

z/VM



General Information

Version 6 Release 4

Note:

Before you use this information and the product it supports, read the information in “Notices” on page 117.

This edition applies to version 6, release 4, modification 0 of IBM z/VM (product number 5741-A07) and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces GC24-6193-11.

© **Copyright IBM Corporation 1990, 2016.**

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

Contents

Tables	vii
-------------------------	------------

About this document	ix
--------------------------------------	-----------

Intended audience	ix
-----------------------------	----

Where to find more information	ix
--	----

Links to other documents and websites	ix
---	----

How to send your comments to IBM	xi
---	-----------

Chapter 1. z/VM overview	1
---	----------

z/VM components, facilities, and optional features	1
--	---

z/VM provides proven system integrity, security, availability, and reliability	3
--	---

z/VM supports application development and deployment	4
--	---

z/VM is accessible by people with disabilities	5
--	---

Chapter 2. What is new or changed in z/VM V6.4	7
---	----------

Hypervisor efficiency and scalability	7
---	---

HyperPAV technology exploitation	7
--	---

Guest large page support	7
------------------------------------	---

Guest Transactional Execution support.	8
--	---

Enhanced real memory and guest virtual memory support	8
---	---

Memory scalability efficiency improvements.	8
---	---

Surplus CPU power distribution improvement	8
--	---

Support for large HCD Dynamic I/O Activates.	8
--	---

System ease of use improvements	8
---	---

Dynamic simultaneous multithreading level	9
---	---

Control Program environment variables	9
---	---

QUERY SHUTDOWN command.	9
---------------------------------	---

SCSI management queries	9
-----------------------------------	---

SCSI reliability, availability, and serviceability enhancements	9
---	---

CMS Pipelines	10
-------------------------	----

DirMaint RACF connector	10
-----------------------------------	----

RACF automatic control of access list authority	10
---	----

Performance Toolkit enhancements	10
--	----

IBM Wave for z/VM (IBM Wave) optional feature	10
---	----

z/OS V2.2 equivalency	11
---------------------------------	----

Hardware currency and exploitation	12
--	----

IBM z13 (Driver D27) and z13s I/O support	12
---	----

IBM Dynamic Partition Manager (DPM) administrative mode for Linux extended to z/VM V6.4	13
---	----

Installation, migration, and serviceability	13
---	----

Enhanced upgrade in place	13
-------------------------------------	----

Determine installed service	14
---------------------------------------	----

3590 and 3592 tape not supported for the installation of z/VM V6.4	14
--	----

Additional enhancements and improvements	14
--	----

IBM OpenStack for z Systems	16
---------------------------------------	----

Statements of general direction	17
---	----

Chapter 3. Technical information	19
---	-----------

Hardware and storage requirements	19
---	----

Server requirements	19
-------------------------------	----

Installation requirements	20
-------------------------------------	----

Storage requirements	20
--------------------------------	----

Workstation requirement	21
-----------------------------------	----

IBM Wave hardware requirements.	21
---	----

OSA/SF hardware requirements	21
Performance Toolkit hardware requirements	21
SSI cluster hardware requirements.	22
TCP/IP hardware requirements.	22
VM Guest RSA-Assist Support for Linux hardware requirements	22
zHPF guest support hardware requirements	23
zHPF paging hardware requirements	23
PCIe support hardware requirements.	23
Multi-VSwitch Link Aggregation support hardware requirements	23
Device support	23
Supported devices	23
Unsupported devices	24
Software (program) requirements	24
Prerequisite licensed programs	25
High Level Assembler	25
z/VM base components additional program requirements.	26
Supported levels of z/VM base facilities and optional features	26
DFSMS/VM additional program requirements.	27
DirMaint additional program requirements	28
HCD and HCM additional program requirements	28
IBM Wave program requirements	28
IOCP additional program requirements	28
Language Environment additional program requirements	28
OSA/SF additional program requirements	28
Performance Toolkit additional program requirements	29
RACF additional program requirements	29
RSCS additional program requirements	29
Shared-DASD complex additional program requirements	29
SSI clusters additional program requirements	29
TCP/IP additional program requirements	30
z/VM OpenStack support additional program requirements	30
Program requirements for other products running on z/VM	30
Limitations	30
Performance considerations	31
Operating systems supported as guests	31
Other programs supported on z/VM	31
Multicultural support	31
Packaging and ordering information	32
Object code only and limited source availability	33
Integrity and security	33
Data integrity for guests	33
System integrity statement for z/VM.	34
Security, auditability, and control	35
Chapter 4. z/VM base product	39
Control Program (CP)	39
Virtual machine modes	39
Architecture compatibility and support	40
Guest LANs and virtual switches	40
z/VM HiperDispatch	41
Inter-System Facility for Communications	41
VM Dump Tool	41
Support for hardware facilities, features, and architectures.	42
Conversational Monitor System (CMS)	51
Versions of CMS.	51
Shared File System	52
CMS application programming	52
Systems management application programming interface	53
CMS Pipelines	53
CMS application multitasking	53
OpenExtensions	53

Program Management Binder for CMS	54
Reusable Server Kernel	54
Virtual Machine Resource Manager	55
XEDIT	55
HELP facility	55
TCP/IP for z/VM	56
Link protocols	56
Network protocols	56
Transport protocols	57
Applications and protocols	57
TCP/IP requirements and additional information.	58
APPC/VM VTAM Support (AVS)	58
Dump Viewing Facility	58
Group Control System (GCS)	59
Hardware Configuration Definition (HCD) and Hardware Configuration Manager (HCM) for z/VM	59
Language Environment	60
Open Systems Adapter Support Facility (OSA/SF)	60
REXX/VM.	60
Transparent Services Access Facility (TSAF).	61
Virtual Machine Serviceability Enhancements Staged/Extended (VMSES/E).	61
Chapter 5. z/VM optional features	63
Data Facility Storage Management Subsystem for VM (DFSMS/VM)	63
Directory Maintenance Facility for z/VM (DirMaint)	64
IBM Wave for z/VM	65
Performance Toolkit for VM	65
RACF Security Server for z/VM	66
RSCS Networking for z/VM.	67
z/VM Single System Image Feature	68
Chapter 6. z/VM library guide	69
Where to get information about z/VM	69
Documentation supplied with z/VM	69
IBM Knowledge Center	69
IBM Publications Center	69
z/VM Collection	69
z/VM Internet Library.	69
Abstracts of the z/VM base product publications.	70
System overview	70
Installation, migration, and service	70
Planning and administration.	71
Customization and tuning	74
Operation and use	75
Application programming	76
Diagnosis	81
Abstracts of the z/VM optional feature publications.	82
DFSMS/VM publications	82
Directory Maintenance Facility for z/VM publications	83
IBM Wave for z/VM publications	83
Performance Toolkit for VM publications	84
RACF Security Server for z/VM publications	84
RSCS Networking for z/VM publications	85
Appendix A. IBM servers supported by z/VM	87
Appendix B. IBM operating systems supported as guests of z/VM	93
Appendix C. IBM devices supported by z/VM	97
Direct access storage devices (DASD).	98
DASD control units and storage controls	102

Tape units and tape libraries	102
Tape control units	104
Printers	104
Card readers and card punches	105
Terminals, displays, and consoles.	105
Display printers	105
Display control units	105
Communications controllers	106
Switches	106
Server features and miscellaneous devices.	106
Notices	117
Trademarks	119
Terms and conditions for product documentation	119
IBM online privacy statement	120
Glossary	121
Index	123

Tables

1.	IBM-translated parts of z/VM	32
2.	Server support	88
3.	Guest support.	93
4.	Device support – DASD	98
5.	Device support – DASD control units and storage controls	102
6.	Device support – tape units and tape libraries	102
7.	Device support – tape control units	104
8.	Device support – printers	104
9.	Device support – card readers and card punches.	105
10.	Device support – display control units	105
11.	Device support – communications controllers.	106
12.	Device support – switches	106
13.	Device support – server features and miscellaneous devices	106

About this document

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

- Product overview
- What is new or changed in z/VM V6.4
- Hardware and software requirements
- Guide to the product information
- IBM servers, guest operating systems, and devices supported by current z/VM releases

Intended audience

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

Where to find more information

You can obtain more information about z/VM from the documents described in Chapter 6, “z/VM library guide,” on page 69.

Links to other documents and websites

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

How to send your comments to IBM

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

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

1. Send an email to mhvrcfs@us.ibm.com.
2. Go to IBM z/VM Reader's Comments (www.ibm.com/systems/z/os/zvm/zvmforms/webqs.html).

Include the following information:

- Your name
- Your email address
- The publication title and order number:
 z/VM V6.4 General Information
 GC24-6193-12
- The topic name or page number related to your comment
- The text of your comment

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

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

If you have a technical problem

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

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

Chapter 1. z/VM overview

z/VM V6.4 virtualization technology is designed to run hundreds to thousands of Linux servers on a single IBM z Systems™ or LinuxONE server with the highest degrees of efficiency and elasticity.

The ability of z/VM to support multiple machine images and architectures provides a highly flexible production and test environment for z Systems™ and LinuxONE operating systems to:

- Simplify migration from one release to another
- Facilitate the transition to newer applications
- Provide a test environment whenever one is needed
- Consolidate many systems onto one physical server

A fundamental strength of z/VM is the ability for virtual machines to share system resources with very high levels of resource utilization. z/VM V6.4 provides even greater levels of extreme scalability, security, and efficiency to create opportunities for cost savings, while providing a robust foundation for cognitive computing on z Systems and LinuxONE servers

z/VM V6.4 provides support for the IBM z13™, IBM z13s, IBM LinuxONE Rockhopper™, IBM LinuxONE Emperor™, and IBM zEnterprise® (196, 114, EC12, BC12) servers, as well as Red Hat, SUSE, and Ubuntu Linux distributions. Support for simultaneous multithreading (SMT) technology extends per-processor, core capacity growth beyond single-thread performance for Linux on z Systems running on an IBM Integrated Facility for Linux (IFL) specialty engine on a z13™, z13s, or LinuxONE server. z/VM multithreading technology support can enable additional price/performance benefits over previous hardware generations and can meet workload requirements transparently. Improvements made in the areas of reliability, availability, and serviceability allow low-end devices such as IBM Storwize® V7000, V840, and V9000 to be attached to a z/VM host, removing the need for a SAN Volume Controller.

z/VM V6.4 is a supported environment using IBM Dynamic Partition Manager for Linux-only systems with SCSI storage. This simplifies system administration tasks for a more positive experience by those with limited mainframe skills. IBM Wave Version 1 Release 2 (V1.2) is now included in z/VM V6.4 as a priced feature. It can greatly simplify the task of administering a z/VM environment.

z/VM V6.4 can help you extend the business value of IBM z Systems and IBM LinuxONE™ technology across the enterprise by integrating applications and data, while providing exceptional levels of availability, security, and operational ease. World-class virtualization technology offered by z/VM can provide the ability to host a large number of virtual servers running different operating systems on a z Systems server and LinuxONE.

z/VM components, facilities, and optional features

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

Control Program (CP)

CP is a hypervisor and real-machine resource manager.

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

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

Conversational Monitor System (CMS)

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

Dump Viewing Facility

Dump Viewing Facility is a tool that analyzes and manages system software problems interactively under CMS. (CP problems are handled by the VM Dump Tool function of CP.)

Group Control System (GCS)

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

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

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

Language Environment® for z/VM

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

Open Systems Adapter/Support Facility (OSA/SF)

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

REXX/VM

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

Transparent Services Access Facility (TSAF)

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

TCP/IP for z/VM

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

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

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

For more information about these components and facilities, see Chapter 4, “z/VM base product,” on page 39.

z/VM also offers the following optional features:

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

DFSMS/VM controls data and storage resources and provides an interface for tape librarian products to an Automated Tape Library.

Directory Maintenance Facility for z/VM (DirMaint™)

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

IBM Wave for z/VM (IBM Wave)

IBM Wave is a provisioning and productivity management solution for simplifying the control and use of virtual Linux servers and z/VM.

IBM z/VM Single System Image Feature (VMSSI)

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

Performance Toolkit for VM

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

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

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

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

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

For more information about these optional features, see Chapter 5, “z/VM optional features,” on page 63.

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

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

- Integrity and security:
 - z/VM supports guest use of the cryptographic facilities provided by supported IBM servers.
 - z/VM supports drive-based data encryption with IBM System Storage® tape drives. z/VM supports native and guest use of the IBM Encryption Key Manager for encryption of tapes.
 - z/VM supports the use of DASD volumes that reside on data encryption drives (DEDs).
 - Support for Transport Layer Security (TLS) and Secure Sockets Layer (SSL) is provided through TCP/IP for z/VM. The z/VM System SSL cryptographic module has been validated in accordance with Federal Information Processing Standard (FIPS) 140-2. The z/VM TLS/SSL server supports TLS 1.2, SHA-2 certificates, and NIST SP 800-131a compliant configurations.
 - Integrated access control and authentication services can be augmented with the addition of an external security manager (ESM), such as the RACF Security Server for z/VM. RACF can also be used to audit connections to z/VM real devices.
 - z/VM V6.3, with the SSI and RACF Security Server features enabled, has been certified to conform to the Operating System Protection Profile (OSPP) with Virtualization (-VIRT) and Labeled Security (-LS) extensions of the Common Criteria standard for IT security, ISO/IEC 15408, at Evaluation Assurance Level 4 (EAL4+).

z/VM V6.4 is designed to meet the same Common Criteria certification as achieved for z/VM V6.3. For more information, see IBM z/VM Security and Integrity resources (www.ibm.com/vm/security/).

- RSCS TCPNJE traffic can be encrypted by directing the flow through a TLS/SSL server.

IBM will correct any integrity exposures introduced by unauthorized programs into the system.

- Availability and reliability:
 - Application recovery: z/VM provides services which permit recovery of incomplete interactions with resource managers.
 - Automated operations: z/VM offers several levels of automated system management support. One example is the Programmable Operator. For a higher degree of automation, IBM Operations Manager for z/VM can be added.
 - z/VM provides duplexed data with transparent ongoing synchronization between the primary and backup copy, and automatic transparent switching to the backup copy in case of an error in the primary copy.
 - Online configuration changes eliminate many previously required outages.
 - z/VM systems can be connected for improved server and user availability.
 - Fast restart reduces the end user impact of any outage.
 - Setting up a z/VM SSI cluster and using live guest relocation allows z/VM and hardware maintenance to be less disruptive to workloads and allows less disruptive workload balancing.

z/VM supports application development and deployment

Applications for the CMS environment can be written in a variety of languages, including:

- APL2®
- Assembler
- C and C++
- COBOL
- FORTRAN
- Pascal
- PL/I
- REXX

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:

- Certain CP and CMS control blocks
- CP and CMS monitor records
- CP assembler macros (such as IUCV and APPCVM)
- CP DIAGNOSE instructions
- CP system services (such as *MSG)
- Data record formats, such as accounting records, intended to be processed by application programs

- GCS assembler macros
- System and user exits
- TCP/IP APIs (such as C sockets, IUCV sockets, Remote Procedure Calls (RPCs), and monitor records)

Special facilities are available to CMS applications, such as:

- Assembler macros and functions
- Callable services library (CSL)
- CMS multitasking services
- CMS Pipelines
- DOS/VSE support
- Enterprise Systems Architecture/Extended Configuration (ESA/XC) support
- Language Environment
- MQSeries® Client library
- OpenExtensions (POSIX interfaces and shell)
- OS/MVS simulation
- Program Management Binder
- Reusable Server Kernel
- REXX sockets API
- SAA Common Programming Interface - Communications (CPI-C) library
- Systems management API

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

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

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

z/VM is accessible by people with disabilities

The following features support use by people with disabilities:

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

Chapter 2. What is new or changed in z/VM V6.4

The following topics describe the new functions, enhancements, and other changes included in z/VM V6.4:

- “Hypervisor efficiency and scalability”
- “System ease of use improvements” on page 8
- “Hardware currency and exploitation” on page 12
- “Installation, migration, and serviceability” on page 13
- “Additional enhancements and improvements” on page 14
- “IBM OpenStack for z Systems” on page 16

Notes:

1. Some product updates might require additional hardware or software support. See Chapter 3, “Technical information,” on page 19.
2. Some product updates might be available for previous z/VM releases. See IBM: z/VM Service Resources (www.ibm.com/vm/service/).

For planned future changes to z/VM, see “Statements of general direction” on page 17.

For additional information about the changes to external interfaces in z/VM V6.4 and previous releases, see *z/VM: Migration Guide*.

Hypervisor efficiency and scalability

The overall efficiency of the z/VM hypervisor has been enhanced with the following scalability improvements.

HyperPAV technology exploitation

z/VM exploits the ability for an IBM DS8000® device to issue concurrent I/O requests to an ECKD™ paging volume. In HyperPAV mode, if the base volume is busy, z/VM selects a free alias device from a pool, binds the alias to the base device, and starts the I/O. When the I/O completes, the alias device is returned to the pool to be used for another I/O to the same logical subsystem (LSS).

The primary benefit of exploiting HyperPAV is to improve paging throughput during periods of high-volume disk I/O, which will increase the efficiency of the z/VM frame replenishment algorithm used to manage storage overcommitted workloads. HyperPAV paging also enables the management of fewer and larger CPOWNEED volumes.

HyperPAV paging is exploited by the z/VM hypervisor for:

- The SYSRES volume, and volumes containing checkpoint and warm start data
- Volumes used for paging, spooling, and the z/VM user directory
- Minidisk pools, as defined by a guest's use of the MAPMDISK IDENTIFY macro

Guest large page support

z/VM provides support for the enhanced-DAT facility, which allows a guest to exploit large (1 MB) pages. A larger page size decreases the amount of guest

New or Changed

memory needed for dynamic address translation (DAT) tables and also decreases the overhead required to perform address translation. In all cases, guest memory is mapped into 4 KB pages at the host level.

With guest large page support, Linux on z Systems, z/VSE®, and z/OS® virtual machines can benefit from reduced memory footprints and address translation times. This can decrease overhead and improve throughput.

Guest Transactional Execution support

z/VM supports guest exploitation of the Transactional Execution (TX) facility on supported machines. The TX facility allows a program to issue multiple instructions that appear to operate atomically, offering an alternative to more costly mutual-exclusion mechanisms such as software locks. This support can improve the efficiency and scalability of multithreaded software such as Java™ or guest operating system functions.

Enhanced real memory and guest virtual memory support

The maximum amount of real storage (memory) that z/VM exploits is increased from 1 TB to 2 TB. The maximum supported virtual memory for a single guest remains at 1 TB. When keeping the virtual to real overcommitment the same, doubling the real memory used results in doubling the active virtual memory that can be used effectively. This virtual memory can come from an increased number of virtual machines and/or larger virtual machines. This allows for greater leverage of white space.

Memory scalability efficiency improvements

Improvements to memory management algorithms provide a basis for future enhancements that can increase the performance of workloads that experience available list spin lock contention.

Surplus CPU power distribution improvement

Virtual machines that do not consume all of their entitled CPU power, as determined by their share setting, generate “surplus CPU power”. This enhancement distributes the surplus to other virtual machines in proportion to their share settings. This is managed independently for each processor type (General Purpose, IFL, zIIP, and so on) across the virtual machines.

Support for large HCD Dynamic I/O Activates

CP no longer relies on obtaining a large contiguous area of memory to handle Dynamic I/O Activate requests from HCD. In most cases this eliminates the risk of receiving an HCD message that the configuration change is too large for CP to handle. In addition, the default size of the VM HCD virtual machine has been increased to 2 GB, the maximum allowed for an XC mode virtual machine. This gives HCD the maximum amount of virtual memory possible for use in handling large Dynamic I/O Activates.

System ease of use improvements

Systems administrators can experience more effective management of a z/VM system by taking advantage of the following enhancements.

Dynamic simultaneous multithreading level

Support for simultaneous multithreading (SMT) is enhanced with the addition of the SET MULTITHREAD command. Once z/VM V6.4 has been IPLed with multithreading enabled in the system configuration file, this command can be used to nondisruptively switch between one and two activated threads per IFL core. Performance of a system and workload with one active thread per core is comparable to that of the same system and workload with multithreading disabled. Thus, the dynamic SMT level capability allows the benefit of multithreading to be evaluated for a workload without requiring an outage to enable or disable SMT.

The SET MULTITHREAD command is allowed only when the system has been enabled for multithreading in the system configuration file, which can specify activating either one or two threads per core. It is not possible to revert to a non-SMT configuration without an IPL. SMT-enabled configurations are restricted to 32 cores even when operating in single-threaded mode due to the logical processor addressing limit.

This satisfies the statement of direction made in the IBM Software Announcement dated February 16, 2016.

Control Program environment variables

This enhancement allows automation procedures to adapt more easily to changes in operating environments to help simplify the control and testing of a system setup. For example, an operator can indicate at IPL time that the system is running in a disaster recovery or test environment, which in turn enables automation routines to modify the devices used and alter the choice and sequence in which virtual machines are activated, as well as perform other environment-dependent functions.

QUERY SHUTDOWN command

The new QUERY SHUTDOWN command enables a z/VM system programmer or a guest virtual machine to determine whether a system shutdown is in progress and obtain additional information about the shutdown. This can help automate an orderly shutdown of the z/VM system and its virtual servers. This function can be of particular value to virtual machines that coordinate the shutdown of other virtual machines. The coordinating virtual machines would receive the signal that the system is shutting down, issue the new QUERY command to get additional information, and take the appropriate action for an orderly shutdown.

SCSI management queries

This support provides enhancements to the commands for EDEVICES within z/VM. It improves the usability and problem diagnosis for EDEV-intensive environments and provides a clearer end-to-end view of the storage configuration. This simplifies the process of verifying that the storage configuration is consistent between z/VM and the disk storage subsystem.

SCSI reliability, availability, and serviceability enhancements

The following updates are designed to further enhance the reliability of SCSI devices:

- The CP missing interrupt handler is disabled for EDEVICES, allowing the SCSI driver to manage its outstanding requests in a more appropriate manner.
- The SCSI driver is updated to provide additional path recovery.

- Internal timer values are changed to enhance the interoperability between the SCSI driver and SAN Volume Controller (SVC) as well as between devices incorporating SVC technology.
- Debug facilities within the SCSI driver are enhanced, allowing IBM support teams to more quickly diagnose and debug issues in the field.
- Guidelines for multipath configuration are provided for SVC and devices incorporating SVC technology to ensure that path recovery is optimal.
- Concurrent code loads on the SVC and devices incorporating SVC technology are now supported without quiescing EDEVICE I/O. This was previously restricted, and continues to be restricted on releases prior to z/VM V6.4.

CMS Pipelines

CMS Pipelines functionality is modernized by adopting 20 years of development since the original Pipelines integration into z/VM. This upgrade addresses client concerns with using downloaded code, includes fixes not previously integrated into the z/VM product, broadens the ecosystem, enables innovation for clients and ISVs, and includes additional functionality.

The new *z/VM: CMS Pipelines User's Guide and Reference* is based on *CMS/TSO Pipelines: Author's Edition*, SL26-0018, and replaces both of the previous z/VM CMS Pipelines publications.

DirMaint RACF connector

The DirMaint RACF connector is modernized with a collection of functional enhancements that improve how z/VM security is handled in a managed environment. The connector allows appropriate security policy changes to be passed directly to RACF. This allows a z/VM environment managed by IBM Wave or within an OpenStack environment to function properly with RACF installed on the system.

RACF automatic control of access list authority

The ADDCREATOR and NOADDCREATOR options on the RACF SETROPTS command determine whether or not the creator of a RACF profile is automatically added to its access control list. This enhancement removes the need for manual intervention in RACF resource configuration and eliminates a point of potential human error from security policy management.

Performance Toolkit enhancements

Performance Toolkit for VM, function level 640, is enhanced to exploit z/Architecture® and its expanded set of instructions. Consequently, the PERFSVM virtual machine must run on z/Architecture CMS (z/CMS).

IBM Wave for z/VM (IBM Wave) optional feature

IBM Wave for z/VM (IBM Wave) is packaged with z/VM V6.4 as an optional feature but is not preinstalled and can be ordered through the System Delivery Offering (SDO).

IBM Wave is a provisioning and productivity administrative solution that can significantly reduce the learning curve for administering and controlling z/VM and its Linux guests. At the same time, it can also increase the productivity of those who have a deep understanding of z/VM environments. IBM Wave intelligent visualization of the virtual server environment and physical infrastructure provides intuitive administration of physical servers, z/VM, Linux

images running as z/VM guests, and other resources. With IBM Wave, you can rapidly gain insight into your entire virtualized infrastructure topology at a glance and also accelerate the path to a highly virtualized cloud infrastructure.

IBM Wave Release 2 further expands the capabilities by delivering increased support for Linux distributions and devices, reporting and auditing, as well as additional enterprise-grade security and performance enhancements.

The following enhancements have been made available with service updates:

- Enhanced logging allows an IBM Wave administrator to satisfy corporate auditing needs by accessing a consolidated log of specified activities.
- Installation and operation of the IBM Wave server (IBM WAVESRV) are supported on Red Hat RHEL 7 and SUSE SLES 12 distributions.
- Dynamic CPU management enables virtual CPUs to be added dynamically to a guest's configuration, either temporarily or permanently, without requiring an outage that could disrupt their business.
- Canonical Ubuntu on IBM z Systems and LinuxONE is supported as a guest operating system that IBM Wave can provision and manage, giving clients greater freedom of choice in Linux distributions. This satisfies the statement of direction made in the IBM Software Announcement dated February 16, 2016.
- Support for z/VM single system image is enhanced to automatically populate network configuration information and changes across the SSI cluster to help simplify administration.
- Performance is improved and the user interface is simplified through the use of more consistent terminology to facilitate the use of the Automatic Guest Classification facility to recognize guests for grouping purposes based on their characteristics.
- Dynamic memory management enables virtual memory to be added dynamically, either temporarily or permanently, to a guest's configuration, without requiring a disruptive outage.
- Disk storage management is improved to enable better planning for growth by allowing existing storage resources to be used to increase the size of a volume group rather than assuming that new resources must first be allocated.
- Security is improved by eliminating the need for IBM Wave service machines to have logon passwords, using LOGONBY capabilities, allowing authorized users to log on using their own passwords.
- Performance is improved by reducing overhead with IBM Wave detecting when a managed guest is idle and avoiding waking it up unnecessarily.

z/OS V2.2 equivalency

The following z/VM facilities are upgraded to a level equivalent to z/OS V2.2.

- **LDAP server and client utilities**

The updated z/VM support includes:

- Paged and sorted search results
- Group search limits
- 64-bit DB2® support
- SHA-2 and Salted SHA-2
- Listening on INADDR_ANY or in6addr_any interfaces
- Administrative group and roles

- **MPROUTE server**

The updated z/VM support includes:

- Deprecation of the OMPROUTE_OPTIONS=hello_hi environment variable
- Processing inbound OSPF hello packets from neighbors at the highest priority, for the purpose of maintaining OSPF adjacencies
- Modifications to avoid abends when formatting or parsing OSPF packet content
- Enhancements to existing informational and debug messages, to provide more specific information when an IOCTL call fails

- **System SSL and utilities**

The updated z/VM support includes:

- Support for the current level of the System SSL and LDAP servers
- GSKKMAN support for RFC 5280 certificates

- **Program Management Binder**

The updated z/VM support includes:

- Generalized alignment of the text elements and parts of a program
- A new SYMTRACE option to provide new binder messages displaying a trace of progress of binder symbol resolution for a specific symbol
- Updated support for Assembler J-cons
- Various RAS items

Hardware currency and exploitation

z/VM V6.4 incorporates the following improvements in hardware and architecture support.

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

IBM z13 (Driver D27) and z13s I/O support

z/VM supports the following I/O architecture enhancements on the IBM z13 (Driver D27) and z13s:

- **UID (user defined identifier) support for real PCI functions**

With the PTF for APAR VM65865, a UID can be assigned to a real PCI function to be able to more accurately indicate equivalent functions between different LPARs and for exploitation by guest operating systems, Linux in particular. z/VM supports dynamic I/O and guest use for the new PCIe UID support.

Additional program requirements for UID support:

- HCD support on V6.2, V6.3, and V6.4 requires APAR VM65827.
- HCM support on V6.2, V6.3, and V6.4 requires APAR VM65201.
- IOCP support on V6.2 and V6.3 requires APAR VM65817. (IOCP support is included in the V6.4 base.)

- **Shared Memory Communications - Direct Memory Access (SMC-D)**

This technology is the latest networking innovation for the IBM z13 family of processors. It provides support for fast, low-latency LPAR-to-LPAR TCP/IP traffic using the SMC-D software protocol over firmware-provided Internal Shared Memory (ISM) devices.

Supported for z/VM guest exploitation, SMC-D and ISM are designed to use shared memory areas to provide low-latency, high-bandwidth, cross-LPAR connections for applications. This support is intended to provide application-transparent DMA communications to TCP endpoints for sockets-based connections. SMC-D is expected to provide substantial performance, throughput, response time, and CPU consumption benefits compared with standard TCP/IP communications over HiperSockets™. z/VM supports dynamic I/O and guest use of the new Internal Shared Memory (ISM) PCI function type.

See the z/OS subset within the 2964DEVICE or 2965DEVICE PSP bucket for z/OS service required in support of SMC-D connectivity.

IBM Dynamic Partition Manager (DPM) administrative mode for Linux extended to z/VM V6.4

Dynamic Partition Manager (DPM), provided with IBM z13 and LinuxONE servers, is extended to support Linux running on z/VM V6.4 with SCSI storage attached with FCP channels. DPM does not support ECKD-attached DASD devices. DPM provides simplified hardware and virtual infrastructure management including integrated dynamic I/O management.

DPM provides simplified, consumable, and enhanced partition lifecycle and dynamic I/O management capabilities via the Hardware Management Console (HMC) to:

- Create and provision an environment, including the creation of new partitions, assignment of processors and memory, and configuration of I/O adapters (network, FCP storage, crypto, and accelerators)
- Manage the environment by modifying system resources without disrupting running workloads
- Monitor and troubleshoot the environment to identify the source of system failures, conditions, states, or events that may lead to workload degradation

A CPC can be configured in either DPM mode or PR/SM™ mode. The mode is enabled prior to the CPC power-on reset (POR). DPM mode requires two OSA-Express 1000BASE-T Ethernet features for primary and backup connectivity (OSA-Express5S 1000BASE-T Ethernet #0417), along with associated cabling (HW for DPM #0016).

Versions of z/VM prior to V6.4 are not supported. There is no FICON® support in DPM at this time. Therefore, single system image (SSI) clusters are not supported.

Installation, migration, and serviceability

z/VM V6.4 includes enhanced functions and processes to improve the installation, migration, and serviceability of z/VM.

Enhanced upgrade in place

The z/VM upgrade in place process allows upgrading an existing system to a new release of z/VM with minimal impact to the running system. The upgrade in place process has been extended to allow an upgrade from z/VM V6.2 or V6.3 to z/VM V6.4 and positions a system for upgrading to releases beyond z/VM V6.4. Upgrade in place is supported for a member of a z/VM SSI cluster as well as for a nonclustered z/VM system.

Determine installed service

Enhancements to CP and VMSES/E enable you to determine if specific CP service is built into the CP nucleus of a running system. The new CPSERVICE option on the CP QUERY command allows queries based on APAR, PTF, or local modification identifiers.

3590 and 3592 tape not supported for the installation of z/VM V6.4

z/VM V6.4 and the z/VM SDO Licensed Products are not available in 3590 or 3592 tape format. z/VM and the z/VM SDO Licensed Products are available on DVD and via electronic delivery. In addition, z/VM service for all releases will no longer be orderable in 3590 or 3592 tape format. For the current list of z/VM SDO licensed products, see IBM: SDO - System Delivery Offering for IBM z/VM (www.ibm.com/vm/sdo/). This satisfies the statement of direction made in the IBM Hardware Announcement dated January 14, 2015.

Additional enhancements and improvements

z/VM V6.4 includes the following additional enhancements and improvements.

- **Domain Name System (DNS) IPv6 support**
z/VM TCP/IP supports DNS servers that have IPv6 addresses.
- **Encryption of TCPNJE connections**
RSCS TCPNJE traffic can be encrypted by directing the flow through an SSL server. The secure TCP/IP protocols that were previously implemented to support VMCF clients and servers are extended to IUCV clients and servers. Traffic that was previously unencrypted and potentially exposed to sniffers on a network can now be protected. This is especially important for TCPNJE environments, where NJE jobs can include user IDs and passwords.
- **Nondefault tape mode support**
This enhancement provides support for nondefault write formats (modes) on the IBM System Storage 3592 TS1120 (E05), TS1130 (E06), and TS1140 (E07) Tape Drives. This extends the use and interoperability of cartridges between different tape hardware subsystems.
- **Resetting counters for a virtual switch**
This function provides the ability to clear certain counts that are displayed via the QUERY VSWITCH DETAILS command while the virtual switch remains active. The counts that can be cleared are: data, discarded, and errors. These counts can be cleared for the UPLINK port or the BRIDGEPORT.
- **SET DIALDROP command**
The new SET DIALDROP command establishes whether devices dialed to your virtual machine are dropped or stay connected when an implicit or explicit virtual machine reset occurs.
- **LOGON security improvement**
The CP LOGON command has been updated to prevent user ID enumeration without authenticating to the system. This means z/VM will no longer give an error message that indicates a user ID is valid and the password is not. This prevents unauthorized users from determining what might be a valid user ID. Also, many existing error messages will be presented only after a valid user ID and password or password phrase are provided.
- **Default TLS protocol settings changed**

When using the TLS/SSL Server to encrypt TCP/IP traffic, TLS 1.2 and TLS 1.1 are enabled by default. Older versions of the TLS and SSL protocols are disabled by default and require a configuration action to enable. Note that older protocols do not meet modern security requirements and are provided for legacy use only.

- **VLAN access security improvement**

When using an External Security Manager (ESM) to control a z/VM virtual switch, user access to the default VLAN ID is not permitted unless permission has been explicitly granted through appropriate ESM controls.

- **zManager support removed**

The IBM z[™] Unified Resource Manager (zManager) is no longer supported by z/VM. The virtual switch types of IEDN and INMN have been removed from CP and TCP/IP commands and other externals.

- **SMAPI support**

The Systems Management API (SMAPI) interface has been updated with support for z/VM V6.4. These updates include support for specific new functionality, removal of deprecated function, and some long desired minor functional updates and implementation changes.

- Support for SET MACHINE Z
- Changes to address some security concerns
- Removal of Ensemble data from and addition of some new information to SMSTATUS
- Removal of support for Ensembles
- Removal of code specific to the RPC server, last supported in z/VM 5.4

- **Expanded storage (XSTORE) support removed**

IBM z13 and z13s are the last z Systems servers to support expanded storage (XSTORE). z/VM V6.4 does not support XSTORE for either host or guest usage. This satisfies the statements of direction made in the IBM Hardware Announcement dated January 14, 2015, and the IBM Software Announcements dated July 23, 2013, and January 14, 2015.

- **Architecture level set (ALS)**

z/VM V6.4 requires a new architecture level set (ALS) and supports only IBM zEnterprise 196 (z196) and IBM zEnterprise 114 (z114) and later servers. See the appropriate Preventive Service Planning (PSP) bucket for the minimum microcode level (MCL) and any required updates. See Appendix A, “IBM servers supported by z/VM,” on page 87. This satisfies the statement of directory made in the IBM Software Announcement dated February 16, 2016.

- **ESA/390 removal**

z/VM V6.4 enhancements enable hypervisor initialization and termination, the Stand-Alone Program Loader (SAPL), DASD Dump Restore (DDR), Stand-Alone Dump, and other stand-alone utilities to run entirely in z/Architecture mode. The IBM z13 and z13s are planned to be the last z Systems servers to support running an operating system in ESA/390 architecture mode. All future systems will support only operating systems that run entirely in z/Architecture mode.

Execution of architecture-conformant applications on CMS and GCS in ESA/390 architecture mode remains supported, as long as the application does not depend on more advanced functions such as dynamic address translation.

In addition, support has been added to z/VM to simulate a z/Architecture-only environment, by providing a virtual machine environment (MACHINE type Z) that is always in the z/Architecture architectural mode and cannot switch to

ESA/390 mode. This can be useful for testing software in a z/Architecture-only environment, in advance of deploying software on a future z/Architecture-only machine.

The Directory Maintenance Facility (DirMaint) optional feature has been enhanced to handle z/Architecture-mode virtual machines.

- **FlashSystem™ support for FCP-attached SCSI disks**

A z/VM storage administrator can use FlashSystem storage as a z/VM-system-attached DASD, attached to the host without the need for an intermediate SAN Volume Controller (SVC). Previously, while FlashSystem could be used by a Linux virtual machine without an SVC, an external or internal SVC was required to use it for z/VM system volumes, such as EDEVs, or for virtual machine minidisks. This enhancement removes that requirement.

- **IBM Virtualization Engine TS7700 Copy Export support**

APAR VM65789 provides support within the Removable Media Services (RMS) component of DFSMS/VM for the IBM Virtualization Engine TS7700 Copy Export functionality. This function allows a copy of selected logical volumes written to the TS7700 to be removed and taken offsite for disaster recovery purposes. The benefits of volume stacking, which places many logical volumes on a physical volume, are retained with this function. In addition, because the data being exported is a copy of the logical volume, the logical volume data remains accessible by the production host systems.

IBM OpenStack for z Systems

OpenStack is an Infrastructure as a Service (IaaS) cloud computing open source project, managed by the OpenStack Foundation. With the adoption of OpenStack as part of the IBM cloud strategy, virtual servers in the z/VM V6.4 for IBM z System and LinuxONE operating environments can be natively managed using OpenStack open cloud architecture-based interfaces. z/VM drivers provide OpenStack enablement for z/VM, and for z/VM virtual machines running Linux on z Systems and LinuxONE.

Open standards offer a common foundation and compatible interfaces for companies to augment computing resources at a moment's notice to meet changing client demands, such as adding a new mobile application for a new group of clients to drive additional revenue. Open Standards such as OpenStack help enterprises be more nimble and address major client concerns, like vendor lock-in, the need to hire expensive specialized experts, long application development cycles, and security challenges.

With the PTF for APAR VM65893, expected to be available first quarter 2017, the OpenStack Cloud Management Appliance (CMA) packaged with z/VM V6.4 will be enhanced to include the Newton level of OpenStack. Additional enhancements will include provisioning of Ubuntu guests, Flashcopy support for provisioning, and additional RAS and security support.

IBM Cloud Manager with OpenStack for System z has been withdrawn from marketing. The next evolution of z/VM cloud enablement technology is the OpenStack Liberty-based Cloud Management Appliance (CMA) and is available for z/VM 6.3 and 6.4. z/VM installations wanting to move forward with cloud-based solutions beyond Cloud Manager with OpenStack for System z should utilize the cloud enablement support provided by the z/VM OpenStack Liberty-based CMA.

IBM intends to upgrade the z/VM OpenStack CMA to a new level approximately once per year. This new level of the CMA would be based on every other release

of OpenStack. IBM intends to make the new CMA level available only on the most recent release of z/VM. After the CMA upgrade to a new level of OpenStack on a particular release of z/VM, the previous CMA level on that z/VM release will be supported only for another 6 months. During that 6 month period, the previous CMA level will receive only security fixes. Any functional bug fixes and enhancements will be provided only on the latest CMA level available on a z/VM release. After that 6 month period, no service will be provided for previous CMA levels and users should upgrade to the latest CMA level.

For example, when the Newton level ships that will be only for z/VM V6.4, and there would not be an update for z/VM V6.3 for Newton. The Liberty level on z/VM V6.3 will continue to be supported until the end of service date for z/VM V6.3. On z/VM V6.4, when the Newton-based CMA becomes available, only security related fixes will be made available for Liberty and for only the next 6 months.

When a new release of z/VM is made available, IBM intends to freeze the level of the CMA available on the previous z/VM release. This frozen level of the CMA will receive only security fixes until service ends for that z/VM release.

Statements of general direction

- **Stabilization of z/VM support for the IBM zEnterprise 196 (z196) family**
z/VM V6.4 is the last z/VM release planned to support the IBM Enterprise 196 (z196) or IBM zEnterprise 114 (z114) family of servers. Either an IBM zEnterprise EC12 (zEC12) or an IBM zEnterprise BC12 (zBC12) is planned as the required minimum level of server for future z/VM releases. Refer to the IBM Support Portal for the most current support lifecycle information for z/VM.
- **Removal of support for IEEE 802.3 Ethernet frame types**
z/VM V6.4 is planned to be the last z/VM release to support IEEE 802.3 Ethernet frame types. All future z/VM releases are planned to support DIX Version 2 (DIX V2) exclusively. This includes the z/VM virtual switch (VSwitch) and the z/VM TCP/IP server.
- **Removal of support for the IMAP server**
z/VM V6.4 is planned to be the last z/VM release to support IMAP.
- **Removal of support for certain TCP/IP functions**
z/VM V6.4 is planned to be the last z/VM release to support the Graphics Data Display Manager Interface for X Window System (GDDMXD/VM).
- **Installation from 3390 Model 3 DASD**
z/VM V6.4 is planned to be the last release to allow installation using Model 3 3390 DASD (Direct Access Storage Device) volumes. Future z/VM releases will support 3390 installation using only model 9 or model 27 DASD. Installation on SCSI volumes will not be affected.
- **FIPS Certification of z/VM V6.4**
IBM intends to pursue an evaluation of the Federal Information Processing Standard (FIPS) 140-2 using National Institute of Standards and Technology's (NIST) Cryptographic Module Validation Program (CMVP) for the System SSL implementation utilized by z/VM V6.4.
- **Security Evaluation of z/VM V6.4**
IBM intends to evaluate z/VM V6.4 with the RACF Security Server feature, including labeled security, for conformance to the Operating System Protection Profile (OSPP) of the Common Criteria standard for IT security, ISO/IEC 15408, at Evaluation Assurance Level 4 (EAL4+).

- **Removal of support for virtual machines with dedicated processors**
z/VM V6.4 is planned to be the last z/VM release to support dedication of logical to virtual processors via the CP DEDICATE command or with the DEDICATE option on the CPU user directory statement. z/VM running in a logical partition with dedicated processors will continue to be supported.
- **Removal of IBM Wave support for SLES 10**
In a future deliverable, IBM intends to remove IBM Wave support for management of guests that are running the SUSE Linux Enterprise Server (SLES) 10 Linux distribution.
- **Removal of IBM Wave support for second extended filesystem (Ext2)**
In a future deliverable, IBM intends to remove IBM Wave support for the administration of Linux guest file systems that use Ext2.

Statements regarding IBM plans, directions, and intent are subject to change or withdrawal without notice at the sole discretion of IBM. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Chapter 3. Technical information

The following topics provide information about z/VM V6.4 requirements and capabilities:

- “Hardware and storage requirements”
- “Device support” on page 23
- “Software (program) requirements” on page 24
- “Limitations” on page 30
- “Performance considerations” on page 31
- “Operating systems supported as guests” on page 31
- “Other programs supported on z/VM” on page 31
- “Multicultural support” on page 31
- “Packaging and ordering information” on page 32
- “Object code only and limited source availability” on page 33
- “Integrity and security” on page 33

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

Hardware and storage requirements

There are general hardware and storage requirements for installing and operating z/VM V6.4, and some z/VM facilities have additional hardware requirements. These are described in the following topics:

- “Server requirements”
- “Installation requirements” on page 20
- “Storage requirements” on page 20
- “Workstation requirement” on page 21
- “OSA/SF hardware requirements” on page 21
- “Performance Toolkit hardware requirements” on page 21
- “SSI cluster hardware requirements” on page 22
- “TCP/IP hardware requirements” on page 22
- “VM Guest RSA-Assist Support for Linux hardware requirements” on page 22
- “zHPF guest support hardware requirements” on page 23

Some hardware might be supported only on certain servers.

Server requirements

z/VM V6.4 requires a new architecture level set (ALS) and supports only IBM zEnterprise 196 (z196) and IBM zEnterprise 114 (z114) or later, and IBM LinuxONE. See Appendix A, “IBM servers supported by z/VM,” on page 87.

Specific server facilities might also be required. See IBM: z/VM V6.4 - Architecture Level Set (www.ibm.com/vm/zvm640/architecture/).

Installation requirements

z/VM V6.4 is available in the following formats:

- Binary image on DVD
- Electronic delivery (zip files)

For information about receiving the electronic deliverables, see Installation Instructions for Electronically Delivered z/VM System Deliverable and Licensed Products (www.ibm.com/vm/install/prodinst.html).

To install z/VM:

- You need a local non-SNA 3270 terminal or equivalent to use as a console, such as:
 - IBM Open Systems Adapter Integrated Communications Controller (OSA-ICC)
 - Integrated 3270 console provided through the Hardware Management Console (HMC)
 - IBM 2074 Console Support Controller
 - IBM 3174 Control Unit
- Installation from DVD requires the HMC, V1.8 or later.
- Installation of an SSI cluster (1 - 4 z/VM images) requires DASD defined to emulate an IBM 3390 Model 3 or Model 9. For more information about SSI cluster requirements, see *z/VM: CP Planning and Administration*.

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

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

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

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

Storage requirements

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

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

To IPL a z/VM system second level, a minimum of 32 MB of virtual memory is required for the user ID on the first level system.

z/VM can support up to 2 TB (terabytes) of real memory and up to 1 TB of virtual memory for a single guest. The total virtual memory in use by all guests can exceed the amount of real memory. The actual amount of usable real memory is dependent on the amount of real memory in the z/VM logical partition; the

hardware server model, firmware level, and configuration; and the number of guests and their workload. The amount of usable virtual memory is primarily dependent on the amount of DASD space allocated for paging, although it is also influenced by the number of guests and their workload characteristics.

Depending on the workload you plan to run, you might need to add capacity to take full advantage of the 2 TB support. This might mean adding memory and CPUs to your existing configuration, or adding books or drawers to your system, or reconfiguring DASD paging space.

- To assist with planning for z/VM, additional information can be found at IBM: z/VM Technical Resources (www.ibm.com/vm/techinfo/).
- For DASD configuration guidelines, see IBM: VM Performance Resources (www.ibm.com/vm/perf/).
- For guidelines on calculating the required paging space, see *z/VM: CP Planning and Administration*.

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

Workstation requirement

Many z/VM administration tasks require access to a 3270-equivalent workstation.

IBM Wave hardware requirements

For information about hardware requirements for the IBM Wave optional feature, see *IBM Wave for z/VM: Administration and Customization*.

OSA/SF hardware requirements

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

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

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

- *Open Systems Adapter-Express Customer's Guide and Reference*
- *Open Systems Adapter-Express Integrated Console Controller User's Guide*
- *Open Systems Adapter-Express Integrated Console Controller 3215 Support*
- *Open Systems Adapter-Express3 Integrated Console Controller Dual-Port User's Guide*

Performance Toolkit hardware requirements

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

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

SSI cluster hardware requirements

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

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

TCP/IP hardware requirements

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

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

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

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

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

- IBM ESCON Channel-to-Channel Adapter

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

- IBM FICON Channel-to-Channel Adapter

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

- IBM HiperSockets

z/VM programs using TCP/IP can communicate using HiperSockets with other z/VM programs using TCP/IP and with guest operating systems using TCP/IP in other logical partitions within the same central processor complex.

VM Guest RSA-Assist Support for Linux hardware requirements

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

- Corresponding function in Linux on z Systems
- Crypto Express3 or later

zHPF guest support hardware requirements

z/VM provides support for guest use of the High Performance FICON for z Systems (zHPF) feature. The minimum hardware required is determined by the program running in the virtual machine. The minimum hardware requirement varies with different programs. See the publications for the program being used to determine what is required. In addition, a few zHPF features are limited to fullpack minidisks or attached DASD.

zHPF paging hardware requirements

z/VM uses High Performance FICON for z Systems (zHPF) for paging when enabled with the FEATURES ENABLE PAGING_ALIAS system configuration statement or the SET PAGING ALIAS ON command, and the hardware supports the features required.

The minimum hardware required is:

- FICON Express8S
- One of the following or newer:
 - System Storage DS8800 Release 6.2 Bundle 86.20.98.0 (November 2011)
 - System Storage DS8870 Release 7.0 Bundle 87.0.189.0 (October 2012)
 - System Storage DS8880 Release 8.0 Bundle 88.0.153.0 (December 2015)

Additional information on hardware configurations can be found at IBM System Storage Interoperation Center (SSIC) (www.ibm.com/systems/support/storage/ssic/interoperability.wss).

PCIe support hardware requirements

z/VM provides support for guests to dedicate PCIe features to their virtual machines.

The minimum hardware requirement for the zEDC Express and 10GbE RoCE Express features is a zEnterprise EC12 or BC12 at Driver D15 with a minimum bundle level. See the PSP bucket for specific details.

Multi-VSwitch Link Aggregation support hardware requirements

Multi-VSwitch Link Aggregation support allows a port group of OSA-Express features to span multiple virtual switches within a single z/VM system or between multiple z/VM systems.

The minimum hardware requirement for Multi-VSwitch Link Aggregation support is the IBM z13 with the OSA-Express4S or later.

Device support

z/VM handles real devices as either supported or unsupported.

Supported devices

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

- Supported for CP and guest use

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

- Supported for dedicated use by a single guest

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

Unsupported devices

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

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

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

Software (program) requirements

z/VM V6.4 has prerequisite licensed programs, and some z/VM facilities have additional program requirements. These are described in the following topics:

- “Prerequisite licensed programs” on page 25
- “High Level Assembler” on page 25
- “z/VM base components additional program requirements” on page 26
- “Supported levels of z/VM base facilities and optional features” on page 26
- “DFSMS/VM additional program requirements” on page 27
- “DirMaint additional program requirements” on page 28
- “HCD and HCM additional program requirements” on page 28
- “IBM Wave program requirements” on page 28
- “IOCP additional program requirements” on page 28
- “Language Environment additional program requirements” on page 28
- “OSA/SF additional program requirements” on page 28
- “Performance Toolkit additional program requirements” on page 29
- “RACF additional program requirements” on page 29
- “RSCS additional program requirements” on page 29
- “Shared-DASD complex additional program requirements” on page 29
- “SSI clusters additional program requirements” on page 29

- “TCP/IP additional program requirements” on page 30
- “z/VM OpenStack support additional program requirements” on page 30
- “Program requirements for other products running on z/VM” on page 30

Prerequisite licensed programs

z/VM V6.4 requires the following additional IBM licensed programs:

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

The PTFs for the following APARs have been applied to the copy of EREP/VM preinstalled with z/VM V6.4:

- APAR VM63624 provides support for large disk sizes on IBM DS6000™ and DS8000.
- APAR VM64807 provides support for IBM z196.
- APAR VM64928 provides support for IBM z114.
- APAR VM65130 provides support for IBM zEC12.
- APAR VM65279 provides support for IBM zBC12.
- APAR VM65495 provides support for IBM z13.
- APAR VM65704 provides support for IBM z13s.

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

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

The PTFs for the following APARs have been applied to the copy of ICKDSF preinstalled with z/VM V6.4:

- APAR PQ84848 provides support for large FBA devices.
- APAR PQ95319 provides support for large disk sizes on IBM DS6000.
- APAR PQ96706 provides support for large disk sizes on IBM DS8000.
- APAR PM11220 provides support for the direct attachment of IBM XIV® Storage Systems.
- APAR PK87899 provides support for the CP volume ownership function of z/VM SSI cluster support.
- APAR PM99490 provides support for Multi-Target Peer-to-Peer Remote Copy (PPRC) DASD.
- APAR PI46151 provides support for z/Architecture-only servers.

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

High Level Assembler

All IBM-provided assembler source code is intended to be assembled using the IBM High Level Assembler, V1.6 (5696-234) or later.

In some circumstances, the built-in ASSEMBLE command (called by the VMFASM EXEC) may be used to apply changes to product modules that contain configuration data intended to be modified by the customer.

Use of the High Level Assembler is required for the HCD IOCP import function.

The High Level Assembler has the following additional program requirements:

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

- APAR PM49761 provides support for new instructions provided with the zEC12 and zBC12 servers.
- APAR PM79901 provides support for new instructions provided with the z13 server.

z/VM base components additional program requirements

Attention

The PTF for APAR VM65126 ("FIXES TO FCP DUMP SIMULATION") is required on first level systems where the new stand-alone dump program (SDINST) is going to be used to dump second level systems. The PTF is included on V540 RSU1202, V610 RSU1301, and V620 RSU1202. It is important for the first level system to have this fix, or various abends or incorrect operation of the new stand-alone dump program could occur.

The z/VM V6.4 base components have the following additional program requirements:

- The PTF for APAR VM65207 must be applied to z/VM V6.2 or z/VM V6.3 when z/VM V6.4 is a guest of either release and will host a MACHINE Z guest.
- If you intend to exploit HyperPAV paging from z/VM V6.4 running as a guest of z/VM V6.3, you need to apply the PTF for APAR VM65748 to the z/VM V6.3 first level system.
- The PTF for APAR VM65846 provides infrastructure support in z/VM V6.2, V6.3, and V6.4 for a future IBM z Systems server that satisfies the statement of direction for the removal of ESA/390 architecture support, and must be installed on all the members of an SSI cluster before any member will be running on the new server.
- The PTF for VM65865 is required on z/VM V6.2, V6.3, and V6.4 to support:
 - UIDs for real PCI functions
 - Asynchronous CF Duplexing for Lock Structures
- The PTF for APAR VM65867 is required on all z/VM V6.2 and V6.3 systems in an SSI cluster that includes a z/VM V6.4 system.
- To develop OpenExtensions (POSIX) applications on CMS, you need:
 - One of the following compiler programs, or an equivalent product:
 - IBM XL C/C++ for z/VM, V1.1 (5654-A22) or later
 - IBM C/C++ for z/VM, V1.1 (5654-A22)
 - IBM C for VM/ESA, V3.1 (5654-033)
 - Language Environment, level 620

Supported levels of z/VM base facilities and optional features

z/VM V6.4 supports the following base facilities and optional features. Previous levels of these facilities and features are not supported by z/VM V6.4. Additional PTFs for these facilities and features might be required to support some z/VM V6.4 functions.

- Base facilities:
 - HCD and HCM for z/VM, function level 640 (preinstalled)
 - Language Environment, level 620 (preinstalled)
This is equivalent to z/OS V1.12 Language Environment.
 - OSA/SF, function level 440 (preinstalled)
 - TCP/IP for z/VM, level 640 (preinstalled)

- Optional features:
 - DFSMS/VM, function level 221 (available with the z/VM SDO V6.4)
 - Directory Maintenance Facility for z/VM, function level 640 (preinstalled disabled)
 - IBM Wave for z/VM, Release 2 (available with the z/VM SDO V6.4)
 - IBM z/VM Single System Image Feature, function level 640 (available disabled with the z/VM SDO V6.4)
 - Performance Toolkit for VM, function level 640 (preinstalled disabled)
 - RACF Security Server for z/VM, function level 640 (preinstalled disabled)
 - RSCS Networking for z/VM, function level 640 (preinstalled disabled)

DFSMS/VM additional program requirements

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

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

Note: ISPF is not required if you are using only the 3495 Tape Library Dataserver support of DFSMS/VM.
- To use the minidisk management functions:
 - Directory Maintenance Facility for z/VM, function level 640
- If an external security manager is needed for authorizations:
 - RACF Security Server for z/VM, function level 640, 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 640
- For remote operations:
 - RSCS Networking for z/VM, function level 640
- If the compiled REXX installation-wide exit or a compiled ACS REXX exit is desired:
 - Compiler for SAA REXX/370 (5695-013)
 - Library for SAA REXX/370 (5695-014)
- Additional enhancements and support:

Note: PTFs for these APARs have been applied to the copy of DFSMS/VM available with the z/VM SDO V6.4.

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

- APAR VM65789 provides Copy Export support.

DirMaint additional program requirements

Directory Maintenance Facility for z/VM, function level 640, has the following additional program requirements:

- To assemble DirMaint exit routines or tailorable files:
 - IBM High Level Assembler, V1.6 (5696-234) or later

HCD and HCM additional program requirements

HCD and HCM for z/VM, function level 640, has the following additional program requirements:

- APAR VM65201 provides HCM support for UIDs for real PCI functions.
- APAR VM65827 provides HCD support for UIDs for real PCI functions.
- To migrate (convert) a file containing I/O configuration program (IOCP) statements into an existing HCD I/O definition file (IODF):
 - IBM High Level Assembler, V1.6 (5696-234) or later

IBM Wave program requirements

For information about program requirements for the IBM Wave optional feature, see *IBM Wave for z/VM: Administration and Customization*.

IOCP additional program requirements

The input/output configuration program (IOCP) supplied with z/VM V6.4 has no additional program requirements.

Language Environment additional program requirements

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

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

OSA/SF additional program requirements

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

Access through the OSA/SF GUI requires:

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

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

- Additional enhancements and support:

Note: PTFs for these APARs have been applied to the copy of OSA/SF preinstalled with z/VM V6.4.

- APAR OA15170 provides support for Layer 2 and for displaying information for OSN features on the server.

- APAR OA37060 provides support for the OSA-Express4S on the z196 and z114 servers.
- APARs OA37060 and OA38418 provide support for the OSA-Express4S on the zEC12 and later servers.
- To support an OSA in TCP/IP Passthru mode if access to one of the OSA's ports is being shared by more than one host program:
 - TCP/IP for z/VM, level 640

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

- *Open Systems Adapter-Express Customer's Guide and Reference*
- *Open Systems Adapter-Express Integrated Console Controller User's Guide*
- *Open Systems Adapter-Express Integrated Console Controller 3215 Support*
- *Open Systems Adapter-Express3 Integrated Console Controller Dual-Port User's Guide*

Performance Toolkit additional program requirements

Performance Toolkit for VM, function level 640, has the following additional program requirements:

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

RACF additional program requirements

RACF Security Server for z/VM, function level 640, has no additional program requirements.

RSCS additional program requirements

RSCS Networking for z/VM, function level 640, has the following additional program requirements:

- Group Control System (GCS) must be configured and activated.
- To communicate with users in an SNA network:
 - VTAM for VM/ESA, V4.2 (5654-010)
- To communicate within an IP network:
 - TCP/IP for z/VM, level 640
- To assemble RSCS exit routines:
 - IBM High Level Assembler, V1.6 (5696-234) or later

Shared-DASD complex additional program requirements

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

- DirMaint, function level 640
 - Within the shared-DASD complex, all z/VM systems must be licensed for the DirMaint FL640 feature and must be running the same DirMaint service level.

SSI clusters additional program requirements

To create and manage a z/VM single system image (SSI) cluster, the following programs and additional support are required:

- IBM z/VM Single System Image Feature (VMSSI), function level 640, licensed for the server processors that will host the members of the SSI cluster.
- Directory Maintenance Facility, function level 640, or an equivalent product.
- APARs VM65317, VM65318, VM65319, and VM65320 are required on z/VM V6.2 systems to support maintaining the highest level of common parts in a mixed release SSI cluster.
- APARs VM65846 and VM65867 are required on z/VM V6.2 and V6.3 systems in an SSI cluster that also includes a z/VM V6.4 system.
- If an external security manager is needed:
 - RACF Security Server for z/VM, function level 640, or an equivalent product

TCP/IP additional program requirements

TCP/IP for z/VM, level 640, has the following additional program requirements:

- If an external security manager is needed:
 - RACF Security Server for z/VM, function level 640, or an equivalent product
- To develop programs in C:
 - One of the following compiler programs, or an equivalent product:
 - IBM XL C/C++ for z/VM, V1.1 (5654-A22) or later
 - IBM C/C++ for z/VM, V1.1 (5654-A22)
 - IBM C for VM/ESA, V3.1 (5654-033)
 - Language Environment, level 620
- To develop programs in Pascal:
 - IBM VS Pascal, V1.2 (5668-767), Compiler and Library, or an equivalent product

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

z/VM OpenStack support additional program requirements

For information about additional program requirements for z/VM OpenStack support, including the z/VM Cloud Manager Appliance (CMA), see the *Enabling z/VM for OpenStack* publication for the OpenStack release you intend to use. For the current list of publications, see “Planning and administration” on page 71.

Program requirements for other products running on z/VM

- Linux on z Systems
Linux is not provided with z/VM. In order to operate Linux as a guest of z/VM, Linux must be obtained from a Linux distribution partner. For current information on Linux distributions, see IBM z Systems - Operating Systems - Linux on IBM z Systems (www.ibm.com/systems/z/os/linux/).
- Other IBM licensed products might also have additional program requirements to run on z/VM V6.4. See the product documentation and the product websites. Also see IBM: z/VM and VM-related products and features (www.ibm.com/vm/related/).

Limitations

z/VM has the following limitations:

- This z/VM release has an Architecture Level Set (ALS) and operates only on specific IBM z Systems servers. See Appendix A, “IBM servers supported by z/VM,” on page 87.
- z/VM supports only ESA/390 mode (on z13 and z13s and earlier servers), General mode (on later servers), LINUX only mode, and z/VM mode LPARs.
- z/VM may run on IFL processors only if, on the IFL processors, z/VM is being used exclusively to run (1) Linux or OpenSolaris workloads and, (2) if required, z/VM applications in support of those Linux or OpenSolaris workloads.
- 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.

Performance considerations

System performance depends on the hardware resources allocated to z/VM and on the level of activity within each guest. For more information, see IBM: VM Performance Resources (www.ibm.com/vm/perf/). In particular, you should closely review the *z/VM V6.4 Performance Report*.

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

Operating systems supported as guests

See Appendix B, “IBM operating systems supported as guests of z/VM,” on page 93.

Other programs supported on z/VM

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

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

Multicultural support

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

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

Table 1. IBM-translated parts of z/VM

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

Packaging and ordering information

To order z/VM V6.4, order the z/VM System Delivery Offering (SDO) V6.4, program number 5741-A08. The SDO is the IBM program offering for z/VM which includes a single deliverable containing the z/VM operating system and a group of fully serviced IBM Licensed Program Products.

The z/VM SDO V6.4 package contains:

- z/VM V6.4 System Image and RSU (recommended service upgrade)
- z/VM SDO optional features and licensed products, if ordered

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

The z/VM V6.4 System Image is available in the following formats:

- Binary image on DVD
- Electronic delivery (zip files)

The z/VM V6.4 System Image contains:

- Prerequisite IBM licensed programs:
 - EREP/VM, V3.5 (5654-260)
 - ICKDSF for VM, V1.17 (5684-042) (includes ICKDSF stand-alone release 17)
- z/VM V6.4 base product:
 - AVS
 - CMS
 - CP
 - Dump Viewing Facility
 - GCS
 - HCD and HCM for z/VM, function level 640
 - Language Environment, level 620
 - OSA/SF, function level 440
 - REXX/VM
 - TCP/IP for z/VM, level 640 (includes source)
 - TSAF
 - VMSES/E
 - SFS and CRR file pools
 - DASD Dump Restore (DDR) program
 - Input/output configuration program (IOCP)
 - Translated CP and CMS message repositories
 - English and translated z/VM HELP files

- 3800 printer image library object code
- 3800 printer image library source
- z/VM optional features:
 - DirMaint, function level 640
 - Performance Toolkit for VM, function level 640
 - RACF Security Server for z/VM, function level 640
 - RSCS Networking for z/VM, function level 640

Notes:

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

Object code only and limited source availability

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

- AVS
- Dump Viewing Facility
- DFSMS/VM
- Performance Toolkit
- IBM Wave for z/VM

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

Integrity and security

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

Data integrity for guests

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

System integrity statement for z/VM

System integrity is an important characteristic of z/VM. 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 "Integrity and security" on page 33. IBM will continue its efforts to enhance the integrity of z/VM and to respond promptly when exposures are identified in the specified operating environment on releases of z/VM that have not reached their End of Support Date, which can be found at IBM Software Support: Support Lifecycle (www.ibm.com/software/support/lifecycle/index_z.html).

Note: IBM reserves the right to change, modify or withdraw its offerings, policies and practices at any time. All products and support obligations are subject to the terms of the applicable license and services agreements.

z/VM system integrity definition

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

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

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

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

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

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

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

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

Customer responsibilities

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

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

DirMaint system integrity

DirMaint uses standard z/VM system facilities to:

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

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.
- Working together, z/VM and the optional RACF feature support the use of passwords longer than eight characters, called password phrases, that may contain mixed-case letters, numbers, blanks, and special characters.
- User class modification 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 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.
- z/VM supports drive-based data encryption with IBM System Storage tape drives. z/VM supports native and guest use of the IBM Encryption Key Manager for encryption of tapes.
- z/VM supports the use of DASD volumes that reside on data encryption drives (DEDs).

Support for cryptographic facilities on IBM servers

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

- Regional Crypto Enablement (RCE) adapters (on certain servers)
RCE adapters are supported for use by z/OS guests.
- Crypto Express3 or later (depends on the server)
Guest operating systems such as Linux and z/OS can be authorized in the z/VM directory for dedicated or shared use of cryptographic features.
- CP Assist for Cryptographic Functions (CPACF)
Authorization is not required to use CPACF.

CMS file pool security

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

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

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

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

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

TCP/IP security

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

- Secure Sockets Layer (SSL) support
SSL provides secure (encrypted) communication between a remote client and a TCP/IP application server. Under SSL protocol, the application server is always

authenticated. To participate in an SSL session, an application server must provide a certificate from a Certifying Authority to prove its identity.

The TCP/IP System SSL cryptographic module has received Federal Information Processing Standard (FIPS) 140-2 certification.

- Transport Layer Security (TLS) support
TLS is the successor to SSL and offers additional security measures, including support for SHA-256 certificates.
- Exits to an ESM
TCP/IP servers can be configured to use RACF Security Server for z/VM or another ESM to:
 - Authenticate and authorize access to resources managed by the server
 - Validate client user IDs and passwords
- LDAP server and client
The Lightweight Directory Access Protocol (LDAP) server and associated client utilities provide user authentication, authorization, and auditing capabilities.

DirMaint security

DirMaint supports the z/VM security strategy:

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

Security provided by RACF Security Server for z/VM

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

Technical Information

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

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

Chapter 4. z/VM base product

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

- “Control Program (CP)”
- “Conversational Monitor System (CMS)” on page 51
- “TCP/IP for z/VM” on page 56
- “APPC/VM VTAM Support (AVS)” on page 58
- “Dump Viewing Facility” on page 58
- “Group Control System (GCS)” on page 59
- “Hardware Configuration Definition (HCD) and Hardware Configuration Manager (HCM) for z/VM” on page 59
- “Language Environment” on page 60
- “Open Systems Adapter Support Facility (OSA/SF)” on page 60
- “REXX/VM” on page 60
- “Transparent Services Access Facility (TSAF)” on page 61
- “Virtual Machine Serviceability Enhancements Staged/Extended (VMSES/E)” on page 61

Control Program (CP)

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

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

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

Virtual machine modes

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

- **ESA** virtual machines process according to ESA/390 architecture (31-bit). ESA virtual machines are also capable of processing according to z/Architecture (64-bit) if switched into that mode by a guest operating system.
- **XA** virtual machines are supported for compatibility and are functionally equivalent to ESA virtual machines. Some CMS applications might require CMS to be running in an XA virtual machine.
- **XC** virtual machines process according to ESA/XC architecture.
- **Z** virtual machines process according to z/Architecture (64-bit) and cannot be switched to ESA/390 mode.

Notes:

1. Although 370-mode (IBM System/370 architecture) virtual machines are not supported, 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), which means the virtual machine's guest real storage is paged into and out of host real storage.

Architecture compatibility and support

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

z/Architecture guests can use z/Architecture facilities in z/VM, such as 64-bit arithmetic and addressing. A guest in a Z virtual machine runs entirely in z/Architecture mode. A guest in an ESA or XA virtual machine can switch the virtual machine from ESA/390 mode to z/Architecture mode.

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 (XC virtual machines). ESA/XC is an architecture unique to virtual machines and is available only in XC 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 XC virtual machines share multiple data spaces. An XC virtual machine can access one or more data spaces of another XC virtual machine if so authorized. This is extremely useful for applications that require one or more virtual machines to serve many users.

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

Guest LANs and virtual switches

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

- 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.
- 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 construct, known as a virtual switch, to an associated real LAN connected by an OSA-Express QDIO feature. A virtual switch allows connectivity to external LAN segments without requiring a router.

In addition, a virtual switch can transparently bridge from a bridge-capable HiperSockets channel to a network connected by an OSA-Express (OSD) feature.

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

z/VM HiperDispatch

z/VM HiperDispatch can improve CPU efficiency by causing CP to run workload in a manner that recognizes and exploits z Systems machine topology to increase the effectiveness of physical machine memory cache. This includes:

- Requesting the LPAR to handle the partition's logical processors in a manner that exploits physical machine topology
- Dispatching virtual servers in a manner that tends to reduce their movement within the partition's topology
- Dispatching multiprocessor virtual servers in a manner that tends to keep the server's virtual CPUs logically close to one another within the partition's topology

z/VM HiperDispatch can also improve CPU efficiency by automatically tuning the partition's use of its logical CPUs to try to reduce multiprocessor (MP) effects. This includes:

- Sensing and forecasting key indicators of workload intensity and of elevated MP effect
- Autonomically tuning the system to reduce MP effect when it seems MP effect is decreasing CPU efficiency

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

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 or DUMPLD2 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 stand-alone 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 58.

Support for hardware facilities, features, and architectures

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

Notes:

1. Some support depends on hardware level and availability. For detailed information on hardware capabilities, requirements, and availability, see the appropriate Hardware Announcements.
2. In some cases, the supported function is available only to guests.

Collaborative Memory Management Assist

CP support for the Collaborative Memory Management Assist (CMMA) machine feature, in conjunction with CMMA exploitation in z/Architecture guest operating systems such as Linux on z Systems, allows the CP host and its guests to communicate attributes for specific 4 KB blocks of guest memory. This exchange of information allows both host and guest to optimize their use and management of memory. For more information, see *z/VM: Performance*.

Concurrent-sense facility

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

CPU-Measurement Counter Facility

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

Cryptographic facilities

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

- Regional Crypto Enablement (RCE) adapters (on certain servers). RCE adapters are vendor products that provide country specific cryptographic functions. The RCE adapters are attached to the CEC via PCIe, and operating systems interact with the adapters via a HiperSockets network.
- Crypto Express3 or later (depends on the server)
- CP Assist for Cryptographic Functions (CPACF), which supports high-performance bulk encryption using clear keys or encrypted keys.

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

Dynamic storage reconfiguration

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

Enhanced-DAT Facility

The enhanced-DAT facility allows a guest to exploit large (1 MB) pages. A larger page size decreases the amount of guest memory needed for dynamic address translation (DAT) tables and also decreases the overhead required to perform address translation. In all cases, guest memory is mapped into 4 KB pages at the host level.

With guest large page support, Linux on z Systems, z/VSE, and z/OS virtual machines can benefit from reduced memory footprints and address translation times. This can decrease overhead and improve throughput.

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.

Environment information interface

The Store Hypervisor Information (STHYI) instruction can be used to access certain system information, including current CPU resources available at the machine, logical partition, hypervisor, and guest levels, as well as any caps (such as CPU pools) that restrict the guest's use of these resources. This information enables an application to determine the maximum capacity of CP and IFL CPUs available to software.

Extended Channel Measurement Data Support

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

Extended-TOD-Clock Facility

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

Fibre Channel Protocol

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

Fibre Connection channels

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

- A FICON Express4 channel is designed to deliver up to 4 Gbps.
- A FICON Express8 or Express8S channel is designed to deliver up to 8 Gbps.
- A FICON Express16S channel is designed to deliver up to 16 Gbps.

FICON Channel-to-Channel Adapter

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

FlashCopy

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

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

Guest ASCII console

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

High Performance FICON for z Systems

z/VM provides guest support for High Performance FICON for z Systems (zHPF). zHPF is a performance and reliability, availability, serviceability (RAS) enhancement of the z/Architecture and the FICON channel architecture implemented in the supported servers and the System Storage DS8000 series. See “zHPF guest support hardware requirements” on page 23.

Exploitation of zHPF by the FICON channel, the operating system, and the control unit is designed to help reduce the FICON channel overhead. zHPF channel programs can be especially useful for I/O workloads that transfer small (4 KB) blocks of fixed-sized data. The supported FICON Express features support the FICON architecture, FICON channel-to-channel (CTC), and the zHPF architecture simultaneously.

HiperSockets

CP supports the z/Architecture HiperSockets function for high-speed TCP/IP communication among virtual machines and logical partitions (LPARs) within the same central processor complex. 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.

In addition, a z/VM virtual switch can transparently bridge from a bridge-capable HiperSockets channel to a network connected by an OSA (OSD) feature.

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

HyperPAV

z/VM exploits the ability for an IBM DS8000 device to issue concurrent I/O requests to an ECKD paging volume. In HyperPAV mode, if the base volume is busy, z/VM selects a free alias device from a pool, binds the alias to the base device, and starts the I/O. When the I/O completes, the alias device is returned to the pool to be used for another I/O to the same logical subsystem (LSS).

The primary benefit of exploiting HyperPAV is to improve paging throughput during periods of high-volume disk I/O, which will increase the efficiency of the z/VM frame replenishment algorithm used to manage storage overcommitted workloads. HyperPAV paging also enables the management of fewer and larger CPOWNER volumes.

HyperPAV paging is exploited by the z/VM hypervisor for:

- The SYSRES volume, and volumes containing checkpoint and warm start data
- Volumes used for paging, spooling, and the z/VM user directory
- Minidisk pools, as defined by a guest's use of the MAPMDISK IDENTIFY macro

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.

Logical channel subsystem

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

LPAR group absolute capacity capping

This function allows an absolute capacity cap to be set by CPU type on a group of LPARs. This allows each of the partitions to consume capacity up to their individual limits as long as the group's aggregate consumption does not exceed the group absolute capacity limit. The z/VM support includes use of this information in functions that exploit or report processor capacity information for the purpose of tolerating its use.

MIDAWs

z/VM supports guest use of Modified Indirect Data Address Words (MIDAWs). This server feature can allow more flexibility and performance in certain channel programs as an alternative to data-chained channel-command words (CCWs). MIDAWs accommodate noncontiguous data areas that cannot be handled by the predecessor indirect-data-address words (IDAWs). z/VM support for guest use of MIDAWs can allow operating systems such as z/OS to use this new aspect of z/Architecture without regard to whether the operating systems are running in a

logical partition or a virtual machine. This allows guest operating systems to exercise their code-paths just as they would on the real machine during, for example, preproduction testing of z/OS systems. Likewise, the provision of the function in a virtual machine allows guest operating systems to benefit from the real machine's added-value function just as though the guests were running directly on the machine.

Move-Page Facility

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

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 hardware capability to define multiple virtual FCP channels, each with its own worldwide-unique Fibre Channel port name (WWPN). By assigning different hardware-defined virtual port names to different guests, the guests can use the virtual FCP channels as if they were dedicated physical FCP channels.

NPIV is applicable to all of the FICON features supported on the server. 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.

CP also supports:

- OSA-Express Integrated Console Controller (OSA-ICC), a function of the OSA-Express 1000BASE-T Ethernet feature that integrates 3270 emulation for console session connections into the server
- Open Systems Adapter for the Network Control Program (OSA NCP), a function of the Gigabit Ethernet feature (OSA-Express2 and OSA-Express3 only) and the 1000BASE-T Ethernet feature (OSA-Express2 and later) that provides channel connectivity from z Systems operating systems to the IBM Communication Controller for Linux (CCL) on the server.
- Virtual-switch-controlled link aggregation (IEEE 802.3ad), which allows you to dedicate an OSA port to z/VM when the port is participating in an aggregated group. Link aggregation (trunking) is designed to allow you to combine multiple physical OSA-Express ports (OSA-Express2 and later) into a single logical link for increased throughput and for nondisruptive failover in the event that a port becomes unavailable.
- Port isolation security, which provides the ability to restrict guest-to-guest communications within a virtual switch, and OSA-Express QDIO data connection isolation, which provides the ability to isolate QDIO data connection communications on a shared OSA port in support of multitier security zones.
- Virtual Edge Port Aggregator (VEPA) mode on the z/VM virtual switch, which provides the capability to take all virtual machine traffic sent by the server and send it to an adjacent network switch. This mode of operation moves all frame relay switching from the hypervisor virtual switch to the (external) adjacent switch. With the adjacent switch handling the frame relay for virtual switch guest port to guest port communications, imbedded network based appliances in the adjacent switch such as firewalls, Access Control Lists (ACLs), Quality of Service (QoS), and port mirroring are available to be deployed for this guest port to guest port switching. VEPA eliminates the need to provide and support these

network based appliances in the hypervisors or LPARs. VEPA mode is supported on OSA-Express3 and later on IBM zEnterprise EC12 and later.

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 (OSA/SF)” on page 60.

Parallel Access Volumes

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

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

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

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

Peer-to-Peer Remote Copy

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

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

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

Peripheral Component Interconnect Express (PCIe)

CP allows guests to dedicate Peripheral Component Interconnect Express (PCIe) features to their virtual machines.

- **IBM zEDC Express**

The IBM zEnterprise Data Compression (zEDC) Express feature supports a data compression function that can provide high-performance, low-latency compression without significant CPU overhead.

- **IBM 10GbE RoCE Express**

The term RoCE refers to Remote Direct Memory Access over Converged Ethernet. The IBM 10 Gigabit Ethernet (10GbE) RoCE Express feature can reduce consumption of CPU resources for applications utilizing the TCP/IP stack (such as WebSphere® Application Server accessing a DB2 database). Use of the 10GbE RoCE Express feature can help reduce network latency with memory-to-memory transfers for guests that support Shared Memory Communications-Remote Direct Memory Access (SMC-R). It is transparent to applications and can be used for LPAR-to-LPAR communication on a single system or server-to-server communication in a multiple-CPC environment.

CP also supports guest usage of virtual PCIe functions (PCI type ISM), which use Shared Memory Communications-Direct (SMC-D) connectivity.

Queued Direct I/O

CP supports guest use of Queued Direct I/O (QDIO), which allows a program to directly exchange data with an I/O device without performing traditional I/O instructions. To exchange data, both the I/O device and the program reference main storage directly through a set of data queues.

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

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

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

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

SAN Volume Controller

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

Server Time Protocol

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

For example, this allows z/VM to participate in a Geographically Dispersed Parallel Sysplex® Metro/z/OS Global Mirror (GDPS/MzGM) environment.

Shared Memory Communications - Direct Memory Access (SMC-D)

This technology is the latest networking innovation for the IBM z13 family of processors. It provides support for fast, low-latency LPAR-to-LPAR TCP/IP traffic using SMC-D software protocol over firmware-provided Internal Shared Memory (ISM) devices.

SMC-D and ISM are designed to use shared memory areas to provide low-latency, high-bandwidth, cross-LPAR connections for applications. This support is intended to provide application-transparent DMA communications to TCP endpoints for sockets-based connections. SMC-D is expected to provide substantial performance, throughput, response time, and CPU consumption benefits compared with standard TCP/IP communications over HiperSockets. z/VM supports dynamic I/O and guest use of the new Internal Shared Memory (ISM) PCI function type.

Simultaneous multithreading (SMT)

Simultaneous multithreading (SMT) technology is available on the IBM z13 and offers intelligently implemented 2-way simultaneous multithreading. SMT allows two active instruction streams per core, each dynamically sharing the core's execution resources. SMT is available for workloads running on IFL and zIIP.

z/VM provides host exploitation support for SMT on the z13, which enables z/VM to dispatch work on up to two threads (logical CPUs) of an IFL processor core. z/VM multithreading support is enabled only for IFL processors in a LINUX only mode or z/VM mode logical partition.

z/VM exploitation of SMT enables z/VM on the z13 to dispatch work on an individual thread of an IFL core, allowing a core to be shared by multiple guest CPUs or z/VM Control Program tasks. This can result in increased work throughput per core from more efficient use of shared core resources.

Multithreading support is available on a z/VM system only if the facility is available on the hardware and enabled on the z/VM system with the MULTITHREADING system configuration statement. The MULTITHREADING statement is optional, and multithreading is disabled if the statement is omitted.

If multithreading is enabled, the SET MULTITHREAD command can be used to nondisruptively switch between one and two activated threads per IFL core. Performance of a system and workload with one active thread per core is comparable to that of the same system and workload with multithreading disabled. Thus, the dynamic SMT level capability allows the benefit of multithreading to be evaluated for a workload without requiring an outage to enable or disable SMT.

z/VM enablement of multithreading requires that z/VM is configured to run with the HiperDispatch vertical polarization mode enabled and with the dispatcher work distribution mode set for reshuffle. Once enabled, multithreading cannot be disabled without a re-IPL. Dedication of virtual CPUs to z/VM processors is not supported with multithreading because it is not allowed when HiperDispatch vertical polarization mode is active.

When multithreading is enabled, prorated core time is used in the pool limiting calculation for all CPU pools and in the consumption limiting calculation for a

virtual machine being consumption limited. If multithreading is not enabled, raw CPU time is used instead. For an explanation of the three measures of CPU time, see *z/VM: Performance*.

z/VM host multithreading exploitation support does not virtualize threads for guest exploitation. However, Linux guests might benefit from the host support because the first level z/VM system is able to get higher throughput from the multithreaded IFL cores.

Specialty processors

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

IBM Integrated Facility for Linux (IFL)

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

IBM Internal Coupling Facility (ICF)

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

IBM zEnterprise Application Assist Processor (zAAP)

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

Note: zAAPs are not supported on the IBM z13 or later, either in real hardware or in the LPAR configuration.

IBM z Integrated Information Processor (zIIP)

zIIPs are designed to help improve resource optimization and lower the cost for eligible workloads. z/OS exploits zIIPs to offload software system overhead from standard central processors (CPs). This includes certain DB2 processing, enhancing the role of the mainframe as the data hub of the enterprise.

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

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

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

Note: z/VM will not allow a guest to define a processor type if the machine does not support it.

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

Transactional Execution facility

z/VM supports guest exploitation of the Transactional Execution (TX) facility on supported machines. The TX facility allows a program to issue multiple instructions that appear to operate atomically, offering an alternative to more costly mutual-exclusion mechanisms such as software locks. This support can improve the efficiency and scalability of multithreaded software such as Java or guest operating system functions.

Vector Facility for z/Architecture (SIMD)

z/VM supports guest exploitation of the Vector Facility for z/Architecture (SIMD). This support enables guest use of Vector-Facility instructions and the 128-bit vector registers used by these instructions. This support also enables the use of the additional floating-point (AFP) registers for a guest at logon time. As a result, any output that includes floating-point registers will report all 16 floating-point registers when they are available to the user.

Virtual machine I/O assist

CP supports the virtual machine I/O assist facility for guests, which forwards an adapter interruption to a guest while the CPU keeps running in guest state. The virtual machine I/O assist can be enabled or disabled for all guests or specific guests.

Conversational Monitor System (CMS)

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

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

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

Versions of CMS

z/VM provides two versions of CMS:

- ESA/390 CMS (CMS)

ESA/390 CMS runs in an ESA/390 architecture (ESA or XA mode) or ESA/XC architecture (XC mode) virtual machine.

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

- z/Architecture CMS (z/CMS)

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

IPLeD in a virtual machine with more than 2 GB of storage, and programs running on z/CMS can allocate storage above 2 GB.

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

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

Unless otherwise indicated in the z/VM publications, “CMS” means either version, and descriptions of CMS functions apply to both ESA/390 CMS and z/CMS.

Shared File System

The Shared File System (SFS) is an extension of the CMS file system that offers 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.
- Files and directories can be shared with users on other systems.

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

For more information, see:

- “CMS file pool security” on page 36
- *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 (DFSMS/VM)” on page 63.

CMS application programming

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

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

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

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

Systems management application programming interface

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

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

CMS Pipelines

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

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

The CMS OpenExtensions environment 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 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 OpenExtensions POSIX.2 utilities additionally conforms to the X/Open Portability Guide, issue 4 (XPG4) for Commands and Utilities.

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

For information about the program requirements for developing OpenExtensions applications, see “z/VM base components additional program requirements” on page 26.

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 be either read directly into virtual storage for execution or stored in a program library. For more information, see *z/VM: Program Management Binder for CMS*.

Reusable Server Kernel

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

The Reusable Server Kernel is an “empty” server program that server writers can use as a starting point for developing and executing server programs on CMS. The Reusable Server Kernel consists of a text library of routines and a macro library of

function prototypes and constant definitions. To construct an actual server program, the server author attaches application-specific code to a set of interfaces in the Reusable Server Kernel.

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

Virtual Machine Resource Manager

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

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

VMRM Cooperative Memory Management (VMRM-CMM) between a z/VM system and Linux guests assists in managing memory constraint in the system. Based on several variables obtained from the system and storage domain CP monitor data, VMRM detects when there is such constraint, and notifies specific Linux virtual guests when this occurs. The guests can then take the appropriate action to adjust their memory utilization in order to relieve this constraint on the system, such as issuing a CP DIAGNOSE X'10' instruction to release pages of storage.

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

XEDIT

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

HELP facility

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

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

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

TCP/IP for z/VM

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

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

Connectivity and gateway functions

Handle the physical interfaces and routing of data.

Server functions

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

Client functions

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

Network status and management functions

Detect and solve network problems.

Application programming interfaces (APIs)

Allow you to write your own client/server applications.

Link protocols

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

Network protocols

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

Internet protocol

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

In TCP/IP for z/VM, a single stack provides support for static routing of IP version 6 (IPv6) traffic and IPv4 traffic. IPv6 uses 128-bit addresses. However, the z/VM TCP/IP stack cannot be configured as a tunnel endpoint for tunneling IPv6 traffic over IPv4 networks.

Transport protocols

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

Transmission Control Protocol

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

User Datagram Protocol

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

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

Applications and protocols

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

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

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

- **File Transfer Protocol (FTP)** allows you to transfer data between local and foreign hosts or between two foreign hosts.
- **GDDMXD** is an interface that allows graphics from the IBM Graphical Data Display Manager/VM to be displayed on workstations that support the X Window System.
- **Internet Message Access Protocol (IMAP)** provides the processing that allows a client to access electronic mail that is kept in an IMAP Mailstore server.
- **Lightweight Directory Access Protocol (LDAP)** provides user authentication, authorization, and auditing capabilities.
- **MPRoute** uses either Open Shortest Path First (OSPF) or Routing Information Protocol (RIP), or both, to dynamically create and maintain network routing tables.
- **Network File System (NFS)** allows you to manipulate files on different TCP/IP hosts as if they reside on your host.
- **Remote Execution Protocol (REXEC)** allows you to execute a command on a foreign host and receive the results on the local host.
- **Remote Printing (LPR and LPD)** provides both client and server support for remote printing.

Note: The native z/VM TCP/IP LPD server is no longer supported, but z/VM still supports the LPD protocol using non-z/VM servers. An LPD link driver is still provided in the z/VM RSCS feature.

- **Remote Procedure Call (RPC)** is a programming interface that calls subroutines to be executed on a foreign host.
- **Simple Mail Transfer Protocol (SMTP)** is an electronic mail protocol that is implemented with the CMS NOTE and SENDFILE commands.
- **Simple Network Management Protocol (SNMP)** provides a means for managing an internet environment.
- **Secure Sockets Layer (SSL)** is a cryptographic protocol that provides communication security over the Internet.
- **Socket interfaces** allow you to write your own applications to supplement those supplied by TCP/IP.
- **Telnet Protocol** provides a standard method to interface terminal devices and terminal-oriented processes with each other.
- **Transport Layer Security (TLS)** is a cryptographic protocol that provides communication security over the Internet.
- **Trivial File Transfer Protocol** reads and writes files to and from a foreign host.
- **X Toolkit** is a collection of basic C language routines for developing a variety of application environments.
- **X Window System** supports network transparent windowing and graphics.

TCP/IP requirements and additional information

For additional information about TCP/IP for z/VM, see:

- “TCP/IP hardware requirements” on page 22
- “TCP/IP additional program requirements” on page 30
- “TCP/IP security” on page 36
- *z/VM: TCP/IP Planning and Customization*
- *z/VM: TCP/IP User's Guide*

APPC/VM VTAM Support (AVS)

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

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

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

Dump Viewing Facility

The Dump Viewing Facility helps you interactively diagnose system problems. Using this facility, you can display, format, and print data interactively from virtual machine dumps, as well as display and format recorded trace data. The BLOCKDEF utility lets you display, format, and print control block information.

The VIEWSYM command lets you display symptom records, making it easier to identify duplicate problems when they occur. For more information, see *z/VM: Dump Viewing Facility*.

Note: Analysis of CP stand-alone 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 (GCS)

The Group Control System (GCS) runs in an XA or XC virtual machine in place of CMS. GCS 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 (HCD) and Hardware Configuration Manager (HCM) for z/VM

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

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

HCD runs in a z/VM server virtual machine and performs the work of actually creating and changing the hardware and software aspects of your I/O configuration. While HCM provides the primary user interface to HCD, HCD also provides a backup user interface on your z/VM host for certain I/O configuration tasks, in case HCM is not available.

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

For more information, see:

- “HCD and HCM additional program requirements” on page 28
- *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:

- “Language Environment additional program requirements” on page 28
- *z/VM: Language Environment User’s Guide*
- *z/OS: Language Environment Concepts Guide*

Open Systems Adapter Support Facility (OSA/SF)

Open Systems Adapter-Express (OSA-Express) is a family of integrated hardware features that allow the z Systems platform to provide industry-standard connectivity directly to clients on local area networks (LANs) and wide area networks (WANs). The Open Systems Adapter Support Facility (OSA/SF) is a host-based tool supplied with z/VM that allows you to customize the modes of operation of an OSA-Express device. 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:

- “OSA/SF hardware requirements” on page 21
- “OSA/SF additional program requirements” on page 28
- “Server features and miscellaneous devices” on page 106
- *Open Systems Adapter-Express Customer’s Guide and Reference*

REXX/VM

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

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

Transparent Services Access Facility (TSAF)

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

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

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

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

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

Chapter 5. z/VM optional features

z/VM offers the following optional features:

- “Data Facility Storage Management Subsystem for VM (DFSMS/VM)”
- “Directory Maintenance Facility for z/VM (DirMaint)” on page 64
- “IBM Wave for z/VM” on page 65
- “Performance Toolkit for VM” on page 65
- “RACF Security Server for z/VM” on page 66
- “RSCS Networking for z/VM” on page 67
- “z/VM Single System Image Feature” on page 68

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

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

- Space Management

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

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

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

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

DFSMS/VM Removable Media Services (RMS) provides native z/VM support for IBM Tape Library Dataservers. For a list of the supported devices, see “Tape units and tape libraries” on page 102.

For more information, see:

- “DFSMS/VM additional program requirements” on page 27
- “DFSMS/VM publications” on page 82

Directory Maintenance Facility for z/VM (DirMaint)

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

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

For more information, see:

- "DirMaint system integrity" on page 35
- "DirMaint security" on page 37
- "DirMaint additional program requirements" on page 28
- "Directory Maintenance Facility for z/VM publications" on page 83

IBM Wave for z/VM

IBM Wave for z/VM (IBM Wave) is a provisioning and productivity management solution for managing both Linux servers and z/VM. Using a graphical environment, IBM Wave provides an innovative approach to the task of managing one or multiple IBM z Systems servers that may be configured with one or many z/VM instances that can each run thousands of virtual Linux servers.

IBM Wave's unique graphical display of the virtual server environment and physical infrastructure includes:

- Physical servers (IBM mainframes)
- z/VM instances (LPARs)
- Virtual Linux server objects
- Virtual Networks (Guest LANs and VSwitches)
- Virtual-servers-to-Virtual Network connections
- Storage volumes and storage groups

Alongside the rich display, the graphical user interface (GUI) provides all the necessary procedures and functions that are needed for routine management and provisioning tasks and special operations. The innovative design and expert system level know-how of both the Linux and the z/VM aspects makes IBM Wave the clear choice for z/VM and Linux provisioning and management.

One of the primary objectives in the design of IBM Wave is to reduce the learning curve that is typically needed to manage and control z/VM and Linux guests. IBM Wave abstracts the z/Architecture and z/VM virtualization infrastructure so that Linux system administrators can continue to manage their servers with the same skill set they currently possess. This convenience allows for day-to-day operations, along with large scale (virtual) hardware configuration changes, to be completed without extra help from z/VM support.

IBM Wave allows IT organizations and service providers to simplify and automate z/VM administration, which makes it the expert solution for medium-to-large scale consolidation projects that are done in the z/VM environment.

For more information, see:

- *IBM Wave for z/VM: Administration and Customization*
- *IBM Wave for z/VM: User Guide and Reference*

Performance Toolkit for VM

Performance Toolkit for VM 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

Optional Features

In addition to analyzing VM performance data, the Performance Toolkit can process performance data collected by Linux guests.

For more information, see:

- “Performance Toolkit hardware requirements” on page 21
- “Performance Toolkit additional program requirements” on page 29
- “Performance Toolkit for VM publications” on page 84

RACF Security Server for z/VM

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

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

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

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

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

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

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

For more information, see:

- “RACF additional program requirements” on page 29
- “RACF Security Server for z/VM publications” on page 84

RSCS Networking for z/VM

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

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

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

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

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

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

RSCS can communicate with station nodes that are:

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

Optional Features

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

For more information, see:

- “RSCS additional program requirements” on page 29
- “RSCS Networking for z/VM publications” on page 85

z/VM Single System Image Feature

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

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

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

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

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

For more information, see:

- “SSI cluster hardware requirements” on page 22
- “SSI clusters additional program requirements” on page 29
- *z/VM: CP Planning and Administration*

Chapter 6. z/VM library guide

This guide includes the following topics:

- “Where to get information about z/VM”
- “Abstracts of the z/VM base product publications” on page 70
- “Abstracts of the z/VM optional feature publications” on page 82

Where to get information about z/VM

Information about z/VM is available from several sources.

Documentation supplied with z/VM

The following documentation is supplied with z/VM:

- One copy of the *z/VM Agreements and License Information* DVD
- One printed copy of the *z/VM: Installation Guide*
- z/VM HELP files (included in the z/VM system image)

IBM Knowledge Center

z/VM product documentation is available in IBM Knowledge Center - z/VM (www.ibm.com/support/knowledgecenter/SSB27U). IBM Knowledge Center is the central location for finding and organizing information about IBM products. You can search for products and terms, and you can change the search scope and filter to search all of IBM Knowledge Center, a set of product versions, or just within a single version.

IBM Publications Center

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

z/VM Collection

The *IBM Online Library: z/VM Collection*, SK5T-7054, includes product information libraries for z/VM and current IBM licensed programs that run on z/VM. The *z/VM Collection* is available as a downloadable zip file that can be obtained from IBM Publications Center (www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss).

z/VM Internet Library

The IBM: z/VM Internet Library (www.ibm.com/vm/library/) provides links to additional sources of z/VM information, including:

- z/VM program directories
- z/VM data areas, control blocks, and monitor records
- IBM Redbooks® publications
- White papers, consultant reports, and performance reports
- Data sheets and brochures
- Reference guides

- Journals, bulletins, and newsletters

Abstracts of the z/VM base product publications

The z/VM base product publications include z/VM publications and other IBM publications that provide related information. The base product publications are grouped into the following task-related categories:

- “System overview”
- “Installation, migration, and service”
- “Planning and administration” on page 71
- “Customization and tuning” on page 74
- “Operation and use” on page 75
- “Application programming” on page 76
- “Diagnosis” on page 81

Notes:

1. A publication identified as “z/VM V6.4” contains new or changed information to support z/VM V6.4. z/VM publications with prior release designations also support z/VM V6.4 but do not contain z/VM V6.4 updates.
2. The updates in a new edition of a non-z/VM publication might not be z/VM related.
3. The indicated edition of a non-z/VM publication might not be the latest edition. Functions described in newer editions might not be supported in z/VM V6.4.

System overview

These publications provide information about z/VM capabilities and requirements, terminology, and specific license terms.

z/VM V6.4 General Information, GC24-6193-12

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

- Product overview
- What is new or changed in the latest z/VM release
- Hardware and software requirements
- Guide to the product information
- IBM servers, guest operating systems, and devices supported by current z/VM releases

z/VM V6.4 Glossary, GC24-6195-05

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

z/VM V6.4 License Information, GC24-6200-11

z/VM: License Information is a supplement to the IBM International Program License Agreement, and includes terms specific to z/VM.

Installation, migration, and service

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

z/VM V6.4 Installation Guide, GC24-6246-03

z/VM: Installation Guide contains step-by-step procedures for installing a z/VM system. The procedures feature an automated installation process using a panel interface to load a prebuilt z/VM System Image packaged in

DASD Dump/Restore (DDR) format on magnetic tape or packaged in binary format on DVD. Installation worksheets are included that you can use to plan your installation.

z/VM V6.4 Migration Guide, GC24-6201-10

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

- An introduction to migration
- Descriptions of system changes and enhancements that you should be aware of before migrating
- Identification of specific external interfaces that have changed, with changes classified as upwardly compatible or incompatible
- Guidance for some migration tasks you might need to perform, such as converting from system definition macros to system configuration files, or migrating your system directory

z/VM V6.4 Service Guide, GC24-6247-02

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

z/VM V6.4 VMSES/E Introduction and Reference, GC24-6243-03

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

Planning and administration

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

z/VM V6.2 CMS File Pool Planning, Administration, and Operation, SC24-6167-01

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

z/VM V6.4 CMS Planning and Administration, SC24-6171-03

z/VM: CMS Planning and Administration provides information on the Conversational Monitor System (CMS) component of z/VM, including:

- 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 V6.4 Connectivity, SC24-6174-06

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

z/VM V6.4 CP Planning and Administration, SC24-6178-10

z/VM: CP Planning and Administration provides information about the Control Program (CP) component of z/VM:

- System planning and administration
This section includes information on creating and updating the system configuration file, setting up service virtual machines, using the stand-alone dump utility, using z/VM HiperDispatch, and other topics.
- User planning and administration
This section includes information on creating and updating a user directory and redefining command privilege classes.
- Storage planning and administration
This section includes information on managing real storage, allocating DASD space, sharing DASD, defining and managing SCSI FCP disks, and other topics.
- Single system image clusters planning and administration
This section includes information on setting up a z/VM single system image (SSI) cluster, preparing for guest relocations in an SSI cluster, and other topics.

z/VM V6.4 Enabling z/VM for OpenStack (Support for OpenStack Liberty Release), SC24-6251-01

z/VM: Enabling z/VM for OpenStack (Support for OpenStack Liberty Release) provides guidance to z/VM customers who wish to configure a product that includes the z/VM plug-in for enabling OpenStack Liberty for z/VM. This document contains planning and configuration information.

z/VM V6.4 Getting Started with Linux on z Systems, SC24-6194-05

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

- z/VM basics
- Planning for Linux virtual servers
- Changing the z/VM system configuration
- Configuring the Directory Maintenance Facility
- Configuring TCP/IP for z/VM
- Creating and cloning Linux virtual machines

- Setting up basic system automation
- Performing runtime tasks
- Monitoring performance and capacity
- Servicing z/VM
- Preparing for live guest relocation
- Using FTP to install Linux from the hardware management console

z/VM V6.4 Group Control System, SC24-6196-03

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

z/VM V6.4 I/O Configuration, SC24-6198-06

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

z/VM V6.4 Running Guest Operating Systems, SC24-6228-05

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 V6.3 Getting Started with Linux on System z®.

z/VM V6.3 Saved Segments Planning and Administration, SC24-6229-02

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

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

z/VM V6.3 Secure Configuration Guide, SC24-6230-05

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

z/VM V6.4 TCP/IP LDAP Administration Guide, SC24-6236-03

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

z/VM V6.4 TCP/IP Planning and Customization, SC24-6238-08

z/VM: TCP/IP Planning and Customization provides information to help you plan and set up TCP/IP networks on your z/VM system. It describes how

to define and configure the virtual machines, servers, and applications available in TCP/IP for z/VM. It also describes how to customize and tune TCP/IP for your specific needs.

z/OS and z/VM: Hardware Configuration Manager User's Guide, SC34-2670-04

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.

Open Systems Adapter-Express Customer's Guide and Reference, SA22-7935-17

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

Open Systems Adapter-Express Integrated Console Controller User's Guide, SA22-7990-01

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 z Systems servers.

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

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

Open Systems Adapter-Express3 Integrated Console Controller Dual-Port User's Guide, SA23-2266-02

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

Customization and tuning

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

z/VM V6.4 CP Exit Customization, SC24-6176-03

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 and controlling dynamically loaded routines
- Defining CP commands and DIAGNOSE codes
- Defining and using CP exit points
- Creating and using CP message repositories

z/VM V6.4 Performance, SC24-6208-08

z/VM: Performance contains information about the planning, managing, measuring, and tuning considerations needed to obtain optimum z/VM system performance. It provides an overview of z/VM system

characteristics and discusses performance methodology, planning measures, monitoring facility tools, and tuning actions.

Operation and use

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

z/VM V6.4 CMS Commands and Utilities Reference, SC24-6166-04

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 V6.1 CMS Primer, SC24-6172-00

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

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

z/VM V6.4 CMS User's Guide, SC24-6173-03

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

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

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

z/VM V6.4 CP Commands and Utilities Reference, SC24-6175-09

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 V6.4 System Operation, SC24-6233-05

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 (SAPL)
- 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 V6.4 TCP/IP User's Guide, SC24-6240-07

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

- Transferring files
- Sending electronic mail

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

z/VM V6.4 Virtual Machine Operation, SC24-6241-02

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

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

z/VM V6.1 XEDIT Commands and Macros Reference, SC24-6244-00

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

z/VM V6.1 XEDIT User's Guide, SC24-6245-00

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

Application programming

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

z/VM V6.4 CMS Application Development Guide, SC24-6162-02

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

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

z/VM V6.4 CMS Application Development Guide for Assembler, SC24-6163-02

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 V6.4 CMS Application Multitasking, SC24-6164-01

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

z/VM V6.4 CMS Callable Services Reference, SC24-6165-03

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 V6.4 CMS Macros and Functions Reference, SC24-6168-02

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 V6.4 CMS Pipelines User's Guide and Reference, SC24-6252-00

z/VM: CMS Pipelines User's Guide and Reference describes the general concepts of CMS Pipelines and how to use CMS Pipelines, and provides reference information on the CMS Pipelines stages, subcommands, and assembler macros.

z/VM V6.4 CP Programming Services, SC24-6179-09

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 V6.1 CPI Communications User's Guide, SC24-6180-00

z/VM: CPI Communications User's Guide provides step-by-step instructions for using Systems Application Architecture® (SAA) Common Programming Interface (CPI) Communications to write communications programs to run in the CMS environment. Sample programs written in REXX/VM show how to use SAA CPI Communications calls and the CMS extensions to CPI Communications.

z/VM V6.1 Enterprise Systems Architecture/Extended Configuration Principles of Operation, SC24-6192-00

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

z/VM V6.2 Language Environment User's Guide, SC24-6199-01

z/VM: Language Environment User's Guide provides information for using Language Environment for z/VM, which is based on z/OS Language Environment. It provides information unique to the z/VM platform and is

intended to be used in conjunction with the z/OS Language Environment documentation (included in the z/VM library). See:

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

z/VM V6.1 OpenExtensions Advanced Application Programming Tools, SC24-6202-00

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

z/VM V6.4 OpenExtensions Callable Services Reference, SC24-6203-02

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

z/VM V6.4 OpenExtensions Commands Reference, SC24-6204-01

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 V6.1 OpenExtensions POSIX Conformance Document, GC24-6205-00

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

z/VM V6.1 OpenExtensions User's Guide, SC24-6206-00

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

z/VM V6.4 Program Management Binder for CMS, SC24-6211-04

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

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

z/VM V6.4 Reusable Server Kernel Programmer's Guide and Reference, SC24-6220-02

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

z/VM V6.4 REXX/VM Reference, SC24-6221-02

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 V6.4 REXX/VM User's Guide, SC24-6222-01

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 V6.4 Systems Management Application Programming, SC24-6234-11

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

z/VM V6.4 TCP/IP Programmer's Reference, SC24-6239-05

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

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

Common Programming Interface Communications Reference, SC26-4399-09

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

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

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

z/OS: IBM Tivoli Directory Server Plug-in Reference for z/OS, SA76-0169-00

z/OS: IBM Tivoli Directory Server Plug-in Reference for z/OS describes the application services routines that can be used to create plug-ins to extend

the capabilities of the z/VM TCP/IP LDAP server. For information about creating LDAP server plug-ins on z/VM, see z/VM V6.3 TCP/IP Programmer's Reference.

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

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

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

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

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

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

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

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

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

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

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

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

z/OS MVS Program Management: Advanced Facilities, SA23-1392-01

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, SA23-1393-01

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

Diagnosis

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

z/VM V6.4 CMS and REXX/VM Messages and Codes, GC24-6161-03

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 V6.4 CP Messages and Codes, GC24-6177-08

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 V6.4 Diagnosis Guide, GC24-6187-06

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

z/VM V6.1 Dump Viewing Facility, GC24-6191-00

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

For information on analyzing CP dumps, see z/VM V6.3 VM Dump Tool.

z/VM V6.4 Other Components Messages and Codes, GC24-6207-03

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.

For messages issued by the Language Environment component of z/VM, see z/OS: Language Environment Run-Time Messages.

z/VM V6.4 TCP/IP Diagnosis Guide, GC24-6235-05

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

z/VM V6.4 TCP/IP Messages and Codes, GC24-6237-08

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

z/VM V6.3 VM Dump Tool, GC24-6242-03

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 stand-alone dumps, CP abend dumps, and virtual machine dumps of a CP system. This document describes the VM Dump Tool subcommands and macros that allow you to display, locate, and format dump data interactively.

z/OS and z/VM: Hardware Configuration Definition Messages, SC34-2668-03

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

Abstracts of the z/VM optional feature publications

There is a set of publications for each of the following z/VM optional features:

- “DFSMS/VM publications”
- “Directory Maintenance Facility for z/VM publications” on page 83
- “IBM Wave for z/VM publications” on page 83
- “Performance Toolkit for VM publications” on page 84
- “RACF Security Server for z/VM publications” on page 84
- “RSCS Networking for z/VM publications” on page 85

There are no separate publications for the IBM z/VM Single System Image Feature (VMSSI). For an introduction to SSI clusters, see *z/VM: CP Planning and Administration*.

Notes:

1. A publication identified as “z/VM V6.4” contains new or changed information to support z/VM V6.4. z/VM publications with prior release designations also support z/VM V6.4 but do not contain updates for z/VM V6.4.

DFSMS/VM publications

These publications provide information about the DFSMS/VM optional feature. DFSMS/VM controls data and storage resources.

z/VM V6.1 DFSMS/VM Customization, SC24-6181-00

z/VM: DFSMS/VM Customization provides information on customizing DFSMS/VM to meet the needs of your installation. Topics include:

- Customizing the DFSMS/VM control file
- Using installation-wide exits
- Customizing the Interactive Storage Management Facility (ISMF)
- Performance and tuning

z/VM V6.1 DFSMS/VM Diagnosis Guide, GC24-6182-00

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

z/VM V6.4 DFSMS/VM Messages and Codes, GC24-6183-01

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

z/VM V6.1 DFSMS/VM Planning Guide, SC24-6184-00

z/VM: DFSMS/VM Planning Guide provides information on planning for the installation and use of DFSMS/VM. It can help you make decisions about:

- Using the DFSMS/VM storage management functions in your system
- Authorizing storage administrators
- Setting up the server machines
- Generating accounting records
- Backing up secondary storage

z/VM V6.4 DFSMS/VM Removable Media Services, SC24-6185-03

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 V6.1 DFSMS/VM Storage Administration, SC24-6186-00

z/VM: DFSMS/VM Storage Administration provides conceptual information on System Managed Storage (SMS) and guides you through the DFSMS/VM storage management and minidisk management operations. It also contains a DFSMS/VM command reference, an ISMF reference, exit information, a DFSMS/VM sample environment, and other useful reference information.

Directory Maintenance Facility for z/VM publications

These publications provide information about the Directory Maintenance Facility for z/VM (DirMaint) optional feature. DirMaint provides interactive facilities for managing the z/VM user directory.

z/VM V6.4 Directory Maintenance Facility Commands Reference, SC24-6188-05

z/VM: Directory Maintenance Facility Commands Reference provides detailed reference information on DirMaint commands. Most of the DirMaint commands have the same name and format as the z/VM directory statements they support. This document also describes the additional utilities that DirMaint provides for minidisk management and command security.

z/VM V6.4 Directory Maintenance Facility Messages, GC24-6189-04

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

z/VM V6.4 Directory Maintenance Facility Tailoring and Administration Guide, SC24-6190-05

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.

IBM Wave for z/VM publications

These publications provide information about the IBM Wave for z/VM (IBM Wave) optional feature. IBM Wave is a provisioning and productivity management solution for simplifying the control and use of virtual Linux servers and z/VM.

IBM Wave for z/VM V1R2 Administration and Customization, SC27-6118-08

IBM Wave for z/VM: Administration and Customization describes how to maintain and customize IBM Wave to suit the specific requirements. Topics in this publication include:

- Installation prerequisites, installation concepts, and server installation
- Site management and other administrative actions
- Security subsystem and user management
- Parameter settings for customization
- User exits

IBM Wave for z/VM V1R2 User Guide and Reference, SC27-6119-08

IBM Wave for z/VM: User Guide and Reference describes how to use IBM Wave. Topics in this publication include:

- IBM Wave viewers
- Available tasks for the objects IBM Wave manages
- Reporting subsystem

- Script management subsystem
- Command line interface (CLI)
- Examples and samples

Performance Toolkit for VM publications

These publications provide information about the Performance Toolkit for VM optional feature. Performance Toolkit provides tools for analyzing z/VM and Linux performance data.

z/VM V6.4 Performance Toolkit Guide, SC24-6209-07

z/VM: Performance Toolkit Guide describes how the Performance Toolkit for VM can assist system operators and systems programmers or analysts in the following areas:

- **System operator console operation in full screen mode.** Features are provided to facilitate the operation of VM systems, thereby improving operator efficiency and productivity.
- **Performance monitoring on VM.** An enhanced real-time performance monitor allows systems programmers to monitor system performance and to analyze bottlenecks.

z/VM V6.4 Performance Toolkit Reference, SC24-6210-07

z/VM: Performance Toolkit Reference describes the subcommands and screens provided with the Performance Toolkit for VM. It also describes the Performance Toolkit messages.

RACF Security Server for z/VM publications

These publications provide information about the RACF Security Server for z/VM optional feature. RACF provides data security for an installation by controlling access.

z/VM V6.4 RACF Security Server Auditor's Guide, SC24-6212-02

z/VM: RACF Security Server Auditor's Guide describes the role of the RACF auditor and explains the auditing tools that RACF provides. Reports on system and resource use can provide the auditor with information about the basic system-security environment of an installation.

z/VM V6.4 RACF Security Server Command Language Reference, SC24-6213-03

z/VM: RACF Security Server Command Language Reference describes the syntax and functions of RACF commands.

z/VM V6.3 RACF Security Server Diagnosis Guide, GC24-6214-02

z/VM: RACF Security Server Diagnosis Guide explains how to diagnose problems that might occur in RACF.

z/VM V6.3 RACF Security Server General User's Guide, SC24-6215-01

z/VM: RACF Security Server General User's Guide teaches general users how to use RACF to protect their own minidisks, SFS files, SFS directories, and other general resources. It contains an introduction to RACF and sections that guide the user through basic security tasks.

z/VM V6.4 RACF Security Server Macros and Interfaces, SC24-6216-03

z/VM: RACF Security Server Macros and Interfaces describes the syntax and functions of RACF customization macros. It also contains information about other interfaces such as RACF SMF records and subcodes of CP DIAGNOSE code X'A0'.

For information on the RACROUTE macro and the independent RACF system macros, see *z/VM V6.2 Security Server RACROUTE Macro Reference*.

z/VM V6.4 RACF Security Server Messages and Codes, GC24-6217-04

z/VM: RACF Security Server Messages and Codes describes RACF messages, abend codes, manager return codes, and utility return codes.

z/VM V6.4 RACF Security Server Security Administrator's Guide, SC24-6218-05

z/VM: RACF Security Server Security Administrator's Guide provides information to help the RACF security administrator to:

- Plan how to use RACF, which includes deciding which resources are protected and which users and groups are known to RACF
- Perform daily administration tasks, such as giving users access to the system (assigning user IDs and passwords) and giving users access to system resources or functions
- Coordinate with other administrators, such as the tape librarian

z/VM V6.4 RACF Security Server System Programmer's Guide, SC24-6219-05

z/VM: RACF Security Server System Programmer's Guide provides information for system programmers or installation personnel responsible for:

- Maintaining RACF databases
- Writing, testing, and installing RACF exits
- Modifying RACF to satisfy an installation's particular needs

z/VM V6.4 Security Server RACROUTE Macro Reference, SC24-6231-03

z/VM: Security Server RACROUTE Macro Reference describes the full-function RACROUTE external security interface for z/VM, the requests that can be invoked by it, their respective syntax, and related information. This document is intended for programmers who are writing applications that need to invoke RACF (or another external security manager) from z/VM, or programmers who are writing an external security manager for z/VM.

RSCS Networking for z/VM publications

These publications provide information about the RSCS Networking for z/VM optional feature. RSCS enables users on a z/VM system to send messages, files, commands, and jobs to other users within a network.

z/VM V6.1 RSCS Networking Diagnosis, GC24-6223-00

z/VM: RSCS Networking Diagnosis describes the functions of RSCS and RSCS diagnostic aids and facilities. This document is intended for people who are responsible for isolating and diagnosing any problems that might occur in RSCS operation.

z/VM V6.1 RSCS Networking Exit Customization, SC24-6224-00

z/VM: RSCS Networking Exit Customization describes the RSCS exit facilities. This document describes how you can customize RSCS processing to meet the functional requirements of your facility.

z/VM V6.4 RSCS Networking Messages and Codes, GC24-6225-02

z/VM: RSCS Networking Messages and Codes provides explanations of the messages and abend codes for RSCS and RSCS facilities.

z/VM V6.4 RSCS Networking Operation and Use, SC24-6226-02

z/VM: RSCS Networking Operation and Use explains how to operate RSCS and how to use it to enable your z/VM system to communicate with other systems. This document describes how to send data across the network, how to define and manage links, and how to automate your network. It explains how to operate and use the RSCS server, RSCS domain name server, RSCS dynamic authorization, and the RSCS Data Interchange

Manager. This document also describes the syntax and functions of the RSCS commands and link operational parameters and the commands for the other RSCS facilities.

z/VM V6.4 RSCS Networking Planning and Configuration, SC24-6227-03

z/VM: RSCS Networking Planning and Configuration provides a high-level description of RSCS and provides information to help you plan for your RSCS network. This document describes planning tasks for RSCS configuration, operation, administration, and diagnosis. It also describes the syntax and functions of the configuration file statements for the various types of RSCS servers.

Appendix A. IBM servers supported by z/VM

This topic provides information about the IBM servers supported by current releases of z/VM. For guest considerations, see Appendix B, “IBM operating systems supported as guests of z/VM,” on page 93.

The following general notes apply to the support matrix:

- z/VM V6.2 or later can be selected as the hypervisor for IBM LinuxONE systems:
 - IBM LinuxONE Emperor is based on the IBM z13.
 - IBM LinuxONE Rockhopper is based on the IBM z13s.

Note: Prior to the general availability of the z13s, Rockhopper was based on the IBM zEnterprise BC12.

Additional z/VM and other APARs might be required to support these servers, as identified in Table 2 on page 88.

- The Processor Resource/Systems Manager™ (PR/SM) hardware feature supports the creation of multiple logical partitions (LPARs) on a single central processor complex (CPC), dividing server resources across the LPARs. Each LPAR supports an independent operating system. z/VM supports the following LPAR types (modes):
 - ESA/390 (on z13 and z13s and earlier servers)
 - General (on later servers)
 - LINUX only
 - z/VM (available only on the IBM System z10™ and later servers)
- z/VM supports guest use of IBM mainframe specialty processors:
 - IBM Integrated Facility for Linux (IFL)
 - IBM Internal Coupling Facility (ICF)
 - IBM zEnterprise Application Assist Processor (zAAP)

Note: zAAPs are not supported on the IBM z13 and later servers.

- IBM z Integrated Information Processor (zIIP)
- A LINUX only LPAR, whether the allocated processors are IFL processors or general purpose processors, will not support any of the IBM traditional operating systems (such as z/OS, z/TPF, or z/VSE). Only Linux on z Systems, OpenSolaris, or z/VM can run in a LINUX only LPAR.
- An IFL processor can be defined only in a LINUX only or z/VM LPAR.
- z/VM may run on IFL processors only if, on the IFL processors, z/VM is being used exclusively to run (1) Linux or OpenSolaris workloads and, (2) if required, z/VM applications in support of those Linux or OpenSolaris workloads.
- z/VM can exploit the ESCON Multiple Image Facility (EMIF) when running in an LPAR.
- To use the integrated 3270 console task on the Hardware Management Console as a system operator console for z/VM, the server must be running at the current Licensed Internal Code (LIC) level with the Hardware Management Console at V1.8.0 or later.
- Before installing z/VM, refer to the Preventive Service Planning (PSP) bucket for your z Systems server for the minimum MCL level and any required updates.

Supported Servers

Server	PSP bucket
IBM z13 and IBM LinuxONE Emperor	Upgrade 2964DEVICE, Subset 2964/ZVM
IBM z13s and IBM LinuxONE Rockhopper	Upgrade 2965DEVICE, Subset 2965/ZVM
IBM zEnterprise EC12	Upgrade 2827DEVICE, Subset 2827/ZVM
IBM zEnterprise BC12	Upgrade 2828DEVICE, Subset 2828/ZVM
IBM zEnterprise 196	Upgrade 2817DEVICE, Subset 2817/ZVM
IBM zEnterprise 114	Upgrade 2818DEVICE, Subset 2818/ZVM
IBM System z10 Enterprise Class	Upgrade 2097DEVICE, Subset 2097/ZVM
IBM System z10 Business Class	Upgrade 2098DEVICE, Subset 2098/ZVM
IBM System z9 [®] Enterprise Class	Upgrade 2094DEVICE, Subset 2094/ZVM
IBM System z9 Business Class	Upgrade 2096DEVICE, Subset 2096/ZVM
IBM eServer [™] zSeries 990	Upgrade 2084DEVICE, Subset 2084/ZVM
IBM eServer zSeries 890	Upgrade 2086DEVICE, Subset 2086/ZVM
IBM eServer zSeries 900	Upgrade 2064DEVICE, Subset 2064/ZVM
IBM eServer zSeries 800	Upgrade 2066DEVICE, Subset 2066/ZVM

The PSP buckets can be found by logging on to IBM Resource Link (www.ibm.com/servers/resourcelink/). A Resource Link user ID and password are required.

- The inclusion of a server in this matrix does not imply that the server is available in all countries.
- For information about IBM servers and z/VM releases not listed in this matrix, see the previous editions of this document.

Table 2. Server support

Server	z/VM	Support requirements
z13 (or LinuxONE Emperor) [Machine Type 2964]	V6.4	1, 2, 3, 4, 5, 6, 11
	V6.3	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11
	V6.2	1, 2, 3, 4, 5, 6, 7, 10, 11
z13s (or LinuxONE Rockhopper) [Machine Type 2965]	V6.4	1, 2, 3, 4, 5, 6, 11
	V6.3	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11
	V6.2	1, 2, 3, 4, 5, 6, 7, 10, 11
General support notes for z13 and z13s (also see support requirements): <ul style="list-style-type: none"> • Supported LPAR modes: ESA/390, LINUX only, z/VM. • Up to 60 LPARs. • Up to 1 TB (2 TB for z/VM V6.4) of real storage (memory) per LPAR (z/VM image) and up to 1 TB of virtual memory for a single guest. The total virtual memory in use by all guests can exceed the amount of real memory. • Up to 64 logical processors per z/VM image and up to 64 virtual processors per virtual machine in each z/VM image. • z/Architecture guest use of Crypto Express5S. • The IBM zEnterprise Application Assist Processor (zAAP) CPU type is not supported either in real hardware or in the LPAR configuration. Similarly, z/VM will not allow guests to define this CPU type. 		

Table 2. Server support (continued)

Server	z/VM	Support requirements
Support requirements for z13 and z13s:		
If an APAR is required for a z/VM release to support a function on a particular server, that APAR is also required for the same z/VM release to support the same function on a newer server.		
1		EREP support requires APAR VM65495 for z13, VM65704 for z13s.
2		CMS IOCP support requires: APAR VM65568 for z13; VM65736 for z13 (Driver D27), z13s, and LinuxONE..
3		HCD support requires APAR VM65489 for z13; VM65729 for z13 (Driver D27), z13s, and LinuxONE.
4		HCM support requires APAR VM64437 for z13; VM64844 for z13 (Driver D27), z13s, and LinuxONE.
5		HLASM support requires APAR PM79901.
6		OSA/SF support requires APARs OA37060 and OA38418 for OSA-Express4S.
7		Support for guests to exploit zEC12 function supported by z/VM on the z13 requires APAR VM65577. This APAR also provides support for Crypto Express5S and enhanced domain support on the z13. Performance Toolkit support for the z13 requires APAR VM65527. DirMaint support for enhanced domain support requires APAR VM65588.
8		Support for Multi-VSwitch Link Aggregation requires APARs VM65583 (CP) and PI21053 (TCP/IP). Systems management API support requires APAR VM65670. Performance Toolkit support requires APAR VM65528. Also see "Multi-VSwitch Link Aggregation support hardware requirements" on page 23.
9		z/VM host exploitation support for simultaneous multithreading (SMT) requires APARs VM65586 and VM65696. These APARs also provide z/VM support for up to 64 logical processors on the z13: 64 cores with multithreading disabled, or 32 cores (up to 2 threads per core) with multithreading enabled. Stand-alone dump support for SMT requires APARs VM65676 and VM65677. Performance Toolkit support for SMT requires APAR VM65529.
10		Support for z/VM V6.3 and V6.2 guests to exploit z/VM supported z13 functions on the z13 (Driver D27) and the z13s requires APAR VM65716. This APAR also provides support for the following additional functions on z13 (Driver D27) and z13s: <ul style="list-style-type: none"> • LPAR group absolute capacity capping • Dynamic memory management • Shared Memory Communications - Direct Memory Access (SMC-D) • Regional Crypto Enablement (RCE) adapters (also requires APAR VM65577) Performance Toolkit support for SMC-D requires APAR VM65698.
11		Support for assigning UIDs to real PCI functions on z/VM V6.2, V6.3, and V6.4 requires APARs VM65865 (CP), VM65827 (HCD), and VM65201 (HCM). IOCP support on V6.2 and V6.3 requires APAR VM65817. (IOCP support is included in the V6.4 base.)
zEnterprise EC12 (zEC12) [Machine Type 2827]	V6.4	3, 7, 8
	V6.3	3, 4, 5, 6, 7, 8, 12
	V6.2	2, 3, 4, 5, 6, 7, 8, 10, 11
	V5.4	2, 3, 4, 5, 6, 7, 8, 9

Supported Servers

Table 2. Server support (continued)

Server	z/VM	Support requirements
zEnterprise BC12 (zBC12) [Machine Type 2828]	V6.4	3, 7, 8
	V6.3	3, 4, 5, 6, 7, 8, 12
	V6.2	3, 4, 5, 6, 7, 8, 10, 11
	V5.4	3, 4, 5, 6, 7, 8, 9
zEnterprise 196 (z196) [Machine Type 2817]	V6.4	8
	V6.3	8, 13
	V6.2	3, 8, 10, 11, 13
	V5.4	1, 3, 4, 5, 6, 7, 8, 9
zEnterprise 114 (z114) [Machine Type 2818]	V6.4	8
	V6.3	8, 13
	V6.2	3, 8, 10, 11, 13
	V5.4	1, 3, 4, 5, 6, 7, 8, 9
General support notes for zEnterprise servers (also see support requirements): <ul style="list-style-type: none"> • Supported LPAR modes: ESA/390, LINUX only, z/VM. • Up to 60 LPARs. • Up to 1 TB of real storage (memory) per LPAR (z/VM image) and up to 1 TB of virtual memory for a single guest. The total virtual memory in use by all guests can exceed the amount of real memory. • Up to 32 logical processors per z/VM image and up to 64 virtual processors per virtual machine in each z/VM image. • Guest use of Crypto Express3. • z/Architecture guest use of Crypto Express4S (zEC12 and zBC12 only). 		

Table 2. Server support (continued)

Server	z/VM	Support requirements
Support requirements for zEnterprise servers:		
If an APAR is required for a z/VM release to support a function on a particular server, that APAR is also required for the same z/VM release to support the same function on a newer server.		
1		Support for guest exploitation of selected z196 or z114 features requires APARs VM64798, VM64879, VM64881, and VM64947.
2		Support for guest exploitation of selected zEC12 features and Crypto Express4S requires APAR VM65007. Performance Toolkit support requires APAR VM65046.
3		EREP support requires: APAR VM64807 for z196; VM64807 and VM64928 for z114; VM65130 for zEC12; VM65130 and VM65279 for zEC12 (Driver D15F) and zBC12.
4		CMS IOCP support requires: APAR VM64799 for z196; VM64940 for z114; VM65131 for zEC12; VM65131 and VM65442 for zEC12 (Driver D15F) and zBC12.
5		HCD support requires: APAR VM64672 for z196; VM64856 for z114; VM65047 for zEC12; VM65047 and VM65239 for zEC12 (Driver D15F) and zBC12.
6		HCM support requires: APAR VM64747 for z196; VM64880 for z114; VM65236 for zEC12 (Driver D15F) and zBC12.
7		HLASM support requires: APAR PK97799 for z196 and z114; PM49761 for zEC12 and zBC12.
8		OSA/SF support requires: APAR OA37060 for OSA-Express4S on z196 and z114; APARs OA37060 and OA38418 for OSA-Express4S on zEC12 and zBC12.
9		Support for protected key CPACF requires APAR VM64793.
10		Guest support for High Performance FICON requires APARs VM65041, VM65144, and VM65180. Performance Toolkit requires APAR VM65044. Also see “zHPF guest support hardware requirements” on page 23.
11		Support for the z/VM virtual switch HiperSockets bridge requires APARs VM65042 and PM46988. Performance Toolkit support requires APAR VM65044. On z196 and z114, Driver D93 is required (see the PSP bucket for the minimum MCL level).
12		Guest support for zEDC Express and 10GbE RoCE Express features requires APARs VM65417 (CP), VM65572 (CP TRACERED support) and PI20509 (TCP/IP). Performance Toolkit support requires APAR VM65399. Also see “PCIe support hardware requirements” on page 23.
13		The PTF for APAR VM65207 must be applied to z/VM 6.2 or z/VM 6.3 when z/VM 6.4 is a guest of either release and will host a SET MACHINE Z guest.
System z10 Enterprise Class (z10 EC) [Machine Type 2097]	V6.3	
	V6.2	6
	V5.4	2, 3, 4, 5
System z10 Business Class (z10 BC) [Machine Type 2098]	V6.3	
	V6.2	6
	V5.4	1, 2, 3, 4, 5

Supported Servers

Table 2. Server support (continued)

Server	z/VM	Support requirements
General support notes for System z10 servers (also see support requirements): <ul style="list-style-type: none"> Supported LPAR modes: ESA/390, LINUX only, z/VM. Up to 60 LPARs. Up to 1 TB of real storage (memory) per LPAR (z/VM image) and up to 1 TB of virtual memory for a single guest. The total virtual memory in use by all guests can exceed the amount of real memory. Up to 32 real processors per z/VM image and up to 64 virtual processors per virtual machine in each z/VM image. Guest use of Crypto Express2 and Crypto Express3. 		
Support requirements for System z10 servers: <p>If an APAR is required for a z/VM release to support a function on a particular server, that APAR is also required for the same z/VM release to support the same function on a newer server.</p> <ol style="list-style-type: none"> 1 EREP support for the z10 BC requires APAR VM64475. 2 CMS IOCP support for the z10 BC and the enhanced z10 EC requires APAR VM64474. 3 HCD/HCM support for the z10 BC and the enhanced z10 EC requires APAR VM64410. 4 OSA/SF support for the z10 BC and the enhanced z10 EC requires APAR OA26286. 5 Support for protected key CPACF requires APAR VM64793. 6 Guest support for High Performance FICON requires APARs VM65041, VM65144, and VM65180. Performance Toolkit requires APAR VM65044. Also see “zHPF guest support hardware requirements” on page 23. 		
System z9 Enterprise Class (z9 EC) [Machine Type 2094]	V5.4	
System z9 Business Class (z9 BC) [Machine Type 2096]	V5.4	
General support notes for System z9 servers: <ul style="list-style-type: none"> Supported LPAR modes: ESA/390, LINUX only. z/VM supports guest use of the configurable Crypto Express2. 		
zSeries 990 (z990) [Machine Type 2084]	V5.4	
zSeries 890 (z890) [Machine Type 2086]	V5.4	
zSeries 900 (z900) [Machine Type 2064]	V5.4	
zSeries 800 (z800) [Machine Type 2066]	V5.4	
General support notes for zSeries servers: <ul style="list-style-type: none"> Supported LPAR modes: ESA/390, LINUX only, Basic (z900 and z800 only). z/VM supports guest use of the PCI Cryptographic Accelerator, PCI Cryptographic Coprocessor, PCIX Cryptographic Coprocessor, and Crypto Express2 features, where available. z/VM supports guest use of the CMOS Cryptographic Processor Feature (CCF) for the z900 and z800, except in a LINUX only LPAR. 		

Appendix B. IBM operating systems supported as guests of z/VM

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

In general, an operating system is supported as a guest of z/VM only where support has been announced for that operating system to run on the server. For example, if an operating system is not supported to run in a LINUX only LPAR, that operating system is not supported as a guest of z/VM running in a LINUX only LPAR.

This topic 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 87.

The following general notes apply to the support matrix:

- All listed guests are supported in ESA virtual machines. XA virtual machines are supported for compatibility and are functionally equivalent to ESA virtual machines.
- Guests may be uniprocessor or multiprocessor.
- For information about product releases not listed, see the previous editions of this document.

Table 3. Guest support

Guest	z/VM host	Support notes
Linux on z Systems	V6.4	1
	V6.3	1
	V6.2	1
	V5.4	1
Support notes for Linux guests:		
1	IBM z Systems: Linux on z Systems - Resources - Tested platforms (www.ibm.com/systems/z/os/linux/resources/testedplatforms.html) shows IBM tested and supported Linux environments on IBM servers. A supported Linux distribution is also supported as a guest of z/VM if the z/VM release supports that server.	
z/OS V1.12, V1.13, V2.1 or later	V6.4	1, 2
	V6.3	1, 2
	V6.2	1, 2
	V5.4	1, 2
Support notes for z/OS guests:		
1	Exploitation of hardware function by z/OS might require specific levels of z/VM. For details, see <i>z/OS: Planning for Installation</i> , GA22-7504.	
2	z/VM supports only virtual coupling.	

Supported Guests

Table 3. Guest support (continued)

Guest	z/VM host	Support notes
z/TPF V1.1	V6.4	1
	V6.3	1
	V6.2	1
	V5.4	1
Support 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. Use of the Emulation Program for 37xx communication controllers is not supported.		
z/VSE V4.3 or later	V6.4	
	V6.3	
	V6.2	
	V5.4	
z/VM V6.4	V6.4	2, 3
	V6.3	2, 3, 6, 7
	V6.2	2, 3, 7
	V5.4	2, 3, 7
z/VM V6.3	V6.4	2, 3
	V6.3	2, 3
	V6.2	2, 3
	V5.4	1, 2, 3
z/VM V6.2	V6.4	2, 3
	V6.3	2, 3
	V6.2	2, 3
	V5.4	1, 2, 3
z/VM V5.4	V6.4	2, 3, 4, 5
	V6.3	2, 3, 4, 5
	V6.2	2, 3, 4, 5
	V5.4	2, 3

Table 3. Guest support (continued)

Guest	z/VM host	Support notes
Support notes for guests:		
1		Supported only when the z/VM V5.4 host is running on a System z10 or later.
2		Running z/VM in a virtual machine (z/VM as a guest of z/VM, also known as “second-level”) is functionally supported but is intended only for testing purposes for the second-level z/VM system and its guests (called “third-level” guests). Performance of the third-level guests will not compare favorably to what would be achieved running those guest workloads second-level.
3		z/VM Guest Coupling Simulation will not work in the z/VM guest.
4		APAR VM64980 should be installed on the z/VM V5.4 guest system. The fix addresses a rare problem that can occur when V5.4 is running as a second-level system on V6.2 or later and a guest of the second-level system is using floating-point instructions. The problem occurs only if the guest executes certain floating-point instructions for which the guest has been told by STLFE facilities that bits are not available to it.
5		z/VM V5.4 is not supported as a guest if the z/VM V6.2 or later host is running on z13 or later.
6		If a z/VM V6.4 guest intends to exploit HyperPAV paging support, APAR VM65748 is required on the z/VM V6.3 host.
7		A z/VM V6.4 guest is supported only when the z/VM host is running on a z196 or z114 or later.

Supported Guests

Appendix C. IBM devices supported by z/VM

This topic provides information about the IBM devices supported by current releases of z/VM.

The device support matrix is intended to provide a quick reference for support of various devices on z/VM. You should also check the hardware device support list for your server to be sure the hardware supports a particular device. Other devices not found in the device support matrix might be supported; consult hardware announcements and device documentation for this information.

Device support might not be in the base z/VM release, but might have been delivered through the service stream by the PTF for an APAR.

Devices are fully supported by z/VM unless otherwise indicated. z/VM supports devices at any of the following three levels of function:

1. **Full support:** z/VM uses the device itself and makes it available to guests. z/VM will usually manage the device on behalf of all users.
2. **Dedicated to a guest or virtual machine:** z/VM does not use the device itself but enables guests or applications in virtual machines to use the device. There are usually no restrictions on the use of the device by a guest or application, except that the device may not be dedicated to more than one virtual machine at a time.
3. **Dedicated to a guest or virtual machine as an unsupported device:** z/VM does not use the device itself but enables guests or applications in virtual machines to use the device. The guest or application is completely responsible for the management and use of the device. Usage restrictions might apply; the major common usage restriction for this type of support is that the device cannot be used as the guest's IPL device.

The device support matrix might show that support for a device was not provided after a particular z/VM release. The following clarifications on the meaning of the programming support withdrawal might be helpful:

- Where programming support for a *real* device has been withdrawn as of some z/VM release, the corresponding *virtual* device remains supported in the same or later z/VM releases.
- There are some device types which can be defined as, or use many of the device attributes of, the real devices for which programming support is being terminated. z/VM support for these currently supported devices remains unchanged.

Additional device restrictions might apply. There might be exceptions to supported devices and releases depending on your specific configuration. See *z/VM: CP Planning and Administration* for your release and applicable hardware announcements.

It is sometimes possible to define a device differently depending on whether the definition is done using the system configuration file or dynamic I/O configuration support. For instance, a 3745 Communications Controller would be defined as a 3705 Communications Controller using the system configuration file but as a 3745 Communications Controller using dynamic I/O configuration. For information

Supported Devices

about the system configuration file, see *z/VM: CP Planning and Administration*. For information about dynamic I/O configuration, see *z/VM: I/O Configuration*.

The following general notes apply to the matrix:

- For information about IBM devices and z/VM releases not listed in this matrix, see the previous editions of this document.
- If a z/VM-related problem occurs with an OEM device that is equivalent to a supported IBM device listed in this section, IBM will fix the problem if it can be recreated with the corresponding IBM device. Customers should always contact the OEM manufacturer first about any problem encountered with an OEM device to find out if it is a known OEM problem.

Direct access storage devices (DASD)

Table 4. Device support – DASD

Device	z/VM support	Support notes
TotalStorage Enterprise Storage Server® (2105) [ESCON or FICON attached]	V5.4, V6.2, V6.3, V6.4	<p>z/VM automatically virtualizes the controller mode that a guest supports.</p> <p>Models 750 and 800 supported in 2105 native controller mode.</p> <p>Supported in emulation mode as a 3390 Model 3 or 6 Storage Control with the following DASD device types:</p> <ul style="list-style-type: none">• 3390 Model 2, 3, and 9 DASD (including large Model 9s known as Model 27 and Model 54)• 3390 Model 2 and 3 DASD in 3380 track compatibility mode <p>Only 3390 Model 3 or 9 is supported for installation of z/VM. Mixed 3390 models are not supported for installation of z/VM.</p> <p>Performance-oriented track level commands supported for guest use only.</p> <p>Parallel Access Volumes (PAVs) supported as minidisks for guests that exploit PAV architecture.</p> <p>DFSMS/VM FL221 support for 3390 Model 9 disks larger than 10017 cylinders requires APAR VM63004.</p> <p>For guest use of FlashCopy, the device must be dedicated to the guest or defined as a fullpack minidisk.</p> <p>PPRC supported for guest use only. Native use of PPRC requires ICKDSF V1.17.</p> <p>For ICKDSF support information, see the ICKDSF Program Directory.</p>

Table 4. Device support – DASD (continued)

Device	z/VM support	Support notes
TotalStorage Enterprise Storage Server (2105) [FCP attached]	V5.4, V6.2, V6.3, V6.4	<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>

Supported Devices

Table 4. Device support – DASD (continued)

Device	z/VM support	Support notes
System Storage DS8000 series (2107) [ESCON or FICON attached]	V5.4, V6.2, V6.3, V6.4	<p>z/VM automatically virtualizes the controller mode that a guest supports.</p> <p>Supported in 2107 native controller mode.</p> <p>Supported in emulation mode as a 3390 Model 3 or 6 Storage Control with the following DASD device types:</p> <ul style="list-style-type: none"> • 3390 Model 2, 3, and 9 DASD (including large Model 9s known as Model 27 and Model 54) • 3390 Model 2 and 3 DASD in 3380 track compatibility mode <p>Only 3390 Model 3 or 9 is supported for installation of z/VM. Mixed 3390 models are not supported for installation of z/VM.</p> <p>Dynamic volume expansion is supported.</p> <p>Extended address 3390-A devices are supported. z/VM V5.4 requires APARs VM64709 and VM64711. A 3390-A volume with more than 65,520 cylinders must be dedicated to a guest or defined as a fullpack minidisk.</p> <p>Performance-oriented track level commands supported for guest use only.</p> <p>Parallel Access Volumes (PAVs) supported as minidisks for guests that exploit PAV architecture.</p> <p>On FICON-attached DS8000, HyperPAV volumes supported as fullpack minidisks for guests that exploit this PAV architecture.</p> <p>DFSMS/VM FL221 support for 3390 Model 9 disks larger than 10017 cylinders requires APAR VM63004.</p> <p>For guest use of FlashCopy, the device must be dedicated to the guest or defined as a fullpack minidisk. On z/VM V5.4, support for the FlashCopy/SE feature requires APAR VM64449.</p> <p>PPRC supported for guest use only. Native use of PPRC requires ICKDSF V1.17.</p> <p>Guest support for High Performance FICON (zHPF) on z/VM V6.2 requires APARs VM65041 and VM65180. (zHPF is not supported on z/VM V5.4.) Also see “zHPF guest support hardware requirements” on page 23.</p> <p>Support for Soft Fence and Query Host Access on z/VM V6.3 requires APAR VM65322.</p> <p>For ICKDSF support information, see the ICKDSF Program Directory.</p>

Table 4. Device support – DASD (continued)

Device	z/VM support	Support notes
System Storage DS8000 series (2107) [FCP attached]	V5.4, V6.2, V6.3, V6.4	<p>For native z/VM use, the DS8000 volumes appear as 9336-20 DASD.</p> <p>For guest use, the DS8000 volumes can appear either as 9336-20 DASD or as SCSI LUNs directly accessed through FCP subchannels.</p> <p>DFSMS/VM FL221 support for SCSI LUNs larger than 381 GB requires APAR VM63664.</p>
System Storage DS6000 series (1750) [ESCON or FICON attached]	V5.4, V6.2, V6.3, V6.4	<p>z/VM automatically virtualizes the controller mode that a guest supports.</p> <p>Supported in 1750 native controller mode.</p> <p>Supported in emulation mode as a 3390 Model 3 or 6 Storage Control with the following DASD device types:</p> <ul style="list-style-type: none"> • 3390 Model 2, 3, and 9 DASD (including large Model 9s known as Model 27 and Model 54) • 3390 Model 2 and 3 DASD in 3380 track compatibility mode <p>Only 3390 Model 3 or 9 is supported for installation of z/VM. Mixed 3390 models are not supported for installation of z/VM.</p> <p>Performance-oriented track level commands supported for guest use only.</p> <p>Parallel Access Volumes (PAVs) supported as minidisks for guests that exploit PAV architecture.</p> <p>DFSMS/VM FL221 support for 3390 Model 9 disks larger than 10017 cylinders requires APAR VM63004.</p> <p>For guest use of FlashCopy, the device must be dedicated to the guest or defined as a fullpack minidisk.</p> <p>PPRC supported for guest use only. Native use of PPRC requires ICKDSF V1.17.</p> <p>For ICKDSF support information, see the ICKDSF Program Directory.</p>
System Storage DS6000 series (1750) [FCP attached]	V5.4, V6.2, V6.3, V6.4	<p>For native z/VM use, the DS6000 volumes appear as 9336-20 DASD.</p> <p>For guest use, the DS6000 volumes can appear either as 9336-20 DASD or as SCSI LUNs directly accessed through FCP subchannels.</p> <p>DFSMS/VM FL221 support for SCSI LUNs larger than 381 GB requires APAR VM63664.</p>

Supported Devices

Table 4. Device support – DASD (continued)

Device	z/VM support	Support notes
XIV Storage System	V5.4, V6.2, V6.3, V6.4	Direct attachment for system use on z/VM V5.4 requires APAR VM64708. HCD support on z/VM V5.4 requires APAR VM64672. ICKDSF support on z/VM V5.4 requires APAR PM11220.
FlashSystem	V6.4	Can be directly attached for system and EDEV use without requiring a SAN Volume Controller (SVC).

DASD control units and storage controls

Table 5. Device support – DASD control units and storage controls

Device	z/VM releases	Support notes
3990 Model 6	V5.4, V6.2, V6.3, V6.4	All functions of former 3990 models supported plus increased cache, System Clock Support, and the Control Unit Initiated Reconfiguration feature (CUIR). Concurrent Copy and Extended Remote Copy functions are supported for guest use only.
SAN Volume Controller 2145	V5.4, V6.2, V6.3, V6.4	RESTRICTION - For releases prior to z/VM V6.4, devices must be offline to z/VM when applying a Concurrent Code Load.
Storwize V7000	V5.4, V6.2, V6.3, V6.4	Configured with SET EDEVICE attribute 2145. RESTRICTION - For releases prior to z/VM V6.4, devices must be offline to z/VM when applying a Concurrent Code Load.

Tape units and tape libraries

Table 6. Device support – tape units and tape libraries

Device	z/VM releases	Support notes
3480 Tape Subsystem	V5.4, V6.2, V6.3, V6.4	Not supported for installation of z/VM.
3490 Tape Subsystem	V5.4, V6.2, V6.3, V6.4	Not supported for installation of z/VM.
3490E (Enhanced Capability Model) Tape Subsystem	V5.4, V6.2, V6.3, V6.4	Not supported for installation of z/VM.
TotalStorage Enterprise Tape System (3590)	V5.4, V6.2, V6.3, V6.4	DFSMS/VM FL221 support for the 3590 Model H requires APAR VM63097. Not supported for installation of z/VM V6.4 or later.

Table 6. Device support – tape units and tape libraries (continued)

Device	z/VM releases	Support notes
System Storage 3592 TS1120 (E05), TS1130 (E06), and TS1140 (E07) Tape Drives	V5.4, V6.2, V6.3, V6.4	<p>z/VM native tape functions (such as SPXTAPE, DUMP, CMS TAPE) supported with FICON Controller Attachment only.</p> <p>z/VM supports drive-based data encryption, including rekey support.</p> <p>Not supported for installation of z/VM V6.4 or later.</p> <p>On z/VM V5.4:</p> <ul style="list-style-type: none"> Support for the 3592 Model E06 requires APAR VM64459. Support for the 3592 Model E07 requires APAR VM64979. <p>DFSMS/VM FL221 requires:</p> <ul style="list-style-type: none"> APAR VM63353 for basic device support. APAR VM63460 to support Write Once Read Many (WORM) media. APAR VM64458 to support the 3592 Model E06. APAR VM65005 to support the 3592 Model E07.
System Storage TS3400 with TS1120 (E05) or TS1130 (E06) Tape Drives	V5.4, V6.2, V6.3, V6.4	<p>z/VM native tape functions (such as SPXTAPE, DUMP, CMS TAPE) supported with FICON Controller Attachment only.</p> <p>z/VM supports drive-based data encryption, including rekey support.</p> <p>Not supported for installation of z/VM V6.4 or later.</p> <p>On z/VM V5.4, support for the 3592 Model E06 requires APAR VM64459.</p> <p>DFSMS/VM FL221 requires:</p> <ul style="list-style-type: none"> APAR VM63353 for basic device support. APAR VM63460 to support Write Once Read Many (WORM) media. APAR VM64458 to support the 3592 Model E06.

Supported Devices

Table 6. Device support – tape units and tape libraries (continued)

Device	z/VM releases	Support notes
IBM Virtualization Engine TS7700 (Models TS7720 and TS7740)	V5.4, V6.2, V6.3, V6.4	<p>Native z/VM tape library support is provided by DFSMS/VM FL221 with PTFs for RMS APARs VM64773 and VM65005 (and prerequisite service).</p> <p>DFSMS/VM APAR VM65789 is required for RMS Copy Export support.</p> <p>On z/VM V5.4 and V6.2, CP base support for host and guest tape drive function requires the PTFs for CP APAR VM64979 (and prerequisite service).</p> <p>Not supported for installation of z/VM.</p>

Tape control units

Table 7. Device support – tape control units

Device	z/VM releases	Support notes
TotalStorage Enterprise Tape Controller 3590 Model A60	V5.4, V6.2, V6.3, V6.4	
TotalStorage Enterprise Tape Controller 3592 Model J70	V5.4, V6.2, V6.3, V6.4	
System Storage TS1120 Tape Controller Model C06	V5.4, V6.2, V6.3, V6.4	
System Storage Tape Controller Model C07	V5.4, V6.2, V6.3, V6.4	

Printers

Table 8. Device support – printers

Device	z/VM releases	Support notes
3203 Model 5	V5.4, V6.2, V6.3, V6.4	
3262	V5.4, V6.2, V6.3, V6.4	
3268 Model 2, 2C	V5.4, V6.2, V6.3, V6.4	Must be defined as a 3287.
3287 Model 1, 1C, 2, 2C, 4	V5.4, V6.2, V6.3, V6.4	
3289 Model 1, 3, 4, 8	V5.4, V6.2, V6.3, V6.4	
3800 Model 1	V5.4, V6.2, V6.3, V6.4	
3800 Model 3, 6, 8	V5.4, V6.2, V6.3, V6.4	Full support in Model 1 compatibility mode or using Advanced Function Printing (AFP) programs.
3812	V5.4, V6.2, V6.3, V6.4	Full support through the RSCS feature, using AFP programs; downloading of fonts is not supported. Loading of fonts requires Print Services Facility/VM (PSF/VM) licensed program (program number 5684-141).

Table 8. Device support – printers (continued)

Device	z/VM releases	Support notes
3816 Model 01D, 01S	V5.4, V6.2, V6.3, V6.4	Full support through the RSCS feature, using AFP programs; downloading of fonts is not supported. Loading of fonts requires PSF/VM licensed program.
3820	V5.4, V6.2, V6.3, V6.4	Full support through ACF/VTAM V4.2, using AFP programs.
3825	V5.4, V6.2, V6.3, V6.4	Full support using AFP programs.
3827	V5.4, V6.2, V6.3, V6.4	Full support using AFP programs.
3835	V5.4, V6.2, V6.3, V6.4	Full support using AFP programs.
3900	V5.4, V6.2, V6.3, V6.4	Must be defined as a 3800.
4245 Model 1	V5.4, V6.2, V6.3, V6.4	Supported as a 4245 or in 3262 compatibility mode.
4245 Model 12, 20	V5.4, V6.2, V6.3, V6.4	
4248 Model 1, 2	V5.4, V6.2, V6.3, V6.4	Supported as a 4248 or in 3211 compatibility mode.
6262 Model 14, 22	V5.4, V6.2, V6.3, V6.4	Must be defined as a 4248 Model 1.

Card readers and card punches

Table 9. Device support – card readers and card punches

Device	z/VM releases	Support notes
3505 Model B1, B2	V5.4, V6.2, V6.3, V6.4	
3525 Model P1, P2, P3	V5.4, V6.2, V6.3, V6.4	

Terminals, displays, and consoles

Most terminals are supported as a virtual machine console (in 3215 emulation mode or 3270 mode). 3270-family displays can be defined generically (for example, as 3270s) with dynamic computation of screen size based on information returned from the device.

z/VM also supports the integrated 3270 console and integrated ASCII console functions of the Hardware Management Console.

Display printers

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

Display control units

Table 10. Device support – display control units

Device	z/VM releases	Support notes
3174	V5.4, V6.2, V6.3, V6.4	
3272 Model 2	V5.4, V6.2, V6.3, V6.4	
3274	V5.4, V6.2, V6.3, V6.4	

Supported Devices

Table 10. Device support – display control units (continued)

Device	z/VM releases	Support notes
3276 Display/Control Unit	V5.4, V6.2, V6.3, V6.4	Must be dedicated to a virtual machine as a supported device.

Communications controllers

Table 11. Device support – communications controllers

Device	z/VM releases	Support notes
3745	V5.4, V6.2, V6.3, V6.4	Must be dedicated to a virtual machine as a supported device. Must be defined as a 3705, unless you are using dynamic system configuration.

Switches

Table 12. Device support – switches

Device	z/VM releases	Support notes
2032 FICON Director Model 1	V5.4, V6.2, V6.3, V6.4	
9032 ESCON Director Model 2	V5.4, V6.2, V6.3, V6.4	
9032 ESCON Director Model 3	V5.4, V6.2, V6.3, V6.4	
9032 ESCON Director Model 5	V5.4, V6.2, V6.3, V6.4	
9033 ESCON Director Model 1	V5.4, V6.2, V6.3, V6.4	
9033 ESCON Director Model 4	V5.4, V6.2, V6.3, V6.4	

Notes:

1. 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.
2. For more information about ESCON architecture, see *Introducing Enterprise Systems Connection*, GA23-0383.

Server features and miscellaneous devices

Table 13. Device support – server features and miscellaneous devices

Device	z/VM releases	Support notes
Channel-to-Channel Adapter (CTCA)	V5.4, V6.2, V6.3, V6.4	Must be dedicated to a virtual machine as a supported device.
Crypto Express2	V5.4, V6.2, V6.3	Supported on z10 and z9. Must be configured for shared or dedicated access by a virtual machine.

Table 13. Device support – server features and miscellaneous devices (continued)

Device	z/VM releases	Support notes
Crypto Express3	V5.4, V6.2, V6.3, V6.4	Supported on zEnterprise servers and z10. Must be configured for shared or dedicated access by a virtual machine. z/VM V5.4 requires APAR VM64656.
Crypto Express4S	V5.4, V6.2, V6.3, V6.4	Supported on zBC12 and zEC12. Must be configured for shared or dedicated access by a virtual machine in z/Architecture mode. z/VM V5.4 and V6.2 require APAR VM65007.
Crypto Express5S	V6.2, V6.3, V6.4	Supported on z13 or later. Must be configured for shared or dedicated access by a virtual machine in z/Architecture mode. Guest support and enhanced domain support requires APAR VM65577. DirMaint support for enhanced domain support requires APAR VM65588.
ESCON CTCA	V5.4, V6.2, V6.3, V6.4	Extended mode operation only.
FICON CTCA	V5.4, V6.2, V6.3	Must be dedicated to a guest as a supported device.
FICON Express2	V5.4, V6.2, V6.3	Guest CTCA support requires a dedicated device. Guest support for High Performance FICON on z/VM V6.2 requires APAR VM65041 (not supported on z/VM V5.4). Also see “zHPF guest support hardware requirements” on page 23.
FICON Express4	V5.4, V6.2, V6.3, V6.4	Guest CTCA support requires a dedicated device. Guest support for High Performance FICON on z/VM V6.2 requires APAR VM65041 (not supported on z/VM V5.4). Also see “zHPF guest support hardware requirements” on page 23.
FICON Express8, Express8S	V5.4, V6.2, V6.3, V6.4	Guest CTCA support requires a dedicated device. Guest support for High Performance FICON on z/VM V6.2 requires APAR VM65041 (not supported on z/VM V5.4). Also see “zHPF guest support hardware requirements” on page 23.

Supported Devices

Table 13. Device support – server features and miscellaneous devices (continued)

Device	z/VM releases	Support notes
FICON Express16S	V6.2, V6.3, V6.4	<p>Supported on z13 or later.</p> <p>Must be dedicated to a guest as a supported device.</p> <p>Guest support for High Performance FICON on z/VM V6.2 requires APAR VM65041. Also see “zHPF guest support hardware requirements” on page 23.</p>
HiperSockets	V5.4, V6.2, V6.3, V6.4	<p>Must be dedicated to a virtual machine as a supported device.</p> <p>Support for the z/VM virtual switch HiperSockets bridge on z/VM V6.2 requires APARs VM65042 and PM46988 (not supported on z/VM V5.4).</p>
OSA-2	V5.4	<p>Supported on zSeries.</p> <p>Must be dedicated to a virtual machine as a supported device.</p> <p>OSA-2 is recognized as TYPE=OSA.</p> <p>OSA/SF is required to customize the OSA modes. OSA/SF can be running either in an LPAR in the same system complex or as a z/VM guest. OSA/SF for z/VM can be run as a CMS application.</p>
OSA-Express	V5.4	<p>Supported on z9.</p> <p>Must be dedicated to a virtual machine as a supported device.</p> <p>OSA-Express is recognized as:</p> <ul style="list-style-type: none"> • TYPE=OSE for the Fast Ethernet and 155 ATM modes • TYPE=OSD for the Gigabit Ethernet mode • TYPE=OSC for the 1000BASE-T Ethernet feature with the Integrated Console Controller (ICC) function <p>OSA/SF is required to customize the OSA-Express modes. OSA/SF can be running either in an LPAR in the same system complex or as a z/VM guest. OSA/SF for z/VM can be run as a CMS application.</p>

Table 13. Device support – server features and miscellaneous devices (continued)

Device	z/VM releases	Support notes
OSA-Express2	V5.4, V6.2, V6.3, V6.4	<p>Supported on z114, z196, z10, and z9.</p> <p>Must be dedicated to a virtual machine as a supported device.</p> <p>OSA-Express2 is recognized as:</p> <ul style="list-style-type: none"> • TYPE=OSD for the Gigabit Ethernet and 10 Gigabit Ethernet modes. z/VM supports IEEE 802.3ad link aggregation. • TYPE=OSC for the 1000BASE-T Ethernet feature with the Integrated Console Controller (ICC) function. • TYPE=OSE for the 1000Base-T Ethernet feature for use in LAN Emulation mode using the LAN Channel Station (LCS) protocol (non-QDIO mode). • TYPE=OSN for the Gigabit Ethernet and 1000BASE-T Ethernet features with the Open Systems Adapter for the Network Control Program (OSA NCP) function. <p>z/VM supports isolation of QDIO data connection communications on a shared OSA-Express2 port in support of multitier security zones. z/VM V5.4 requires APARs VM64463 and PK67610.</p> <p>OSA/SF is required to customize the OSA-Express2 modes. OSA/SF can be running either in an LPAR in the same system complex or as a z/VM guest. OSA/SF for z/VM can be run as a CMS application.</p>

Supported Devices

Table 13. Device support – server features and miscellaneous devices (continued)

Device	z/VM releases	Support notes
OSA-Express3	V5.4, V6.2, V6.3, V6.4	<p>Supported on zEnterprise servers and z10.</p> <p>Must be dedicated to a virtual machine as a supported device.</p> <p>OSA-Express3 is recognized as:</p> <ul style="list-style-type: none"> • TYPE=OSD for the Gigabit Ethernet, 10 Gigabit Ethernet, and 1000BASE-T modes. z/VM supports IEEE 802.3ad link aggregation. • TYPE=OSC for the 1000BASE-T Ethernet feature with the Integrated Console Controller (ICC) function. • TYPE=OSE for the 1000Base-T Ethernet feature for use in LAN Emulation mode using the LAN Channel Station (LCS) protocol (non-QDIO mode). • TYPE=OSN for the Gigabit Ethernet and 1000BASE-T Ethernet features with the Open Systems Adapter for the Network Control Program (OSA NCP) function. • TYPE=OSM for the 1000BASE-T Ethernet intranode management network (INMN) in an IBM zEnterprise System. z/VM V5.4 (with APAR VM64794) and z/VM V6.2 support configuring via dynamic I/O, but the device cannot be varied online. Not supported by z/VM V6.3 or later. • TYPE=OSX for the 10 Gigabit Ethernet intraensemble data network (IEDN) in an IBM zEnterprise System. z/VM V5.4 (with APAR VM64794) and z/VM V6.2 support configuring via dynamic I/O, but the device cannot be varied online. Not supported by z/VM V6.3 or later. <p>z/VM supports four ports per feature on the Gigabit Ethernet and 1000BASE-T features on the zEC12, z196, and z10 EC, and two ports per feature on the zBC12, z114, and z10 BC.</p> <p>z/VM supports isolation of QDIO data connection communications on a shared OSA-Express3 port on a zEnterprise or z10 server in support of multitier security zones. z/VM V5.4 requires APARs VM64463 and PK67610.</p> <p>OSA/SF is required to customize the OSA-Express3 modes. OSA/SF can be running either in an LPAR in the same system complex or as a z/VM guest. OSA/SF for z/VM can be run as a CMS application.</p>

Table 13. Device support – server features and miscellaneous devices (continued)

Device	z/VM releases	Support notes
OSA-Express4S	V5.4, V6.2, V6.3, V6.4	<p>Supported on z13s, z13, and zEnterprise servers.</p> <p>Must be dedicated to a virtual machine as a supported device.</p> <p>OSA-Express4S is recognized as:</p> <ul style="list-style-type: none"> • TYPE=OSD for the Gigabit Ethernet, 10 Gigabit Ethernet, and 1000BASE-T Ethernet features configured to support QDIO. z/VM supports IEEE 802.3ad link aggregation. • TYPE=OSC for the 1000BASE-T Ethernet feature with the Integrated Console Controller (ICC) function. • TYPE=OSN for the 1000BASE-T Ethernet feature with the Open Systems Adapter for the Network Control Program (OSA NCP) function. • TYPE=OSE for the 1000Base-T Ethernet feature for use in LAN Emulation mode using the LAN Channel Station (LCS) protocol (non-QDIO mode). • TYPE=OSM for the 1000BASE-T Ethernet intranode management network (INMN) in an IBM zEnterprise System. z/VM V5.4 (with APAR VM64794) and z/VM V6.2 support configuring via dynamic I/O, but the device cannot be varied online. Not supported by z/VM V6.3 or later. • TYPE=OSX for the 10 Gigabit Ethernet intraensemble data network (IEDN) in an IBM zEnterprise System. z/VM V5.4 (with APAR VM64794) and z/VM V6.2 support configuring via dynamic I/O, but the device cannot be varied online. Not supported by z/VM V6.3 or later. <p>z/VM supports two ports per feature on the Gigabit Ethernet and 1000BASE-T Ethernet features, and one port per feature on the 10 Gigabit Ethernet features.</p> <p>z/VM supports isolation of QDIO data connection communications on a shared OSA-Express4S port in support of multitier security zones. z/VM V5.4 requires APARs VM64463 and PK67610.</p> <p>z/VM V6.3 support for Multi-VSwitch Link Aggregation on z13 requires APARs VM65583 (CP) and PI21053 (TCP/IP). Systems management API support requires APAR VM65670. Performance Toolkit support requires APAR VM65528.</p> <p>(continued)</p>

Supported Devices

Table 13. Device support – server features and miscellaneous devices (continued)

Device	z/VM releases	Support notes
OSA-Express4S (continued)		<p>One of the following is required to customize the OSA-Express4S modes:</p> <ul style="list-style-type: none">• OSA/SF on the HMC• OSA/SF facility. APAR OA37060 is required on the z196 and z114. APARs OA37060 and OA38418 are required on the zEC12 or later. OSA/SF facility 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.

Table 13. Device support – server features and miscellaneous devices (continued)

Device	z/VM releases	Support notes
OSA-Express5S	V5.4, V6.2, V6.3, V6.4	<p>Supported on z13s, z13, zBC12, and zEC12.</p> <p>Must be dedicated to a virtual machine as a supported device.</p> <p>OSA-Express5S is recognized as:</p> <ul style="list-style-type: none"> • TYPE=OSD for the Gigabit Ethernet, 10 Gigabit Ethernet, and 1000BASE-T features configured to support QDIO. z/VM supports IEEE 802.3ad link aggregation. • TYPE=OSC for the 1000BASE-T Ethernet feature with the Integrated Console Controller (ICC) function. • TYPE=OSE for the 1000Base-T Ethernet feature for use in LAN Emulation mode using the LAN Channel Station (LCS) protocol (non-QDIO mode). • TYPE=OSN for the 1000BASE-T Ethernet feature with the Open Systems Adapter for the Network Control Program (OSA NCP) function. • TYPE=OSM for the 1000BASE-T Ethernet intranode management network (INMN) in an IBM zEnterprise System. z/VM V5.4 (with APAR VM64794) and z/VM V6.2 support configuring via dynamic I/O, but the device cannot be varied online. Not supported by z/VM V6.3 or later. • TYPE=OSX for the 10 Gigabit Ethernet intraensemble data network (IEDN) in an IBM zEnterprise System. z/VM V5.4 (with APAR VM64794) and z/VM V6.2 support configuring via dynamic I/O, but the device cannot be varied online. Not supported by z/VM V6.3 or later. <p>z/VM supports two ports per feature on the Gigabit Ethernet and 1000BASE-T Ethernet features, and one port per feature on the 10 Gigabit Ethernet features.</p> <p>z/VM supports isolation of QDIO data connection communications on a shared OSA-Express5S port in support of multitier security zones. z/VM V5.4 requires APARs VM64463 and PK67610.</p> <p>(continued)</p>
OSA-Express5S (continued)		<p>z/VM V6.3 support for Multi-VSwitch Link Aggregation on z13 requires APARs VM65583 (CP) and PI21053 (TCP/IP). Systems management API support requires APAR VM65670. Performance Toolkit support requires APAR VM65528.</p> <p>OSA/SF on the HMC is required to customize the OSA-Express5S modes.</p>

Supported Devices

Table 13. Device support – server features and miscellaneous devices (continued)

Device	z/VM releases	Support notes
PCI Cryptographic Accelerator	V5.4	Supported on zSeries. Must be configured for shared or dedicated access by a virtual machine.
PCI Cryptographic Coprocessor	V5.4	Supported on zSeries. Must be configured for shared or dedicated access by a virtual machine.
PCIX Cryptographic Coprocessor	V5.4	Supported on zSeries. Must be configured for shared or dedicated access by a virtual machine.
Regional Crypto Enablement (RCE) adapter	V6.2, V6.3, V6.4	Supported on z13s and z13 (Driver D27), or later. z/VM V6.2 and V6.3 require APARs VM65577 and VM65716.
zEDC Express	V6.3, V6.4	Supported on zEC12 and zBC12 or later. zEC12 and zBC12 require Driver D15 Bundle 21. See the PSP bucket for specific details. Must be dedicated to a virtual machine as a supported device. z/VM V6.3 requires APARs VM65417 (CP), VM65572 (CP TRACERED support), and PI20509 (TCP/IP).
10GbE RoCE Express	V6.3, V6.4	Supported on zEC12 and zBC12 or later. zEC12 and zBC12 require Driver D15 Bundle 21. See the PSP bucket for specific details. Must be dedicated to a virtual machine as a supported device. z/VM V6.3 requires APARs VM65417 (CP), VM65572 (CP TRACERED support), and PI20509 (TCP/IP).
3088	V5.4, V6.2, V6.3, V6.4	Must be dedicated to a virtual machine as a supported device.
3088 Model 61 PSCA Card (common link access to workstations)	V5.4, V6.2, V6.3, V6.4	
3174 Model 12L, 22L	V5.4, V6.2, V6.3, V6.4	
3737 Remote Channel-to-Channel Unit Model 2	V5.4, V6.2, V6.3, V6.4	Must be dedicated to a virtual machine as a supported device. Must be defined as a CTCA.
4753 Network Security Processor	V5.4, V6.2, V6.3, V6.4	Must be dedicated to a guest as a supported device.

Table 13. Device support – server features and miscellaneous devices (continued)

Device	z/VM releases	Support notes
7171 Device Attachment Control Unit	V5.4, V6.2, V6.3, V6.4	The 7171 ASCII Device Attachment Control Unit (DACU) allows the attachment of ASCII terminals locally or (through Start-Stop communication line) remotely. The ASCII terminals appear to the host as 3277s, 3278s, and 3279s, with the DACU itself appearing as a 3274 control unit.
9034 ESCON Converter Model 1	V5.4, V6.2, V6.3, V6.4	
9035 ESCON Converter Model 2	V5.4, V6.2, V6.3, V6.4	

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

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

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those

websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information may contain sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The

sample programs are provided “AS IS”, without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at IBM copyright and trademark information - United States (www.ibm.com/legal/us/en/copytrade.shtml).

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

Windows is a trademark of Microsoft Corporation in the United States, other countries, or both.

Java is a trademark or registered trademark of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED “AS-IS” AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM online privacy statement

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM Online Privacy Statement Highlights at <http://www.ibm.com/privacy> and the IBM Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled “Cookies, Web Beacons and Other Technologies”, and the IBM Software Products and Software-as-a-Service Privacy Statement at <http://www.ibm.com/software/info/product-privacy>.

Glossary

For a list of z/VM terms and their definitions, see *z/VM: Glossary*.

The z/VM glossary is also available through the online z/VM HELP Facility, if HELP files are installed on your z/VM system. For example, to display the definition of the term “dedicated device”, issue the following HELP command:

```
help glossary dedicated device
```

While you are in the glossary help file, you can do additional searches:

- To display the definition of a new term, type a new HELP command on the command line:

```
help glossary newterm
```

This command opens a new help file inside the previous help file. You can repeat this process many times. The status area in the lower right corner of the screen shows how many help files you have open. To close the current file, press the Quit key (PF3/F3). To exit from the HELP Facility, press the Return key (PF4/F4).

- To search for a word, phrase, or character string, type it on the command line and press the Clocate key (PF5/F5). To find other occurrences, press the key multiple times.

The Clocate function searches from the current location to the end of the file. It does not wrap. To search the whole file, press the Top key (PF2/F2) to go to the top of the file before using Clocate.

Index

Numerics

370 Accommodation Facility 40
370 virtual machines not supported 40
370-XA architecture compatibility 40

A

abstracts of z/VM books
 base library - application programming 76
 base library - customization and tuning 74
 base library - installation, migration, and service 70
 base library - operation and use 75
 base library - overview 70
 base library - planning and administration 71
 base library - system diagnosis 81
 DFSMS/VM feature 82
 DirMaint feature 83
 IBM Wave feature 83
 OpenExtensions 78
 Performance Toolkit feature 84
 RACF feature 84
 RSCS feature 85
accessibility 5
ACF/VTAM 59
adapter interruptions 48
analyzing dump data 41
APPC/VM
 AVS 58
 ISFC 41
 TSAF 61
APPC/VM VTAM Support (AVS) 58
application development and deployment, support for
 CMS support 52
 overview 4
application programming interfaces 4
architecture compatibility 40
architecture level set (ALS) 19
architecture support
 370-XA compatibility 40
 dynamic storage reconfiguration 42
 ESA/370 compatibility 40
 ESA/390 39
 ESA/XC 40
 ESCON 43
 FCP 43
 FICON 43
 HiperSockets 44
 IEEE Floating Point 45
 LCSS 45
 POSIX 53
 QDIO 48
 TCP/IP 56
 z/Architecture 40
auditability 35
auditable events 38
AVS 58

B

base product
 AVS 58
 CMS 51
 CP 39
 Dump Viewing Facility 58
 GCS 59
 HCD and HCM for z/VM 59
 Language Environment 60
 OSA/SF 60
 REXX/VM 60
 TCP/IP for z/VM 56
 TSAF 61
 VMSES/E 61
BFS (byte file system) 54
books, z/VM
 abstracts
 base library - application programming 76
 base library - customization and tuning 74
 base library - installation, migration, and service 70
 base library - operation and use 75
 base library - overview 70
 base library - planning and administration 71
 base library - system diagnosis 81
 DFSMS/VM feature 82
 DirMaint feature 83
 IBM Wave feature 83
 OpenExtensions 78
 Performance Toolkit feature 84
 RACF feature 84
 RSCS feature 85
 basic 69
 media 69
 sources 69
 supplied with z/VM 69
byte file system (BFS) 54

C

cables, fiber optic 43
card readers and punches supported by z/VM 105
changes in z/VM V6.4
 3590 and 3592 tape not supported for z/VM
 installation 14
 architecture level set (ALS) 14
 CMS Pipelines 10
 CP environment variables 9
 determine installed service 14
 DirMaint RACF connector 10
 DNS IPv6 support 14
 Dynamic Partition Manager (DPM) administrative mode
 for Linux 13
 dynamic SMT level 9
 encryption of TCPNJE connections 14
 ESA/390 removal 14
 FlashSystem support for FCP-attached SCSI disks 14
 guest large page support 7
 guest Transactional Execution support 8
 hardware currency and exploitation 12
 HCD Dynamic I/O Activates, large 8

- changes in z/VM V6.4 (*continued*)
 - HyperPAV technology exploitation 7
 - hypervisor efficiency and scalability 7
 - IBM Wave optional feature 10
 - installation, migration, and serviceability 13
 - LOGON security improvement 14
 - memory scalability efficiency improvements 8
 - nondefault tape mode support 14
 - Performance Toolkit enhancements 10
 - QUERY SHUTDOWN command 9
 - RACF automatic control of access list authority 10
 - real memory and guest virtual memory support 8
 - reset counters for virtual switch 14
 - SCSI management queries 9
 - SCSI RAS enhancements 9
 - SET DIALDROP command 14
 - Shared Memory Communications - Direct Memory Access (SMC-D) support 12
 - SMAPI support 14
 - statements of general direction 17
 - surplus CPU power distribution improvement 8
 - system ease of use improvements 8
 - UID (user defined identifier) support for real PCI functions 12
 - upgrade in place, enhanced 13
 - VLAN access security improvement 14
 - XSTORE support removed 14
 - z/OS V2.2 equivalency 11
 - zManager support removed 14
- channel subsystem
 - channel program access protection 35
- Cloud Manager Appliance
 - program requirements 30
- CMS (Conversational Monitor System)
 - application programming 52
 - byte file system (BFS) 54
 - ESA/390 CMS 51
 - HELP facility 55
 - multicultural support 32
 - multitasking, application 53
 - OpenExtensions 53
 - overview 51
 - Pipelines 53
 - Program Management Binder for CMS 54
 - Reusable Server Kernel 54
 - Shared File System (SFS) 52
 - systems management API (SMAPI) 53
 - versions 51
 - Virtual Machine Resource Manager (VMRM) 55
 - XEDIT 55
 - z/Architecture CMS (z/CMS) 51
- Collaborative Memory Management Assist (CMMA) 42
- communications controllers supported by z/VM 106
- compatibility
 - 370-XA architecture 40
 - ESA/370 architecture 40
 - System/370 architecture 40
- components of z/VM
 - AVS 58
 - CMS 51
 - CP 39
 - Dump Viewing Facility 58
 - GCS 59
 - Language Environment 60
 - REXX/VM 60
 - TSAF 61
 - VMSES/E 61
- concurrent-sense facility support 42
- consoles supported by z/VM 105
- Control Program (CP)
 - architecture compatibility 40
 - ESA/XC architecture support 40
 - guest LAN 40
 - hardware and architecture support
 - adapter interruptions 48
 - Collaborative Memory Management Assist (CMMA) 42
 - concurrent-sense facility 42
 - CPACF 42
 - CPU-Measurement Counter Facility 42
 - cryptographic facilities 42
 - dynamic storage reconfiguration 42
 - ECMDS 43
 - enhanced-DAT facility 43
 - environment information interface 43
 - ESCON 43
 - Extended-TOD-Clock Facility 43
 - FCP 43
 - FICON channels 43
 - FICON CTCA 43
 - FlashCopy 44
 - guest ASCII console 44
 - HiperSockets 44
 - HyperPAV 45
 - ICF 50
 - IEEE Floating Point 45
 - IFL 50
 - integrated 3270 console 45
 - LCSS 45
 - MIDAWs 45
 - Move-Page Facility 46
 - NPIV 46
 - OSA 46
 - Parallel Access Volumes (PAV) 47
 - PPRC 47
 - QDIO architecture 48
 - simultaneous multithreading (SMT) 49
 - specialty processors 50
 - STP 48
 - Transactional Execution facility 51
 - zAAP 50
 - zHPF 44
 - zIIP 50
- ISFC 41
- multicultural support 32
- overview 39
- virtual machine
 - architectures 39
 - type 40
 - VM Dump Tool 41
 - z/Architecture support 40
 - z/VM HiperDispatch 41
- control units supported by z/VM
 - DASD 102
 - display 105
 - storage 102
 - tape 104
- controllers, communications, supported by z/VM 106
- Conversational Monitor System (CMS)
 - application programming 52
 - byte file system (BFS) 54
 - ESA/390 CMS 51
 - HELP facility 55
 - multicultural support 32

- Conversational Monitor System (CMS) *(continued)*
 - multitasking, application 53
 - OpenExtensions 53
 - overview 51
 - Pipelines 53
 - Program Management Binder for CMS 54
 - Reusable Server Kernel 54
 - Shared File System (SFS) 52
 - systems management API (SMAPI) 53
 - versions 51
 - Virtual Machine Resource Manager (VMRM) 55
 - XEDIT 55
 - z/Architecture CMS (z/CMS) 51
- CP (Control Program)
 - architecture compatibility 40
 - ESA/XC architecture support 40
 - guest LAN 40
 - hardware and architecture support
 - adapter interruptions 48
 - Collaborative Memory Management Assist (CMMA) 42
 - concurrent-sense facility 42
 - CPACF 42
 - CPU-Measurement Counter Facility 42
 - cryptographic facilities 42
 - dynamic storage reconfiguration 42
 - ECMDS 43
 - enhanced-DAT facility 43
 - environment information interface 43
 - ESCON 43
 - Extended-TOD-Clock Facility 43
 - FCP 43
 - FICON channels 43
 - FICON CTCA 43
 - FlashCopy 44
 - guest ASCII console 44
 - HiperSockets 44
 - HyperPAV 45
 - ICF 50
 - IEEE Floating Point 45
 - IFL 50
 - integrated 3270 console 45
 - LCSS 45
 - MIDAWs 45
 - Move-Page Facility 46
 - NPIV 46
 - OSA 46
 - Parallel Access Volumes (PAV) 47
 - PCIe 47
 - PPRC 47
 - QDIO architecture 48
 - simultaneous multithreading (SMT) 49
 - specialty processors 50
 - STP 48
 - Transactional Execution facility 51
 - Vector Facility for z/Architecture (SIMD) 51
 - virtual machine I/O assist 51
 - zAAP 50
 - zHPF 44
 - zIIP 50
- ISFC 41
- multicultural support 32
- overview 39
- virtual machine
 - architectures 39
 - type 40
- VM Dump Tool 41

- CP (Control Program) *(continued)*
 - z/Architecture support 40
 - z/VM HiperDispatch 41
- CPU-Measurement Counter Facility 42
- cryptographic facilities, support for 42
- customer responsibilities for data security 35

D

- DASD (direct access storage device)
 - control units supported by z/VM 102
 - HyperPAV support 47
 - Parallel Access Volumes (PAV) support 47
 - required to install z/VM 20
 - supported by z/VM 98
- data integrity
 - customer responsibilities 35
 - for guests 33
- dedicated-only device, definition of 24
- device support, types of 23
- devices supported by z/VM
 - card readers and punches 105
 - communications controllers 106
 - consoles 105
 - DASD 98
 - DASD control units 102
 - display control units 105
 - display printers 105
 - displays 105
 - miscellaneous 106
 - printers 104
 - server features 106
 - storage controls 102
 - switches 106
 - tape control units 104
 - tape libraries 102
 - tape units 102
 - terminals 105
- DFSMS/VM
 - books, abstract of 82
 - description 63
 - multicultural support 32
 - program requirements 27
- diagnosing system problems 58
- direct access storage device (DASD)
 - control units supported by z/VM 102
 - HyperPAV support 47
 - Parallel Access Volumes (PAV) support 47
 - required to install z/VM 20
 - supported by z/VM 98
- DirMaint (Directory Maintenance Facility)
 - books, abstract of 83
 - description 64
 - multicultural support 32
 - program requirements 28
 - security 37
 - system integrity 35
- display control units supported by z/VM 105
- display printers supported by z/VM 105
- displays supported by z/VM 105
- DOS/VSE support 52
- dump data, analyzing 41
- Dump Viewing Facility 58
- DVD
 - installing z/VM from 20
 - z/VM System Image 32
- dynamic storage reconfiguration 42

E

- ECMDS (extended channel measurement data support) 43
- editor, z/VM 55
- enhanced-DAT facility 43
- Enterprise Systems Connection Architecture (ESCON)
 - description 43
 - devices supported 106
 - ESCON Multiple Image Facility (EMIF) 87
- environment information interface 43
- ESA virtual machine 39
- ESA/370 architecture compatibility 40
- ESA/390 CMS 51
- ESA/XC architecture 40
- ESCON (Enterprise Systems Connection Architecture)
 - description 43
 - devices supported 106
 - ESCON Multiple Image Facility (EMIF) 87
- extended channel measurement data support (ECMDS) 43
- Extended Configuration (ESA/XC) architecture 40
- Extended-TOD-Clock Facility support 43

F

- FCP (Fibre Channel Protocol) support 43
- features, optional
 - DFSMS/VM 63
 - DirMaint 64
 - IBM Wave for z/VM 65
 - IBM z/VM Single System Image Feature 68
 - Performance Toolkit for VM 65
 - RACF Security Server for z/VM 66
 - RSCS Networking for z/VM 67
- features, server, supported by z/VM 106
- Federal Information Processing Standard (FIPS) support 37
- fiber optic cables 43
- FICON (fibre connection) channel support 43
- FICON CTCA support 43
- FICON Express16S 43
- FICON Express4 43
- FICON Express8 43
- file pool, CMS
 - description 52
 - management using DFSMS/VM 63
 - security 36
- FIPS support 37
- first-level storage
 - minimum and maximum 20
- FlashCopy support 44
- floating point support 45

G

- GCS (Group Control System) 59
- guest ASCII console 44
- guest LAN 40
- guest operating system support
 - architecture compatibility 40
 - matrix 93

H

- hardware architectures and facilities, support for
 - adapter interruptions 48
 - Collaborative Memory Management Assist (CMMA) 42
 - concurrent-sense facility 42

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

- CPACF 42
- CPU-Measurement Counter Facility 42
- cryptographic facilities 42
- ECMDS 43
- enhanced-DAT facility 43
- environment information interface 43
- ESCON 43
- Extended-TOD-Clock Facility 43
- FCP 43
- FICON channels 43
- FICON CTCA 43
- FlashCopy 44
- guest ASCII console 44
- HiperSockets 44
- HyperPAV 45
- ICF 50
- IEEE Floating Point 45
- IFL 50
- integrated 3270 console 45
- LCSS 45
- MIDAWs 45
- Move-Page Facility 46
- NPIV 46
- OSA 46
- Parallel Access Volumes (PAV) 47
- PCIe 47
- PPRC 47
- QDIO architecture 48
- simultaneous multithreading (SMT) 49
- specialty processors 50
- STP 48
- Transactional Execution facility 51
- Vector Facility for z/Architecture (SIMD) 51
- virtual machine I/O assist 51
- zAAP 50
- zHPF 44
- zIIP 50
- Hardware Configuration Definition and Hardware Configuration Manager for z/VM
 - description 59
 - program requirements 28
- Hardware Management Console 20
- hardware requirements
 - IBM Wave 21
 - Multi-VSwitch Link Aggregation support 23
 - OSA/SF 21
 - PCIe support 23
 - Performance Toolkit 21
 - server 19
 - SSI cluster 22
 - TCP/IP for z/VM 22
 - VM Guest RSA-Assist Support for Linux 22
 - workstation 21
 - z/VM installation 20
 - zHPF guest support 23
 - zHPF paging 23
- HCD and HCM for z/VM
 - description 59
 - program requirements 28
- HELP facility 55
- High Level Assembler 25, 28, 29
- High Performance FICON (zHPF)
 - description 44
 - guest support hardware requirements 23
 - paging hardware requirements 23
- HiperDispatch 41

HiperSockets bridge 41
HiperSockets support 44
HyperPAV
 description 45
HyperPAV support 47, 100

I

I/O (input/output)
 channel program access protection 35
 ESCON 43
 FICON channels 43
 FlashCopy 44
IBM Knowledge Center 69
IBM Publications Center 69
IBM Wave
 program requirements 28
IBM Wave for VM
 publications 83
IBM z Integrated Information Processor (zIIP), support for 50
IBM zEnterprise z Application Assist Processor (zAAP),
 support for 50
ICF (Internal Coupling Facility), support for 50
IEEE Floating Point support 45
IFL (Integrated Facility for Linux), support for 50
input/output (I/O)
 channel program access protection 35
 ESCON 43
 FICON channels 43
 FlashCopy 44
installation requirements 20
integrated 3270 console 45
Integrated Facility for Linux (IFL), support for 50
integrity
 data 33
 DirMaint 35
 system 34
Internal Coupling Facility (ICF), support for 50
IOCP program requirements 28
IPv6 support 56
ISFC (Inter-System Facility for Communications) 41

J

journaling 35

L

Language Environment
 description 60
 program requirements 28
language support 31
LCSS (logical channel subsystems) 45
library, z/VM
 abstracts of books
 base library - application programming 76
 base library - customization and tuning 74
 base library - installation, migration, and service 70
 base library - operation and use 75
 base library - overview 70
 base library - planning and administration 71
 base library - system diagnosis 81
 DFSMS/VM feature 82
 DirMaint feature 83
 IBM Wave feature 83
 OpenExtensions 78

library, z/VM (*continued*)
 abstracts of books (*continued*)
 Performance Toolkit feature 84
 RACF feature 84
 RSCS feature 85
 documentation supplied with z/VM 69
licensed programs supported on z/VM 31
limitations 30
Linux on z Systems
 guest support for 93
 program requirements 30
logical channel subsystems (LCSS) 45
LPAR (logical partition) 87
LPAR group absolute capacity capping 45

M

main storage
 minimum and maximum 20
memory
 minimum and maximum 20
MIDAW support 45
Move-Page Facility support 46
Multi-VSwitch Link Aggregation
 hardware requirements 23
multicultural support 31
multitasking, CMS application 53

N

N_Port ID virtualization (NPIV) 46
network interface card (NIC) 40
new or changed in z/VM V6.4
 3590 and 3592 tape not supported for z/VM
 installation 14
 architecture level set (ALS) 14
 CMS Pipelines 10
 CP environment variables 9
 determine installed service 14
 DirMaint RACF connector 10
 DNS IPv6 support 14
 Dynamic Partition Manager (DPM) administrative mode
 for Linux 13
 dynamic SMT level 9
 encryption of TCPNJE connections 14
 ESA/390 removal 14
 FlashSystem support for FCP-attached SCSI disks 14
 guest large page support 7
 guest Transactional Execution support 8
 hardware currency and exploitation 12
 HCD Dynamic I/O Activates, large 8
 HyperPAV technology exploitation 7
 hypervisor efficiency and scalability 7
 IBM Wave optional feature 10
 installation, migration, and serviceability 13
 LOGON security improvement 14
 memory scalability efficiency improvements 8
 nondefault tape mode support 14
 Performance Toolkit enhancements 10
 QUERY SHUTDOWN command 9
 RACF automatic control of access list authority 10
 real memory and guest virtual memory support 8
 reset counters for virtual switch 14
 SCSI management queries 9
 SCSI RAS enhancements 9
 SET DIALDROP command 14

new or changed in z/VM V6.4 (*continued*)

- Shared Memory Communications - Direct Memory Access (SMC-D) support 12
- SMAPI support 14
- statements of general direction 17
- surplus CPU power distribution improvement 8
- system ease of use improvements 8
- UID (user defined identifier) support for real PCI functions 12
- upgrade in place, enhanced 13
- VLAN access security improvement 14
- XSTORE support removed 14
- z/OS V2.2 equivalency 11
- zManager support removed 14

O

- object code only (OCO) 33
- online books
 - IBM Knowledge Center 69
 - IBM Publications Center 69
 - z/VM Collection 69
- Open Systems Adapter Support Facility (OSA/SF)
 - description 60
 - hardware requirements 21
 - program requirements 28
- OpenExtensions
 - books, abstracts of 78
 - byte file system (BFS) 54
 - overview 53
 - program requirements 26
- OpenStack
 - program requirements 30
- OpenStack support 16
- operating systems supported as guests of z/VM 93
- Operations Manager for z/VM 4
- optional features
 - DFSMS/VM 63
 - DirMaint 64
 - IBM Wave for z/VM 65
 - IBM z/VM Single System Image Feature 68
 - Performance Toolkit for VM 65
 - RACF Security Server for z/VM 66
 - RSCS Networking for z/VM 67
- ordering information 32
- OS/MVS simulation 52
- OSA (Open Systems Adapter) support 46
- OSA-Express 60
- OSA/SF (Open Systems Adapter Support Facility)
 - description 60
 - hardware requirements 21
 - program requirements 28

P

- packaging, product
 - distribution media 32
 - System Image contents 32
- Parallel Access Volumes (PAV) support 47
- PCIe
 - hardware requirements 23
- PCIe support 47
- Peer-to-Peer Remote Copy (PPRC) support 47
- Performance Toolkit for VM
 - book, abstract of 84
 - description 65

Performance Toolkit for VM (*continued*)

- hardware requirements 21
- program requirements 29
- Peripheral Component Interconnect Express (PCIe) support 47
- POSIX support 53
- PR/SM (Processor Resource/Systems Manager) 87
- prerequisite licensed programs 25
- Preventive Service Planning (PSP) buckets 87
- printers supported by z/VM 104
- problems, diagnosing 58
- product information
 - IBM Knowledge Center 69
- Program Management Binder for CMS 54
- program requirements
 - base components 26
 - Cloud Manager Appliance 30
 - DFSMS/VM 27
 - DirMaint 28
 - facilities and optional features, supported levels of 26
 - HCD and HCM for z/VM 28
 - High Level Assembler 25
 - IBM Wave 28
 - installing z/VM from DVD 20
 - IOCP 28
 - Language Environment 28
 - Linux 30
 - OpenExtensions 26
 - OpenStack 30
 - OSA/SF 28
 - other products 30
 - Performance Toolkit for VM 29
 - POSIX 26
 - prerequisite licensed programs 25
 - RACF Security Server for z/VM 29
 - RSCS Networking for z/VM 29
 - shared-DASD complex 29
 - SSI cluster 29
 - TCP/IP for z/VM 30
- Programmable Operator 4
- programming interfaces 4
- programming languages and environments 4
- programs supported on z/VM 31
- PSP buckets 87
- publications, z/VM
 - abstracts
 - base library - application programming 76
 - base library - customization and tuning 74
 - base library - installation, migration, and service 70
 - base library - operation and use 75
 - base library - overview 70
 - base library - planning and administration 71
 - base library - system diagnosis 81
 - DFSMS/VM feature 82
 - DirMaint feature 83
 - IBM Wave feature 83
 - OpenExtensions 78
 - Performance Toolkit feature 84
 - RACF feature 84
 - RSCS feature 85
- basic 69
- media 69
- sources 69
- supplied with z/VM 69

Q

QDIO (queued direct I/O) architecture support 48
queued direct I/O (QDIO) architecture support 48

R

RACF Security Server for z/VM (RACF Security Server for z/VM) 37
 books, abstract of 84
 description 66
 program requirements 29
real storage
 minimum and maximum 20
Remote Spooling Communications Subsystem (RSCS)
 Networking for z/VM (RSCS)
 books, abstract of 85
 description 67
 program requirements 29
requirements
 customer responsibilities for data security 35
 hardware
 installing z/VM 20
 Multi-VSwitch Link Aggregation support 23
 OSA/SF 21
 PCIe support 23
 Performance Toolkit 21
 server 19
 SSI cluster 22
 storage 20
 TCP/IP for z/VM 22
 workstation 21
 zHPF guest support 23
 zHPF paging 23
 installing z/VM 20
 program
 base components 26
 Cloud Manager Appliance 30
 DFSMS/VM 27
 DirMaint 28
 facilities and optional features, supported levels of 26
 HCD and HCM for z/VM 28
 High Level Assembler 25
 IBM Wave 28
 installing z/VM from DVD 20
 IOCP 28
 Language Environment 28
 Linux 30
 OpenExtensions 26
 OpenStack 30
 OSA/SF 28
 other products 30
 Performance Toolkit for VM 29
 POSIX 26
 prerequisite licensed programs 25
 RACF Security Server for z/VM 29
 RSCS Networking for z/VM 29
 shared-DASD complex 29
 SSI clusters 29
 TCP/IP for z/VM 30
 server 19
 storage requirements 20
 virtual machine size, minimum 21
restrictions
 virtual machine 31
Reusable Server Kernel 54
REXX language 60

REXX/VM 60
RSA-Assist support 42
RSCS (Remote Spooling Communications Subsystem (RSCS)
 Networking for z/VM)
 books, abstract of 85
 description 67
 program requirements 29

S

SCSI disks, installing z/VM on 20
Secure Sockets Layer (SSL) support 36
security
 DirMaint 37
 file pool 36
 general 35
 TCP/IP for z/VM 36
server features supported by z/VM 106
Server Time Protocol (STP) 48
servers
 requirements 19
 supported by z/VM 87
SFS (Shared File System) 52
Shared Memory Communications - Direct Memory Access (SMC-D) 49
SIMD 51
simultaneous multithreading (SMT) 49
single system image (SSI) cluster
 hardware requirements 22
 program requirements 29
single system image (SSI) clusters
 description 68
SMAPI 53
SNA (Systems Network Architecture)
 AVS 58
 GCS 59
 ISFC 41
software requirements
 base components 26
 Cloud Manager Appliance 30
 DFSMS/VM 27
 DirMaint 28
 facilities and optional features, supported levels of 26
 HCD and HCM for z/VM 28
 High Level Assembler 25
 IBM Wave 28
 installing z/VM from DVD 20
 IOCP 28
 Language Environment 28
 Linux 30
 OpenExtensions 26
 OpenStack 30
 OSA/SF 28
 other products 30
 Performance Toolkit for VM 29
 POSIX 26
 prerequisite licensed programs 25
 RACF Security Server for z/VM 29
 RSCS Networking for z/VM 29
 shared-DASD complex 29
 SSI cluster 29
 TCP/IP for z/VM 30
source code availability 33
spanned HiperSockets channels 44
specialty processors, support for 50
SSI (single system image) cluster
 hardware requirements 22

- SSI (single system image) cluster *(continued)*
 - program requirements 29
- SSI (single system image) clusters
 - description 68
- SSL support 36
- statements of general direction 17
- STHYI instruction 43
- storage
 - management using DFSMS/VM 63
 - minimum and maximum 20
- storage controls supported by z/VM 102
- storage reconfiguration 42
- Store Hypervisor Information (STHYI) instruction 43
- supported device, definition of 23
- switches supported by z/VM 106
- System Delivery Offering (SDO) 32
- system integrity
 - customer responsibilities for data security 35
 - DirMaint 35
 - z/VM, general statement for 34
- system-managed storage 63
- System/370 architecture compatibility 40
- systems management API (SMAPI) 53
- Systems Network Architecture (SNA)
 - AVS 58
 - GCS 59
 - ISFC 41

T

- tape control units supported by z/VM 104
- tape libraries supported by z/VM 102
- tape units supported by z/VM 102
- TCP/IP for z/VM
 - hardware requirements 22
 - overview 56
 - program requirements 30
 - security 36
- terminals supported by z/VM 105
- TLS support 37
- Transactional Execution facility 51
- Transport Layer Security (TLS) support 37
- TSAF (Transparent Services Access Facility) 61

U

- unsupported device, definition of 24

V

- V=V virtual machine 40
- Vector Facility for z/Architecture 51
- virtual machine
 - architectures 39
 - restrictions 31
 - type 40
- virtual machine I/O assist 51
- Virtual Machine Resource Manager (VMRM) 55
- virtual network adapters 40
- virtual switch 41
- VM Dump Tool 41
- VMRM (Virtual Machine Resource Manager) 55
- VMSES/E 61
- VTAM 59

X

- XA virtual machine 39
- XC virtual machine 39, 40
- XEDIT 55

Z

- Z virtual machine 39
- z/Architecture
 - ESA virtual machine 40
 - Z virtual machine 40
- z/Architecture CMS (z/CMS) 51
- z/OS, guest support for 93
- z/TPF, guest support for 94
- z/VM
 - devices supported 97
 - distribution media 32
 - guest operating systems supported 93
 - hardware requirements
 - installation 20
 - Multi-VSwitch Link Aggregation support 23
 - OSA/SF 21
 - PCIe support 23
 - Performance Toolkit 21
 - server 19
 - SSI cluster 22
 - storage 20
 - TCP/IP for z/VM 22
 - workstation 21
 - zHPF guest support 23
 - zHPF paging 23
 - multicultural support 31
 - new functions, enhancements, and changes 7
 - ordering information 32
 - program requirements
 - base components 26
 - Cloud Manager Appliance 30
 - DFSMS/VM 27
 - DirMaint 28
 - facilities and optional features, supported levels of 26
 - HCD and HCM for z/VM 28
 - High Level Assembler 25
 - installing from DVD 20
 - IOCP 28
 - Language Environment 28
 - OpenExtensions 26
 - OpenStack 30
 - OSA/SF 28
 - Performance Toolkit for VM 29
 - POSIX 26
 - prerequisite licensed programs 25
 - RACF Security Server for z/VM 29
 - RSCS Networking for z/VM 29
 - shared-DASD complex 29
 - SSI cluster 29
 - TCP/IP for z/VM 30
 - programs supported 31
 - servers supported 87
 - source code availability 33
 - System Image contents 32
- z/VM HiperDispatch 41
- z/VM Internet Library 69
- z/VM virtual switch 41
- z/VSE, guest support for 94
- zAAP (IBM zEnterprise Application Assist Processor), support for 50

- zHPF (High Performance FICON)
 - description 44
 - guest support hardware requirements 23
 - paging hardware requirements 23
- zIIP (IBM z Integrated Information Processor), support for 50



Product Number: 5741-A07

Printed in USA

GC24-6193-12

