 Tape Subsystem	V4.4, V5.1, V5.2	Non-XF not supported for installation of z/VM.
3490 Tape Subsystem	V4.4, V5.1, V5.2	
3490E (Enhanced Capability Model) Tape Subsystem	V4.4, V5.1, V5.2	
TotalStorage Enterprise Automated Tape Library (3494)	V4.4, V5.1, V5.2	Native support is provided in conjunction with DFSMS/VM. Not supported for installation of z/VM.
TotalStorage Virtual Tape Server (3494 VTS)	V4.4, V5.1, V5.2	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	V4.4, V5.1, V5.2	Native support is provided in conjunction with DFSMS/VM. Not supported for installation of z/VM.

Table 8. Device Support List—Tape Units and Tape Libraries (continued)

Device	z/VM Releases Providing Support	Support Notes
TotalStorage Enterprise Tape System (3590)	V4.4, V5.1, V5.2	To support the 3590 Model H, DFSMS/VM FL221 requires APAR VM63097.
TotalStorage Enterprise Tape Drive (3592)	V4.4, V5.1, V5.2	Not supported for installation of z/VM. For basic device support: <ul style="list-style-type: none"> z/VM V4.4 requires APAR VM63325. DFSMS/VM FL221 requires APAR VM63353. To support Write Once Read Many (WORM) media: <ul style="list-style-type: none"> z/VM V4.4 requires APAR VM63461. DFSMS/VM FL221 requires APAR VM63460.
9348 Model 11, 12	V4.4	Not supported for installation of z/VM.

Tape Control Units

Table 9. Device Support List—Tape Control Units

Device	z/VM Releases Providing Support	Support Notes
TotalStorage Enterprise Tape Controller 3590 Model A60	V4.4, V5.1, V5.2	
3803 Model 1, 2	V4.4	

Printers

Table 10. Device Support List—Printers

Device	z/VM Releases Providing Support	Support Notes
3203 Model 5	V4.4, V5.1, V5.2	
3262	V4.4, V5.1, V5.2	
3268 Model 2, 2C	V4.4, V5.1, V5.2	Must be defined as a 3287.
3287 Model 1, 1C, 2, 2C, 4	V4.4, V5.1, V5.2	
3289 Model 1, 3, 4, 8	V4.4, V5.1, V5.2	
3800 Model 1	V4.4, V5.1, V5.2	
3800 Model 3, 6, 8	V4.4, V5.1, V5.2	Full support in Model 1 compatibility mode or using Advanced Function Printing™ (AFP™) licensed programs.
3812	V4.4, V5.1, V5.2	Full support through RSCS V3.1.1 or later, using Advanced Function Printing licensed programs; downloading of fonts is not supported. Loading of fonts requires Print Services Facility™/VM (PSF/VM) licensed program (program number 5684-141).

Device Support

Table 10. Device Support List—Printers (continued)

Device	z/VM Releases Providing Support	Support Notes
3816 Model 01D, 01S	V4.4, V5.1, V5.2	Full support through RSCS V3.1.1 or later, using AFP licensed programs; downloading of fonts is not supported. Loading of fonts requires PSF/VM licensed program.
3820	V4.4, V5.1, V5.2	Full support through VM/VTAM, using AFP licensed programs.
3825	V4.4, V5.1, V5.2	Full support using AFP licensed programs.
3827	V4.4, V5.1, V5.2	Full support using AFP licensed programs.
3835	V4.4, V5.1, V5.2	Full support using AFP licensed programs.
3900	V4.4, V5.1, V5.2	Must be defined as a 3800.
4245 Model 1	V4.4, V5.1, V5.2	Supported as a 4245 or in 3262 compatibility mode.
4245 Model 12, 20	V4.4, V5.1, V5.2	
4248 Model 1, 2	V4.4, V5.1, V5.2	Supported as a 4248 or in 3211 compatibility mode.
6262 Model 14, 22	V4.4, V5.1, V5.2	Must be defined as a 4248 Model 1.

Card Readers and Card Punches

Table 11. Device Support List—Card Readers and Card Punches

Device	z/VM Releases Providing Support	Support Notes
3505 Model B1, B2	V4.4, V5.1, V5.2	
3525 Model P1, P2, P3	V4.4, V5.1, V5.2	

Terminals/Displays/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 function of the Hardware Management Console.

Display Printers

Most display printers are supported using Advanced Function Printing (AFP) licensed programs through RSCS.

Display Control Units

Table 12. Device Support List—Display Control Units

Device	z/VM Releases Providing Support	Support Notes
3174	V4.4, V5.1, V5.2	
3272 Model 2	V4.4, V5.1, V5.2	
3274	V4.4, V5.1, V5.2	
3276 Display/Control Unit	V4.4, V5.1, V5.2	Must be dedicated to a virtual machine as a supported device.

Communications Controllers

Table 13. Device Support List—Communications Controllers

Device	z/VM Releases Providing Support	Support Notes
3705	V4.4, V5.1	Must be dedicated to a virtual machine as a supported device.
3720	V4.4, V5.1	Must be dedicated to a virtual machine as a supported device. Must be defined as a 3705, unless you are using dynamic system configuration.
3725	V4.4, V5.1	Must be dedicated to a virtual machine as a supported device. Must be defined as a 3705, unless you are using dynamic system configuration.
3745	V4.4, V5.1, V5.2	Must be dedicated to a virtual machine as a supported device. Must be defined as a 3705, unless you are using dynamic system configuration.
9221 Integrated Communications Attachment	V4.4	

Switches

Table 14. Device Support List—Switches

Device	z/VM Releases Providing Support	Support Notes
2032 FICON Director Model 1	V4.4, V5.1, V5.2	
9032 ESCON Director Model 2	V4.4, V5.1, V5.2	
9032 ESCON Director Model 3	V4.4, V5.1, V5.2	
9032 ESCON Director Model 5	V4.4, V5.1, V5.2	
9033 ESCON Director Model 1	V4.4, V5.1, V5.2	
9033 ESCON Director Model 4	V4.4, V5.1, V5.2	
<p>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.</p> <p>For more information about ESCON architecture, see <i>Introducing Enterprise Systems Connection</i>, GA23-0383.</p>		

Miscellaneous Devices and Server Features

Table 15. Device Support List—Miscellaneous Devices

Device/Feature	z/VM Releases Providing Support	Support Notes
Channel-to-Channel Adapter (CTCA)	V4.4, V5.1, V5.2	Must be dedicated to a virtual machine as a supported device.
Crypto Express2	V5.1, V5.2	<p>Must be dedicated to a Linux on System z9 and zSeries or z/OS guest as a supported device.</p> <p>Configurable Crypto Express2 feature (Coprocessor or Accelerator) supported on z9-109; z/VM V5.1 support requires APAR VM63646.</p>
ESCON CTCA	V4.4, V5.1, V5.2	Extended mode operation only.
FICON CTCA	V4.4, V5.1, V5.2	Must be dedicated to a guest as a supported device.
FICON Express2	V5.1, V5.2	Must be dedicated to a guest as a supported device.
HiperSockets	V4.4, V5.1, V5.2	Must be dedicated to a virtual machine as a supported device.
OSA-2	V4.4, V5.1, V5.2	<p>Must be dedicated to a virtual machine as a supported device.</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> <p>OSA-2 is recognized as TYPE=OSA.</p>

Table 15. Device Support List—Miscellaneous Devices (continued)

Device/Feature	z/VM Releases Providing Support	Support Notes
OSA-Express	V4.4, V5.1, V5.2	<p>Must be dedicated to a virtual machine as a supported device.</p> <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> <p>OSA-Express is recognized as TYPE=OSE for the Fast Ethernet and 155 ATM modes.</p> <p>OSA-Express is recognized as TYPE=OSD for the Gigabit Ethernet mode. Support for CHPID type OSD on z/VM V4.4 or V5.1 requires APAR VM63524.</p> <p>OSA-Express is recognized as TYPE=OSC for the 1000BASE-T Ethernet feature with the Integrated Console Controller (ICC) function.</p> <p>Support for more TCP/IP stacks on z/VM V4.4 or V5.1 requires APARs PQ91421 and VM63524.</p>

Device Support

Table 15. Device Support List—Miscellaneous Devices (continued)

Device/Feature	z/VM Releases Providing Support	Support Notes
OSA-Express2	V4.4, V5.1, V5.2	<p>Must be dedicated to a virtual machine as a supported device.</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> <p>OSA-Express2 is recognized as TYPE=OSD for the Gigabit Ethernet and 10 Gigabit Ethernet modes. Support for CHPID type OSD on z/VM V4.4 or V5.1 requires APAR VM63524.</p> <p>OSA-Express2 is recognized as TYPE=OSC for the 1000BASE-T Ethernet feature with the Integrated Console Controller (ICC) function.</p> <p>OSA-Express2 is recognized as TYPE=OSN for the Gigabit Ethernet and 1000BASE-T Ethernet features with the Open Systems Adapter for the Network Control Program (OSA NCP) function on the z9-109. Not supported on z/VM V4.4; z/VM V5.1 requires APAR VM63722.</p> <p>Support for more TCP/IP stacks on z/VM V5.1 requires APARs PQ91421 and VM63524. Not supported on z/VM V4.4.</p>
PCI Cryptographic Accelerator	V4.4, V5.1, V5.2	<p>Must be dedicated to a Linux on System z9 and zSeries or z/OS guest as a supported device.</p> <p>DirMaint Function Level 410 requires APAR VM62933.</p>
PCI Cryptographic Coprocessor	V4.4, V5.1, V5.2	<p>Must be dedicated to a Linux on System z9 and zSeries or z/OS guest as a supported device.</p> <p>DirMaint Function Level 410 requires APAR VM62933.</p>
PCIX Cryptographic Coprocessor (Crypto Express)	V5.1, V5.2	<p>Must be dedicated to a Linux on System z9 and zSeries or z/OS guest as a supported device.</p>
3088	V4.4, V5.1, V5.2	<p>Must be dedicated to a virtual machine as a supported device.</p>
3088 Model 61 PSCA Card (common link access to workstations)	V4.4, V5.1, V5.2	
3172 Model 1, 2, 3	V4.4, V5.1, V5.2	<p>Must be dedicated to a virtual machine as a supported device.</p>
3174 Model 12L, 22L	V4.4, V5.1, V5.2	

Table 15. Device Support List—Miscellaneous Devices (continued)

Device/Feature	z/VM Releases Providing Support	Support Notes
3737 Remote Channel-to-Channel Unit Model 2	V4.4, V5.1, V5.2	Must be dedicated to a virtual machine as a supported device. Must be defined as a CTCA.
3890™	V4.4, V5.1, V5.2	Must be dedicated to a guest as a supported device.
4753 Network Security Processor	V4.4, V5.1, V5.2	Must be dedicated to a guest as a supported device.
7171 Device Attachment Control Unit	V4.4, V5.1, V5.2	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.
8232 LAN Channel Station	V4.4, V5.1	Must be dedicated to a virtual machine as a supported device. Supported by TCP/IP for z/VM.
9034 ESCON Converter Model 1	V4.4, V5.1, V5.2	
9035 ESCON Converter Model 2	V4.4, V5.1, V5.2	

Notices

IBM may not offer the products, services, or features discussed in this document in all 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 give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, New York 10594-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites 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:

IBM Corporation
Mail Station P300
2455 South Road
Poughkeepsie, New York 12601-5400
U.S.A.
Attention: Information Request

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 information 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 illustrates 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 IBM's application programming interfaces. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, or other countries, or both:

Advanced Function Printing	Print Services Facility
AFP	Processor Resource/Systems Manager
APL2	RACF
BookManager	RAMAC
DB2	Redbooks
DFSMS/VM	RMF
DirMaint	RS/6000
e(logo)server	S/390
ECKD	S/390 Parallel Enterprise Server
Enterprise Storage Server	SAA
Enterprise Systems Architecture/390	System/370
Enterprise Systems Connection Architecture	Systems Application Architecture
ESCON	SystemView
eServer	System z9
FICON	Tivoli
FlashCopy	TotalStorage
GDDM	Virtualization Engine
HiperSockets	VisualAge
IBM	VM/ESA
IBMLink	VSE/ESA
Language Environment	VTAM
MQSeries	WebSphere
Multiprise	z/Architecture
MVS	z/OS
NetView	z/TPF
OpenExtensions	z/VM
OS/390	z/VSE
Parallel Sysplex	zSeries
Performance Toolkit for VM	3890
PR/SM	

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks 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.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Index

Numerics

370 Accommodation Facility 39
370 virtual machine 39
370-XA architecture compatibility 40
64-bit addressing (z/Architecture) exploitation 9

A

abstracts of z/VM books 62
accessibility 4
ACF/VTAM 53
adapter interruptions 45
analyzing dump data 41
APIs 3
APPC/VM
 AVS 52
 ISFC 41
 TSAF 55
APPC/VM VTAM Support (AVS) 52
application development and deployment, support for
 CMS support 46
 enhancements 18
 overview 2
architecture compatibility 40
architecture support
 370-XA compatibility 40
 ESA/370 compatibility 40
 ESA/390 39
 ESA/XC 40
 ESCON 42
 FCP 42
 FICON 42
 HiperSockets 43
 IEEE Floating Point 43
 LCSS 43
 POSIX 47
 QDIO 45
 TCP/IP 49
 z/Architecture 40
auditability 35
auditable events 38
AVS 52

B

base facilities
 AVS 52
 CMS 45
 CP 39
 Dump Viewing Facility 52
 GCS 53
 HCD and HCM for z/VM 53
 Language Environment 54
 OSA/SF 54
 REXX/VM 54
 TCP/IP for z/VM 49
 TSAF 55

base facilities (*continued*)
 VMSES/E 55
benefits offered by z/VM V5 5
BFS (byte file system) 47
block paging 22
books, z/VM
 abstracts
 base library - application programming 67
 base library - customization 66
 base library - end use 71
 base library - installation, migration, and service 63
 base library - operation 67
 base library - planning and administration 64
 base library - system diagnosis 73
 base library - system overview 62
 DFSMS/VM 74
 DirMaint 75
 OpenExtensions 68
 Performance Toolkit for VM 75
 RACF for z/VM 76
 basic (automatically shipped) 62
 editions 77
 media 77
 online
 CD-ROM and DVD 61
 program requirements 29
byte file system (BFS) 47

C

cables, fiber optic 42
card readers and punches supported by z/VM 102
CD-ROM
 online books 61
 z/VM System Image 32
changes in z/VM V5.2
 60 LPARs 16
 64-bit addressing exploitation 9
 application development and deployment 18
 automation of local modification process 19
 CD-ROM not available 20
 CMS binder upgrade 18
 Crypto Express2 Accelerator 14
 device support withdrawn 19
 directory management performance
 improvements 18
 DS4000 support 16
 dynamic LPAR naming 16
 guest support 9
 hardware and architecture support 14
 installation, migration, and service improvements 18
 LAN sniffer 13
 migration of preinstalled z/VM products 18
 MPRoute server 13
 N_Port ID Virtualization (NPIV) 15
 networking enhancements 13
 OSA-Express2 OSN 10

- changes in z/VM V5.2 *(continued)*
 - packaging modifications 19
 - performance assists for z/VM guests, enhanced 11
 - SCSI disk I/O performance improvements 10
 - security enhancements 13
 - self-management 9
 - simplified user administration 17
 - SSL server enhancements 12
 - statement of direction 20
 - System z9 109 support 14
 - systems management improvements 17
 - technology exploitation 14
 - Vector Facility support removed 19
 - virtual image management enhancements 12
 - virtualization technology 9
 - withdrawn functions or support 19
 - xyz 16
- channel subsystem
 - channel program access protection 35
- CMS (Conversational Monitor System)
 - application programming 46
 - byte file system (BFS) 47
 - HELP Facility 49
 - multitasking, application 47
 - national language support 32
 - OpenExtensions 47
 - overview 45
 - Pipelines 46
 - Program Management Binder for CMS 48
 - Reusable Server Kernel 48
 - Shared File System (SFS) 46
 - systems management APIs 48
 - Virtual Machine Resource Manager (VMRM) 48
 - XEDIT 49
- communications controllers supported by z/VM 103
- compatibility
 - 370-XA architecture 40
 - ESA/370 architecture 40
 - System/370 architecture 39
- components, z/VM
 - AVS 52
 - CMS 45
 - CP 39
 - Dump Viewing Facility 52
 - GCS 53
 - Language Environment 54
 - REXX/VM 54
 - TSAF 55
 - VMSES/E 55
- concurrent-sense facility support 41
- consoles supported by z/VM 102
- Control Program (CP)
 - architecture compatibility 40
 - CSE 40
 - ESA/XC architecture support 40
 - guest LAN 40
 - hardware and architecture support
 - adapter interruptions 45
 - concurrent-sense facility 41
 - Crypto Express2 42
 - cryptographic facilities 42
- Control Program (CP) *(continued)*
 - hardware and architecture support *(continued)*
 - ECMDS 42
 - ESCON 42
 - expanded storage 42
 - Extended-TOD-Clock Facility 42
 - FCP 42
 - FICON channels 42
 - FICON CTCA 43
 - FlashCopy 43
 - HiperSockets 43
 - IEEE Floating Point 43
 - integrated 3270 console 43
 - Integrated Facility for Linux 43
 - LCSS 43
 - Move-Page Facility 44
 - NPIV 44
 - OSA 44
 - Parallel Access Volumes 44
 - PCICA 42
 - PCICC 42
 - PCIXCC 42
 - PPRC 44
 - PR/SM 45
 - QDIO architecture 45
 - ISFC 41
 - national language support 32
 - overview 39
 - virtual machine
 - architectures 39
 - type 40
 - VM Dump Tool 41
 - z/Architecture support 40
- control units supported by z/VM
 - DASD 99
 - display 103
 - storage 99
 - tape 101
- controllers, communications, supported by z/VM 103
- Conversational Monitor System (CMS)
 - application programming 46
 - byte file system (BFS) 47
 - HELP Facility 49
 - multitasking, application 47
 - national language support 32
 - OpenExtensions 47
 - overview 45
 - Pipelines 46
 - Program Management Binder for CMS 48
 - Reusable Server Kernel 48
 - Shared File System (SFS) 46
 - systems management APIs 48
 - Virtual Machine Resource Manager (VMRM) 48
 - XEDIT 49
- CP (Control Program)
 - architecture compatibility 40
 - CSE 40
 - ESA/XC architecture support 40
 - guest LAN 40
 - hardware and architecture support
 - adapter interruptions 45

- CP (Control Program) *(continued)*
 - hardware and architecture support *(continued)*
 - concurrent-sense facility 41
 - Crypto Express2 42
 - cryptographic facilities 42
 - ECMDS 42
 - ESCON 42
 - expanded storage 42
 - Extended-TOD-Clock Facility 42
 - FCP 42
 - FICON channels 42
 - FICON CTCA 43
 - FlashCopy 43
 - HiperSockets 43
 - IEEE Floating Point 43
 - integrated 3270 console 43
 - Integrated Facility for Linux 43
 - LCSS 43
 - Move-Page Facility 44
 - NPIV 44
 - OSA 44
 - Parallel Access Volumes 44
 - PCICA 42
 - PCICC 42
 - PCIXCC 42
 - PPRC 44
 - PR/SM 45
 - QDIO architecture 45
 - ISFC 41
 - national language support 32
 - overview 39
 - virtual machine
 - architectures 39
 - type 40
 - VM Dump Tool 41
 - z/Architecture support 40
- cross system extensions (CSE)
 - description 40
 - hardware requirements 22
 - program requirements 26
- Crypto Express2 Accelerator 14
- Crypto Express2, support for 42
- cryptographic facilities, support for 42
- CSE (cross system extensions)
 - description 40
 - hardware requirements 22
 - program requirements 26
- customer responsibilities for data security 35

D

- DASD (direct access storage device)
 - control units supported by z/VM 99
 - CSE requirement 22
 - Parallel Access Volumes support 44
 - required to install z/VM 22
 - supported by z/VM 92
- data integrity
 - customer responsibilities 35
 - for guests 34
- dedicated-only device, definition of 24

- device support withdrawn 19
- device support, types of 24
- devices supported by VM
 - card readers and punches 102
 - communications controllers 103
 - consoles 102
 - DASD 92
 - DASD control units 99
 - display control units 103
 - display printers 102
 - displays 102
 - miscellaneous 104
 - printers 101
 - storage controls 99
 - switches 103
 - tape control units 101
 - tape libraries 100
 - tape units 100
 - terminals 102
- DFSMS/VM
 - books, abstract of 74
 - description 57
 - national language support 32
 - program requirements 26
- diagnosing system problems 52
- direct access storage device (DASD)
 - control units supported by z/VM 99
 - CSE requirement 22
 - Parallel Access Volumes support 44
 - required to install z/VM 22
 - supported by z/VM 92
- DirMaint (Directory Maintenance Facility)
 - books, abstract of 75
 - description 57
 - national language support 32
 - program requirements 27
 - restrictions 31
 - security 37
 - system integrity 35
- display control units supported by z/VM 103
- display printers supported by z/VM 102
- displays supported by z/VM 102
- DOS/VSE support 46
- DS4000 support 16
- dump data, analyzing 41
- Dump Viewing Facility 52
- DVD
 - installing z/VM from 22
 - online books 61
 - z/VM System Image 32

E

- ECMDS (extended channel measurement data support) 42
- editions, book 77
- editor, z/VM 49
- Enterprise Systems Connection Architecture (ESCON)
 - description 42
 - devices supported 103
 - ESCON Multiple Image Facility (EMIF) 45, 83

- ESA virtual machine 39
- ESA/370 architecture compatibility 40
- ESA/XC architecture 40
- ESCON (Enterprise Systems Connection Architecture)
 - description 42
 - devices supported 103
 - ESCON Multiple Image Facility (EMIF) 45, 83
- expanded storage support 42
- extended channel measurement data support (ECMDS) 42
- Extended Configuration (ESA/XC) architecture 40
- Extended-TOD-Clock Facility support 42

F

- FCP (Fibre Channel Protocol) support 42
- features, optional
 - DFSMS/VM 57
 - DirMaint 57
 - Performance Toolkit for VM 58
 - RACF for z/VM 59
- fiber optic cables 42
- FICON (fibre connection) channel support 42
- FICON CTCA support 43
- FICON Express 42
- FICON Express2 42
- file pool, CMS
 - description 46
 - management using DFSMS/VM 57
 - security 36
- first-level storage
 - minimum and maximum 22
- FlashCopy support 43
- floating point support 43

G

- GCS (Group Control System) 53
- German, support for 32
- guest LAN 40
- guest operating system support
 - architecture compatibility 40
 - description 1
 - matrix 87

H

- hardware architectures and facilities, support for
 - adapter interruptions 45
 - concurrent-sense facility 41
 - Crypto Express2 42
 - cryptographic facilities 42
 - ECMDS 42
 - ESCON 42
 - expanded storage 42
 - Extended-TOD-Clock Facility 42
 - FCP 42
 - FICON channels 42
 - FICON CTCA 43
 - FlashCopy 43
 - HiperSockets 43

- hardware architectures and facilities, support for (*continued*)

- IEEE Floating Point 43
- integrated 3270 console 43
- Integrated Facility for Linux 43
- LCSS 43
- Move-Page Facility 44
- NPIV 44
- OSA 44
- Parallel Access Volumes 44
- PCICA 42
- PCICC 42
- PCIXCC 42
- PPRC 44
- PR/SM 45
- QDIO architecture 45
- Hardware Configuration Definition and Hardware Configuration Manager for z/VM
 - description 53
- Hardware Management Console 22
- hardware requirements
 - CSE 22
 - OSA/SF 22
 - server 21
 - TCP/IP for z/VM 23
 - z/VM installation 21
- HCD and HCM for z/VM
 - description 53
- HELP Facility 49
- High Level Assembler 25, 27, 28
- HiperSockets support 43
- Host Management Facilities/VM 4

I

- I/O (input/output)
 - channel program access protection 35
 - ESCON 42
 - FICON channels 42
 - FlashCopy 43
- IBM Publications Center 61
- IEEE Floating Point support 43
- IFL (Integrated Facility for Linux) support 43
- input/output (I/O)
 - channel program access protection 35
 - ESCON 42
 - FICON channels 42
 - FlashCopy 43
- installation requirements 21
- integrated 3270 console 43
- integrity
 - data 34
 - DirMaint 35
 - system 34
- IPv6 support 50
- ISFC (Inter-System Facility for Communications) 41

J

- Japanese, support for 32
- journaling 36

L

- Language Environment
 - description 54
 - program requirements 27
- LCSS (logical channel subsystems) 43
- library, z/VM
 - abstracts of books
 - base library - application programming 67
 - base library - customization 66
 - base library - end use 71
 - base library - installation, migration, and service 63
 - base library - operation 67
 - base library - planning and administration 64
 - base library - system diagnosis 73
 - base library - system overview 62
 - DFSMS/VM 74
 - DirMaint 75
 - OpenExtensions 68
 - Performance Toolkit for VM 75
 - RACF for z/VM 76
- basic (automatically shipped) documentation 62
- book editions and media 77

LibraryCenter 61

licensed programs supported on z/VM 31

limitations 30

Linux

- guest support for Linux on S/390 88
- guest support for Linux on System z9 and zSeries 88
- program requirements 30
- z/VM enhancements for guest support 9

Linux distributions supported by SSL server 12

logical channel subsystems (LCSS) 43

LookAt 61

LPAR (logical partition) 45

M

- main storage
 - minimum and maximum 22
- media, book 77
- memory
 - minimum and maximum 22
- Move-Page Facility support 44
- multitasking, CMS application 47

N

- N_Port ID Virtualization (NPIV) 15, 44
- national language support 32
- NetView 4
- Network Interface Card (NIC) 40
- networking
 - enhancements 13
- new or changed in z/VM V5.2
 - 60 LPARs 16
 - 64-bit addressing exploitation 9
 - application development and deployment 18
 - automation of local modification process 19

new or changed in z/VM V5.2 *(continued)*

- CD-ROM not available 20
- CMS binder upgrade 18
- Crypto Express2 Accelerator 14
- device support withdrawn 19
- directory management performance improvements 18
- DS4000 support 16
- dynamic LPAR naming 16
- guest support 9
- hardware and architecture support 14
- installation, migration, and service improvements 18
- LAN sniffer 13
- migration of preinstalled z/VM products 18
- MPRoute server 13
- N_Port ID Virtualization (NPIV) 15
- networking enhancements 13
- OSA-Express2 OSN 10
- packaging modifications 19
- performance assists for z/VM guests, enhanced 11
- SCSI disk I/O performance improvements 10
- security enhancements 13
- self-management 9
- simplified user administration 17
- SSL server enhancements 12
- statement of direction 20
- System z9 109 support 14
- systems management improvements 17
- technology exploitation 14
- Vector Facility support removed 19
- virtual image management enhancements 12
- virtualization technology 9
- withdrawn functions or support 19
- xyz 16

O

- object code only (OCO) 33
- online books
 - CD-ROM and DVD 61
 - program requirements 29
 - shipped with z/VM 62
- Open Systems Adapter Support Facility (OSA/SF)
 - description 54
 - hardware requirements 22
 - program requirements 28
- OpenExtensions
 - books, abstracts of 68
 - byte file system (BFS) 47
 - overview 47
 - program requirements 27
- operating systems supported as guests of z/VM 87
- ordering information 32
- OS/MVS simulation 46
- OSA (Open Systems Adapter) support 44
- OSA-2 54
- OSA-Express 54
- OSA-Express2 OSN 10
- OSA/SF (Open Systems Adapter Support Facility)
 - description 54
 - hardware requirements 22

OSA/SF (Open Systems Adapter Support Facility)
(*continued*)
program requirements 28

P

packaging modifications 19
packaging, product
distribution media 32
System Image contents 32
Parallel Access Volumes support 44
PCI Cryptographic Accelerator (PCICA), support for 42
PCI Cryptographic Coprocessor (PCICC), support for 42
PCIX Cryptographic Coprocessor (PCIXCC), support for 42
Peer-to-Peer Remote Copy (PPRC) support 44
Performance Toolkit for VM
book, abstract of 75
description 58
POSIX support 47
PR/SM (Processor Resource/Systems Manager)
support 45
preferred virtual machine 40
prerequisite licensed programs 25
prerequisites
block paging 22
customer responsibilities for data security 35
hardware
CSE 22
installing z/VM 21
OSA/SF 22
server 21
storage 22
TCP/IP for z/VM 23
installing z/VM 21
program
additional base facilities and optional features 26
CSE 26
DFSMS/VM 26
DirMaint 27
installing z/VM from DVD 22
Language Environment 27
Linux 30
online books 29
OpenExtensions 27
OSA/SF 28
prerequisite licensed programs 25
RACF for z/VM 28
tasks that require the High Level Assembler 25
TCP/IP for z/VM 28
VM-related products 30
server 21
storage requirements 22
virtual machine size, minimum 22
price/performance benefits 5
printers supported by z/VM 101
problems, diagnosing 52
Program Management Binder for CMS 48
program requirements
additional base facilities and optional features 26

program requirements (*continued*)
CSE 26
DFSMS/VM 26
DirMaint 27
installing z/VM from DVD 22
Language Environment 27
Linux 30
online books 29
OpenExtensions 27
OSA/SF 28
prerequisite licensed programs 25
RACF for z/VM 28
tasks that require the High Level Assembler 25
TCP/IP for z/VM 28
VM-related products 30
Programmable Operator 4
programming interfaces 3
programming languages supported 2
programs supported on z/VM 31
publications, z/VM
abstracts
base library - application programming 67
base library - customization 66
base library - end use 71
base library - installation, migration, and service 63
base library - operation 67
base library - planning and administration 64
base library - system diagnosis 73
base library - system overview 62
DFSMS/VM 74
DirMaint 75
OpenExtensions 68
Performance Toolkit for VM 75
RACF for z/VM 76
basic (automatically shipped) 62
editions 77
media 77
online
CD-ROM and DVD 61
program requirements 29

Q

QDIO (queued direct I/O) architecture support 45
queued direct I/O (QDIO) architecture support 45

R

RACF for z/VM (Resource Access Control Facility for z/VM) 37
books, abstract of 76
description 59
program requirements 28
real storage
minimum and maximum 22
requirements
block paging 22
customer responsibilities for data security 35
hardware
CSE 22

- requirements *(continued)*
 - hardware *(continued)*
 - installing z/VM 21
 - OSA/SF 22
 - server 21
 - storage 22
 - TCP/IP for z/VM 23
 - installing z/VM 21
 - program
 - additional base facilities and optional features 26
 - CSE 26
 - DFSMS/VM 26
 - DirMaint 27
 - installing z/VM from DVD 22
 - Language Environment 27
 - Linux 30
 - online books 29
 - OpenExtensions 27
 - OSA/SF 28
 - prerequisite licensed programs 25
 - RACF for z/VM 28
 - tasks that require the High Level Assembler 25
 - TCP/IP for z/VM 28
 - VM-related products 30
 - server 21
 - storage requirements 22
 - virtual machine size, minimum 22
- Resource Access Control Facility for VM (RACF for z/VM) 37
 - books, abstract of 76
 - description 59
 - program requirements 28
- restrictions
 - DirMaint 31
 - reconfiguration of partitionable processors 30
 - virtual machine 30
- Reusable Server Kernel 48
- REXX language 54
- REXX/VM 54
- RSA-Assist support 42

S

- SCSI disks, installing z/VM on 22
- security
 - DirMaint 37
 - file pool 36
 - general 35
 - TCP/IP for z/VM 37
- servers
 - requirements 21
 - supported by z/VM 83
- SFS (Shared File System) 46
- SNA (Systems Network Architecture)
 - AVS 52
 - GCS 53
 - ISFC 41
- softcopy books
 - CD-ROM and DVD 61
 - program requirements 29
 - shipped with z/VM 62

- software requirements
 - additional base facilities and optional features 26
 - CSE 26
 - DFSMS/VM 26
 - DirMaint 27
 - installing z/VM from DVD 22
 - Language Environment 27
 - Linux 30
 - online books 29
 - OpenExtensions 27
 - OSA/SF 28
 - prerequisite licensed programs 25
 - RACF for z/VM 28
 - tasks that require the High Level Assembler 25
 - TCP/IP for z/VM 28
 - VM-related products 30
- source code availability 33
- spanned HiperSockets channels 43
- statement of direction 20
- storage
 - management using DFSMS/VM 57
 - minimum and maximum 22
- storage controls supported by z/VM 99
- supported device, definition of 24
- switches supported by z/VM 103
- System Delivery Offering (SDO) 32
- system integrity
 - customer responsibilities for data security 35
 - DirMaint 35
 - z/VM, general statement for 34
- System z9 109 support 14
- system-managed storage 57
- System/370 architecture compatibility 39
- systems management APIs 48
- Systems Network Architecture (SNA)
 - AVS 52
 - GCS 53
 - ISFC 41
- SystemView Host Management Facilities/VM 4

T

- tape control units supported by z/VM 101
- tape libraries supported by z/VM 100
- tape units supported by z/VM 100
- TCP/IP for z/VM
 - hardware requirements 23
 - overview 49
 - program requirements 28
 - security 37
- terminals supported by z/VM 102
- Tivoli Storage Manager 20
- TPF, guest support for 88
- TSAF (Transparent Services Access Facility) 55

U

- unsupported device, definition of 24

V

- V=F virtual machine 40
- V=R virtual machine 40
- V=V virtual machine 40
- Vector Facility 19
- virtual machine
 - architectures 39
 - restrictions 30
 - type 40
- Virtual Machine Resource Manager (VMRM) 48
- virtual network adapters 40
- virtual switch 41
- VM Dump Tool 41
- VMRM (Virtual Machine Resource Manager) 48
- VSE/ESA, guest support for 88
- VTAM 53

W

- withdrawn functions or support 19

X

- XA virtual machine 39
- XC virtual machine 39, 40
- XEDIT 49
- xyz 16

Z

- z/Architecture
 - ESA virtual machine 40
- z/Architecture (64-bit addressing) exploitation 9
- z/OS, guest support for 88
- z/OS.e, guest support for 89
- z/TPF, guest support for 88
- z/VM Internet Library 61
- z/VM LibraryCenter 61
- z/VM V4.4
 - devices supported 91
 - guest operating systems supported 87
 - servers supported 83
- z/VM V5.1
 - devices supported 91
 - guest operating systems supported 87
 - servers supported 83
- z/VM V5.2
 - book editions and media 77
 - devices supported 91
 - distribution media 32
 - guest operating systems supported 87
 - hardware requirements
 - CSE 22
 - installation 21
 - OSA/SF 22
 - server 21
 - storage 22
 - TCP/IP for z/VM 23
 - languages supported 32
 - new functions, enhancements, and changes 9

z/VM V5.2 (*continued*)

- ordering information 32
- program requirements
 - additional base facilities and optional features 26
 - CSE 26
 - DFSMS/VM 26
 - DirMaint 27
 - installing from DVD 22
 - Language Environment 27
 - online books 29
 - OpenExtensions 27
 - OSA/SF 28
 - prerequisite licensed programs 25
 - RACF for z/VM 28
 - tasks that require the High Level Assembler 25
 - TCP/IP for z/VM 28
- programs supported 31
- servers supported 83
- source code availability 33
- System Image contents 32
- z/VM virtual switch 41
- z/VSE, guest support for 88

Readers' Comments — We'd Like to Hear from You

z/VM

built on IBM Virtualization Technology

General Information

version 5 release 2

Publication No. GC24-6095-03

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? ☐ Yes ☐ No

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.

Name

Address

Company or Organization

Phone No.



Cut or Fold
Along Line

Fold and Tape

Please do not staple

Fold and Tape




NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation
Department 55JA, Mail Station P384
2455 South Road
Poughkeepsie, New York 12601-5400



Fold and Tape

Please do not staple

Fold and Tape

Cut or Fold
Along Line



Program Number: 5741-A05

Printed in USA

GC24-6095-03



Spine information:



z/VM

General Information

version 5 release 2