z/VM

Migration Guide

Version 6 Release 4

GC24-6201-12

Note:
Before you use this information and the product it supports, read the information in "Notices" on page 471.

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 GC24-6201-11.

Contents

Figures

Tables

About This Document

This document provides information to help you migrate to IBM® z/VM® V6.4 from:

- A previous release of z/VM (version 6, version 5, version 4, or version 3)
- VM/ESA, version 2

If your current system is older than VM/ESA V2.1, see "Migrating from Older VM Products" on page 3.

This document provides three types of information:

- It describes enhancements and changes to the z/VM system that you should be aware of before migrating.
- It identifies specific external interfaces that have changed and provides an assessment of the compatibility of each change – upwardly compatible or incompatible.
- It provides guidance and procedures for some migration tasks that you might need to perform.

Intended Audience

This document is intended for personnel who are responsible for planning and completing a system migration. Some of this information could also be helpful to programmers who will be migrating applications to the new system.

You should be familiar with z/VM. However, depending on what VM product you are migrating from, z/VM V6.4 might be quite different from your current system. Also, the hardware and software requirements for installing and running z/VM might have changed since your VM release.

Where to Find More Information

For a general description of z/VM V6.4 and information about the current hardware and software requirements, see *z/VM: General Information*. For more information about z/VM functions, see the other books listed in the "Bibliography" on page 477.

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 number:
  z/VM V6.4 Migration Guide
  GC24-6201-12
- 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/).

Summary of Changes

This document contains terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change. Product updates are provided as program temporary fixes (PTFs) for authorized program analysis reports (APARs).

GC24-6201-12, z/VM Version 6 Release 4 (August 2017)

This edition includes information about product changes provided or announced after the general availability of z/VM V6.4.

- Installation, Migration, and Service
  - "[V6.4] Improved upgrade installation planning information" on page 49
- Support and exploitation of hardware and architectures
  - "[V6.4 APAR] z/VM Support for IBM z14" on page 121
  - "[V6.4 APAR] Extended Address Volume Minidisk Support" on page 124
  - "[V6.4 APAR] Processor Scalability Efficiency Improvements" on page 124
  - "[V6.4 APAR] Alternate Subchannel Set Dynamic I/O Support" on page 125
  - "[V6.4 APAR] Coupling over RoCE (CL5) Channel Path Support" on page 125
- Connectivity and networking
  - "[V6.4 APAR] Firewall Friendly FTP" on page 161
- System administration and operation
  - "[V6.4 APAR] z/VM HELP Quality Improvement" on page 238
- System diagnosis
  - "[V6.4 APAR] DUMP Processing Enhancements" on page 269

GC24-6201-11, z/VM Version 6 Release 4 (March 2017)

This edition includes information about product changes provided or announced after the general availability of z/VM V6.4.

- Product packaging
  - "[V6.4] z/VM Omnibus PKIT" on page 11
- Support and exploitation of hardware and architectures
  - [V6.3 APAR] ESA/390 Removal Support Available for z/VM V6.3
  - "[V6.4 APAR] Concurrent I/O Support for IBM XIV Storage System" on page 121
- Connectivity and networking
  - "[V6.4 APAR] Crypto Express APVIRT Support for z/VM TLS/SSL Server and LDAP/VM" on page 159
  - "[V6.4 APAR] Distributed IUCV Enhancements" on page 159
  - "[V6.4 APAR] NICDEF Security Controls" on page 160
- System administration and operation
  - "[V6.4 APAR] z/VM V6.4 Support for OpenStack Newton" on page 236
  - "[V6.4 APAR] RACF Security Policy Enhancements" on page 237
- Application development and deployment

Chapter 1. Introduction to Migration

Migration is the transfer and adjustment of information when upgrading from one product release to a later release. Rather than starting from scratch when you upgrade from your current system, you probably want to transfer and adjust a lot of your current information to use on the new system. Information you may want to transfer includes:

- I/O configurations
- Networking configurations
- Saved segment definitions
- Spool files
- User directory entries
- SFS file pool servers
- Application programs

Before the migration, you need to determine your hardware requirements for the new system, including your DASD, real storage, and performance needs. Preparing for the increased memory scalability of z/VM V6.4 (compared with releases prior to z/VM V6.3) will depend on your current hardware configuration for memory and CPUs. 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 V6.4, additional information can be found at IBM: z/VM V6.4 Resources (www.ibm.com/vm/zvm640/). In particular, you should closely review the *z/VM V6.4 Performance Report*.

Before you begin to transfer any information, you should use this document to determine if there are differences between your current system and the new system that might affect the migration, for example functions that work differently on the new system.

How Migration Information Is Presented

Migration information is presented two ways:

- The topics in Chapter 2, "System Changes," on page 5 provide brief descriptions of significant changes, enhancements, and additions to z/VM system functions. The information is organized according to the general functional areas where the changes have occurred, such as product packaging, installation and service, hardware and architecture support, connectivity and networking, and so on.
- The topics in Chapter 3, "Changes to External Interfaces," on page 279 identify changes to specific external interfaces in the z/VM components. External interfaces are commands, routines, macros, DIAGNOSE codes, directory control statements, and so on. Each change is identified as either ***upwardly compatible*** or ***incompatible***. (For definitions of these terms, see "Compatibility Terms" on page 2.) The information is organized by component, and within each component by interface type.

For each change, the release in which it occurred is indicated in brackets. For example, changes in z/VM V6.4 are indicated like this: **[V6.4]**. Changes are listed from oldest release to newest release.

Notes:

1. The indicated release is either:
   - The first release that included the change in the base product.
   - The last release for which the change was available through service as a program temporary fix (PTF) for an authorized program analysis report (APAR), if no later release existed when the change was added to this publication.
2. A change might be available as an APAR for some releases prior to the indicated release.
3. A change might be superseded by another change, or support might be withdrawn, in a later release.
4. For types of changes that occur every release, such as responses from commands that query the level of the system, only the latest release is indicated.

The topics in Chapter 4, "Migration Tasks," on page 415 provide guidance for some migration tasks you might have to do, such as converting from using system definition macros (HCPRIO, HCPSYS, and HCPBOX) to using system configuration files, or migrating your user directory and your spool files.

Compatibility Terms

In Chapter 3, "Changes to External Interfaces," on page 279, the following terms are used to convey the degree of compatibility for each change:

Upwardly compatible
: The syntax, function, or response of the external interface has been changed, but not significantly. Invocations and applications using the external interface on the new system in the same manner as on the current system should continue to execute unchanged.

  Note: New function, if not exploited, is also upwardly compatible.

Incompatible
: The syntax, function, or response of the external interface has been changed significantly. Some invocations and applications using the external interface on the new system in the same manner as on the current system might execute differently, incorrectly, or not at all.

  Note: Depending on how you use the interface, a change identified as incompatible might be upwardly compatible for you.

Cross-References to Other Documents

Many of the descriptions of system changes contain cross-references to other documents for more information about those functions. Over the course of product releases, the titles of some VM documents have changed. In most cases, the cross-reference points to the *current* title of the appropriate document.

Migrating from Older VM Products

If you are migrating from a VM product older than VM/ESA V2.1, you need an additional document to get information about all of the changes to VM between your release and z/VM V6.4, as shown in the following table.

Table 1. Documents You Need for Migrating from Older VM Products

If you are migrating from:	Use these documents:
VM/SP 5 or 6 VM/SP HPO 5 or 6 VM/ESA V1.1.0 (370 Feature)	1. Use *VM/ESA V2.1: Conversion Guide and Notebook for VM/SP, VM/SP HPO, and VM/ESA (370 Feature)*, SC24-5754, to determine the changes between your VM product and VM/ESA V2.1 2. Then use this document to determine the changes between VM/ESA V2.1 and z/VM V6.4
VM/XA SP 2 or 2.1 VM/ESA V1.1.0 (ESA Feature) VM/ESA V1.1.1	1. Use *VM/ESA V2.1: Conversion Guide and Notebook for VM/XA SP and VM/ESA*, SC24-5753, to determine the changes between your VM product and VM/ESA V2.1 2. Then use this document to determine the changes between VM/ESA V2.1 and z/VM V6.4
VM/ESA V1.1.5 370 Feature VM/ESA V1.2.0 VM/ESA V1.2.1 VM/ESA V1.2.2	1. Use *VM/ESA V2.4: Conversion Guide and Notebook*, GC24-5839, to determine the changes between your VM product and VM/ESA V2.4 2. Then use this document to determine the changes between VM/ESA V2.4 and z/VM V6.4

These additional documents can be obtained from the following sources:

- IBM: z/VM Internet Library (www.ibm.com/vm/library/)
- IBM Publications Center (www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss)

Considerations for Future Migration Automation

To take advantage of any automated migration aids supplied in the future by IBM, you must adhere to the following:

- If you need to alter, modify, or customize any IBM supplied parts, always use the VMSES/E local modification procedures unless it is specifically documented to do otherwise. The LOCALMOD tool is supplied with VMSES/E to simplify the creation of local modifications.
- If you want to use an IBM supplied shared file system (SFS) for your data, use VMSYSU. Do not use VMSYS, VMSYSR, or VMPSFS.
- Each customizable file must reside on the disk specified for that part in the product documentation.
- Additional minidisks should not be defined on the *xxx*RES, *xxx*W01, and *xxx*W02 DASD volumes for the new release.

Chapter 2. System Changes

These topics describe significant changes, enhancements, and additions to z/VM system functions. Changes made since your current release can affect your migration to the new z/VM V6.4 system. For information about changes to specific external interfaces, see Chapter 3, "Changes to External Interfaces," on page 279.

The system changes are described in the following major topics:

Product Packaging

These topics describe changes in what facilities are provided in the z/VM base product or offered as optional features, and how z/VM is distributed.

[V4.1] CMS Utilities Feature Integrated into CMS

Most of the functions formerly provided in the optional CMS Utilities Feature have been integrated into CMS.

The following functions are located on the 190 disk:

- BROWSE
- CLRSCRN
- DEPRINT
- DEVTYPE
- DIRMAP
- FILESTCK
- FINDSTAK
- FLIST
- GETFMADR
- OPTIMISE
- REPRINT
- SADT
- SETKEYX
- SHRLDR
- STAG
- USERID
- VMSIZE
- WAKEUP
- XRDR
- YDISK

The following functions, which may require privilege classes not given to general users, are located on the 193 disk:

- ACCOUNT
- AUDITOR
- DCSSBKUP
- DCSSRSAV
- QSYSOWN
- SFPURGER
- SYSWATCH

All of these functions (except OPTIMISE and SETKEYX) are documented in *z/VM: CMS Commands and Utilities Reference*. The documentation for these functions has been substantially updated and improved.

All messages issued by these functions are contained in the CMS message repository, and documentation for the messages has been added to *z/VM: CMS and REXX/VM Messages and Codes*.

OPTIMISE and SETKEY have been included in CMS only for compatibility and are documented only in HELP.

The following functions have not been integrated into CMS:
- CPFMT
- SYSIDT

[V4.1] TCP/IP Included in the z/VM Base

TCP/IP, previously offered as an optional feature of VM, is now supplied with the z/VM base. The TCP/IP Network File System (NFS) Feature has been integrated into the TCP/IP base. TCP/IP source, previously offered as a feature, is also supplied with the z/VM base.

[V4.1] Directory Maintenance Facility Feature

Directory Maintenance Facility (DirMaint™), based on the Directory Maintenance VM/ESA licensed program, is offered as an optional feature of z/VM.

[V4.1] REXX/EXEC Migration Tool for VM/ESA Deleted

The REXX/EXEC Migration Tool for VM/ESA (ESAMIGR) is no longer supplied with z/VM.

[V4.1] Distributed Computing Environment Feature Deleted

The Distributed Computing Environment (DCE) Feature is not supported on z/VM V4.1 and can no longer be ordered.

[V4.1] LANRES/VM Feature Deleted

LANRES/VM is no longer offered as a feature of z/VM.

[V4.1] VM BookManager Library Feature Deleted

The VM BookManager® Library Feature, which provided BookManager versions of VM publications on the product delivery medium, is no longer offered with z/VM. BookManager versions of VM publications are still included on the *Online Library Omnibus Edition: VM Collection* CD-ROM. The CD-ROM includes the Softcopy Receiver Tool for transferring BookManager files to the workstation or host.

[V4.3] Resource Access Control Facility Feature

Resource Access Control Facility (RACF®) is offered as an optional feature of z/VM. RACF works with existing z/VM system functions to provide improved data security for an installation.

The RACF feature is functionally equivalent to the RACF for VM V1R10.0 licensed program (5740-XXH). However, only the RACF feature is licensed to run on z/VM in an Integrated Facility for Linux (IFL) LPAR.

[V4.4] Hardware Configuration Definition and Hardware Configuration Manager for z/VM

Hardware Configuration Definition and Hardware Configuration Manager for z/VM is a new facility included in the z/VM base and pre-installed. For more information about this facility, see "[V4.4] Hardware Configuration Definition and Hardware Configuration Manager for z/VM" on page 173.

[V4.4] Language Environment Integrated as a Component

Language Environment® is now a component of z/VM. The C/C++, COBOL, and PL/I run-time libraries included in the Language Environment component have been upgraded to the level shipped with version 1, release 4.0 of z/OS®.

The Language Environment component is installed and serviced under the MAINT user ID and uses a set of disks that belong to MAINT. Language Environment is included in the ZVM $PPF file. In addition, files with file types of LOCALE, LOCALEX, CHARMAP, GENXLT, UCMAP, UCONVTBL and a few EDC$ ASSEMBLE files now will reside on the Language Environment 49B disk.

Attention: Do not migrate Language Environment 1.8 or earlier to z/VM V4.4. The only level of Language Environment supported on z/VM V4.4 is the new Language Environment component. The files for the z/VM Language Environment component are installed on the MAINT 19E disk. Depending on how you plan to migrate other files from your old 19E disk to the new system, you may need to remove the old Language Environment files. See "Removing the Old Level of Language Environment."

Removing the Old Level of Language Environment

The files for the z/VM Language Environment component are installed on the MAINT 19E disk. If you are migrating files from your old 19E to the new 19E, do not copy any Language Environment files. If you intend to use your old 19E disk with the new z/VM V4.4 system, and want to copy the Language Environment component from the new 19E to the old 19E, you must first remove all Language Environment files from your old 19E.

To remove the Language Environment files from your old 19E disk:

1. Log on as MAINT.
2. Access the 19E disk in R/W mode:

   ```
   access 19e y
   ```

3. Use the VMSES/E VMFERASE command to erase the Language Environment files from Y. Enter:

   ```
   vmferase prod prodid%compname from y
   ```

   where *prodid* and *compname* identify the level of Language Environment on your system. See Table 2 on page 8.

Table 2. Language Environment level identifiers

Language Environment level	*prodid*	*compname*
1.5 (C run-time only), included with VM/ESA V2.1	5688198E	POSIX
1.5 (complete), available for VM/ESA V2.1	5688198E	LE370
1.6 (C run-time only), included with VM/ESA V2.2	5688198F	POSIX
1.6 (complete), available for VM/ESA V2.2	5688198F	LE370
1.8, available for VM/ESA V2.3 and later	5688198H	LE370

4. If the VMFERASE command returns the following message,

   ```
   VMFERA2738I No files to erase from Y for product prodid%compname
   ```

   then the files are not listed in the VMSES PARTCAT file and cannot be erased with the VMFERASE command. You must use the CMS ERASE command to manually erase all files on Y whose file names begin with the characters EDC, CEE, IBM, IGZ, or SCEE.
5. After the Y-disk is cleaned up, you must remove the Language Environment shared segments. Enter the following commands:

   ```
   cp purge nss name scee
   cp purge nss name sceex
   ```
6. Because the file status table for the Y-disk is included in the CMS saved system, and you have changed the contents of the Y-disk, you must resave CMS. Use the SAMPNSS command to create a new skeleton system data file, then IPL and save the rebuilt CMS saved system:

   ```
   sampnss cms
   ipl 190 clear parm savesys cms
   ```
7. Now that you have removed the old Language Environment files and segments from your system, you must decided what to do with the user ID to which the Language Environment files were initially loaded (P688198E, P688198F, or P688198H, depending on the level of Language Environment) You can either delete this user ID from the CP directory and reuse all the disk space assigned to it, or you can leave the user ID on your system as a means of returning to the old Language Environment level if an application requires it.

[V4.4] Performance Toolkit Feature

The Performance Toolkit for VM is offered as an optional feature of z/VM. It is supplied pre-installed but disabled, and is serviced through VMSES/E. The Performance Toolkit is intended as a replacement for the RealTime Monitor (RTM) and Performance Reporting Facility (PRF) optional features, and provides enhanced functions. For more information, see "[V4.4] Performance Toolkit for VM" on page 176.

[V5.1] Packaging Modifications

The packaging of the z/VM product and the z/VM System Delivery Offering (SDO) have been modified as follows:

- The 3270 PC File Transfer product (5664-281), which previously could be ordered with the z/VM V4 SDO, is now available as part of the z/VM V5.1 product. It is shipped as a sample program on the system disk, with no support available.
- The restricted source feature, available as a no-charge feature that could be ordered with z/VM V4, and the PL/X source, provided with the installation

media in z/VM V4, are not shipped with z/VM V5.1 and cannot be ordered with z/VM V5.1. Both are available as no-charge downloads from IBM Resource Link (www.ibm.com/servers/resourcelink/).

Note: If you are not registered with Resource Link®, you will be required to register for a user ID and password. You must also be licensed for z/VM V5.1, and entitlement will be verified when you request the source code. After approval, you will receive instructions describing how to download the code.

- DFSMS/VM, previously provided automatically with z/VM V4, is no longer shipped automatically. It can be ordered as a no-charge feature with the z/VM V5 SDO.
- The national language features for Interactive System Productivity Facility (ISPF), program number 5684-043, have been removed from the z/VM SDO. These features can still be ordered using the standalone ordering process.
- Tivoli® Storage Manager (TSM) for VM, Version 4 (5697-TS9) Release 4.2, which was packaged on the z/VM V4 system DDRs, is not provided with z/VM V5.1 and cannot be ordered with the z/VM V5 SDO. To protect data on z/VM, IBM recommends TSM for z/OS and z/VM, V5.2 (5698-A13), or TSM Extended Edition for z/OS and z/VM, V5.2 (5698-A11). If you want to run TSM in a Linux guest environment on z/VM V5.1, IBM recommends TSM, V5.2 (5608-ISM), or TSM Extended Edition, V5.2 (5608-ISX).
- The Java™ and NetRexx programs on z/VM, previously available as downloads from the z/VM Web site, are not supported with z/VM V5.1. Java and NetRexx functionality can be obtained from Linux running on z/VM.
- The RealTime Monitor (RTM) and Performance Reporting Facility (PRF) features cannot be ordered with z/VM V5.1, nor can they be licensed with z/VM V5.1.
- The IBM Debug Tool for z/VM, V4.1 (5654-A23), is available as a chargeable feature of the z/VM V5 SDO.

[V5.2] Packaging Modifications

The packaging of the z/VM product and the z/VM System Delivery Offering (SDO) have been modified as follows:

- z/VM V5.2 will not be available on CD-ROM. z/VM continues to be distributed on 3480/3490 and 3590 tape and on DVD.
- The Tivoli Storage Manager for VM product was removed from the z/VM V5.1 system image and has not been orderable with the z/VM SDO since September 24, 2004. With the announcement of the Tivoli Storage Manager V5.3 family of products in the Tivoli Storage Manager announcement dated December 14, 2004, support for z/VM was not migrated to Tivoli Storage Manager V5.3 nor is it supported.

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

[V5.3] RACF Security Server Feature

IBM RACF Security Server for z/VM, function level 530, is offered as an optional feature of z/VM V5.3. It is the successor to the IBM Resource Access Control Facility for z/VM feature, V1R10. The RACF Security Server for z/VM also provides some enhanced functions. For more information, see "[V5.3] RACF Security Server for z/VM" on page 193.

[V5.3] RSCS Networking Feature

IBM Remote Spooling Communications Subsystem (RSCS) Networking for z/VM, function level 530, is offered as an optional feature of z/VM V5.3. It is supplied preinstalled (disabled) and is serviced through VMSES/E. RSCS Networking for z/VM replaces the IBM RSCS V3.2 product (5684-096), which was previously packaged with z/VM. RSCS Networking for z/VM is licensed under the IPLA terms and conditions for z/VM, and is available for both IFL and standard processor configurations. RSCS Networking for z/VM also provides some enhanced functions. For more information, see "[V5.3] RSCS Networking for z/VM" on page 141.

[V5.3] Availability of 3592 Tape Media for Ordering

Effective September 14, 2007, z/VM V5.3 is available for ordering on 3592 tape media.

[V5.4] 3480 Tape Media Discontinued for Ordering

z/VM V5.4 cannot be ordered on 3480 tape media. z/VM V5.4 is available only on 3590/3592 tape, on DVDs, and by Internet delivery through ShopzSeries.

[V6.1] Changes to Documentation Provided with z/VM

The *IBM Online Library: z/VM Collection* CD-ROM has been replaced by the *IBM Online Library: z/VM Collection* DVD, SK5T-7054. The DVD is provided with z/VM V6.1. The CD-ROM version has been discontinued.

The following printed documents are not provided with z/VM V6.1:

- *z/VM: Getting Started with Linux on z Systems*

  The printed version of this document has been discontinued. PDF and BookManager versions are still available from the same sources as other z/VM publications.
- *z/VM Summary for Automated Installation and Service (DVD Installation)*

  This document has been discontinued (no longer published in any format).
- *z/VM Summary for Automated Installation and Service (Tape Installation)*

  This document has been discontinued (no longer published in any format).

[V6.1] German Files Discontinued

Translation of message repositories and HELP files into German has been discontinued. German files are not shipped with or available for z/VM V6.1. The z/VM installation procedures have been revised to remove the German option.

[V6.2] IBM z/VM Single System Image Feature

The IBM z/VM Single System Image Feature (VMSSI) is offered as an optional feature of z/VM V6.2. VMSSI enables the creation of z/VM single system image (SSI) clusters. A z/VM SSI cluster is a multisystem environment in which the z/VM member systems can be managed as a single resource pool and running Linux guests can be relocated from one member to another. For information about the z/VM SSI environment and setting up z/VM SSI clusters, see *z/VM: CP Planning and Administration*.

VMSSI must be ordered through the z/VM System Delivery Offering (SDO). For information about the licensing requirements for VMSSI, see *z/VM: License Information*.

Various z/VM functions have been added or enhanced to support z/VM SSI clusters. Those additions and enhancements are described in the appropriate sections of this publication.

[V6.2] z/Architecture CMS

z/Architecture® CMS (z/CMS), previously supplied as a sample program, is now a fully supported part of z/VM. z/CMS runs in IBM z/Architecture 31-bit addressing mode in an ESA or XA virtual machine and enables CMS programs to use z/Architecture instructions, including those that operate on 64-bit registers, while permitting existing ESA/390 architecture CMS programs to continue to function without change. Although z/CMS 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 specify the SUBPOOL='USERG' parameter on the CMSSTOR OBTAIN macro to allocate storage above 2 GB. For more information about z/CMS, see *z/VM: CMS Planning and Administration*.

The following functions have been updated:

- CMS QUERY CMSLEVEL command
- CMS SET CMS370AC command
- CMS CMSSTOR macro
- CMS DEFNUC macro
- REXX/VM STORAGE() function

[V6.2] Shipping IBM Systems Director Agents for Linux with z/VM Discontinued

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

[V6.3] xCAT Appliance Packaged with z/VM

The Extreme Cloud Administration Toolkit (xCAT) appliance is packaged with z/VM V6.3. xCAT is an open source scalable distributed computing management and provisioning tool that provides a unified interface for hardware control, discovery, and OS diskful/diskfree deployment. The toolkit can be used for the deployment and administration of Linux clusters. The toolkit is packaged and configured to be ready for use with no modification.

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

[V6.3] z/VM Collection Zip File

The *IBM Online Library: z/VM Collection*, SK5T-7054, is no longer available as a physical DVD and is no longer supplied with the z/VM product. The *z/VM Collection* is now available as a zip file that can be downloaded from IBM Publications Center (www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss).

[V6.4] z/VM Omnibus PKIT

The *IBM Online Library: z/VM Omnibus PKIT*, SK5T-9509, has replaced the *IBM Online Library: z/VM Collection*, SK5T-7054. Like the *z/VM Collection* (CKIT), the *z/VM Omnibus PKIT* contains product information libraries (PKITs) for current z/VM releases and current IBM licensed programs that run on z/VM. It is

available as a zip file that can be downloaded from IBM Publications Center (www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss), and it will be updated quarterly as needed.

Note: The zip file is very large (over 2450 MB), and you need to allow sufficient time for the download.

Installation, Migration, and Service

These topics describe changes to the z/VM installation, migration, and service processes, procedures, and tools.

For more information about installation changes, see *z/VM: Installation Guide*. For more information about service changes, see *z/VM: Service Guide*.

[V2.2] Electronic Delivery of Service

The ITNVTSTR EXEC supports the delivery of IBM service electronically by satellite. This electronic delivery is provided by the IBM Advanced Digital Delivery System product (5799-XDG).

For more information, see:

- *z/VM: VMSES/E Introduction and Reference*
- *Advanced Digital Delivery User's Guide*, SC23-3281

[V2.2] Local Modification of Replacement-Maintained Parts

The VMFREPL EXEC supports the local modification of replacement-maintained parts. VMFREPL can be used to:

- Copy the highest level of a part
- Copy a specified part
- Update a Version Vector Table
- Update a Select Data file
- Display the highest level of a part

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

[V2.3] VMSES/E Enhancements

The following execs have been added:

VMFENRPT
: Creates a report of the products that are enabled, disabled, and installed on your system.

VMFSUFIN
: Installs service from RSU service envelope files, COR service envelope files, or both.

VMFSUFTB
: Builds a table, sysid SYSSUF, that contains a list of all installed products and related data needed to service each product.

VMFINS DISABLE
: Changes a product to a disabled state.

VMFINS ENABLE
: Changes a product to an enabled state.

The following Software Inventory tables have been added:

VM SYSSUF (System-Level Service Update Facility table)
Contains a list of all products that are installed on the system. For each product, it contains the data needed to service the product.

VM SYSREST (System-level Restart table)
Contains records used to restart the VMFSUFIN EXEC.

The following execs have been changed:

ITNVTSTR
The PROD and KEY operands have been added. The PROD operand identifies the products that were ordered for a RSU package. The KEY operand indicates that the RSU package is one of two packages that are to be installed together.

VMFINS DELETE
The DISABLE operand has been added. The DISABLE operand sets up a product as disabled and deleted.

VMFINS MIGRATE
The DISABLE and ENABLE operands have been added. The DISABLE operand sets up a product as disabled. The ENABLE operand sets up a product as enabled.

VMFINS INSTALL
The DISABLE, ENABLE, NOSETUP, and SETUP operands have been added. The DISABLE operand sets up a product as disabled. The ENABLE operand sets up a product as enabled. The NOSETUP operand indicates that a new minidisk or directory access order is not set up. The SETUP operand indicates that a new minidisk or directory access order is set up. It is set up according to entries in the :MDA section of the product parameter file.

The following Software Inventory table has been changed:

VM SYSAPPS
The :ESTAT tags has been added. The :ESTAT tag specifies the enablement status of a product on the system.

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

[V2.4] Tool for Removing PTFs

The VMFREM EXEC removes individual PTFs by "un-applying" them from all service levels (apply disks) and optionally "un-receiving" them. To "un-apply" a PTF means to undo the function previously performed for that PTF by the VMFAPPLY command. To "un-receive" a PTF means to undo the function previously performed for that PTF by the VMFREC command.

VMFREM also removes complete service levels and optionally "un-receives" PTFs that are applied only to the removed levels. In addition, commit support is provided for individual PTFs that have been applied. For more information, see *z/VM: VMSES/E Introduction and Reference*.

[V3.1] $PPF Override File Name

The IBM-supplied override product parameter file (PPF), which contains overrides to the base $PPF files for each component, has changed to ZVM $PPF. The override PPF for VM/ESA was ESA $PPF.

[V4.1] Installation Changes

The following changes and enhancements have been made to the installation procedure:

- A new "express" installation method using predefined defaults is available, making it faster and easier for you to install and service z/VM. There are some restrictions when using the express installation method:
  - Only IBM supplied PPFs are used.
  - VM source code is not installed.
  - Customer local modifications are not allowed.
  - Products and features are installed onto minidisks only. You cannot move them to SFS.
  - Only the SMALL FILEPOOL is provided (no large VMSYS (SFS) filepool).
  - Only one DASD type and model can be used for your installation.
- The new SERVICE and PUT2PROD commands automate the application of an RSU and CORrective service. The SERVICE command installs an RSU or applies CORrective service for z/VM components, features, or products that are installed on the z/VM system DDR. The PUT2PROD command places components, features, or products that were serviced using the SERVICE command into production.

  All customers can use these commands at installation time. However, after installation is complete, they may be used only by express customers.
- DASD types 9345 and FBA are not supported.

[V4.1] VMSES/E Enhancements

The following commands have been added:

PUT2PROD
: Places a component, feature, or product that was serviced by preventive or corrective service into production.

SERVICE
: Installs an RSU or applies CORrective service for the z/VM components, features, or products that are installed on the z/VM system DDR.

[V4.3] Installation Changes

The following changes and enhancements have been made to the installation procedure:

- 3590 tape drive is now supported for installation.
- Non-XF 3480 tape drive is no longer supported for installation.
- 4mm DAT tape is no longer supported for installation.
- 3380 DASD is no longer supported for installation.
- Mixed DASD are no longer supported for installation. All 3390 DASD used for installation must be the same model.
- TSAF and AVS are no longer optionally installed; they are now part of the base installation.

[V4.3] TCP/IP Configuration Wizard

A new utility automates the connection of a newly-installed z/VM system to a TCP/IP-based network. The TCP/IP configuration wizard, IPWIZARD, requires no knowledge of z/VM TCP/IP and is similar to the network configuration utilities used in Linux for zSeries distributions during Linux installation. This easy-to-use configuration wizard assists the z/VM installer in providing desired Internet Protocol (IP) configuration information such as host and domain name, IP addresses, and subnet mask. From that information, the wizard generates an initial TCP/IP configuration (creating the SYSTEM DTCPARMS, TCPIP DATA, and PROFILE TCPIP files) and verifies that connectivity to the network has been established. For more information, see *z/VM: Installation Guide*.

[V4.3] Service Enhancements

The SERVICE tool can now detect local modifications and present that information to you. This allows you to rework your local modifications before the new service is built into executables. The new VMFUPDAT command provides a panel interface that displays which local modifications need to be reworked and allows updates to the System-Level Service Update Facility table. For more information, see *z/VM: VMSES/E Introduction and Reference*.

New functions in the VMFUPDAT command allow you to:

- Change the INSTALL, BUILD, INCLUDE, INSPPF, BLDPPF, and P2PPF tags in the System-Level Status table
- Change the local modification rework status in the System-Level Local Modification table

[V4.4] Integrated 3270 Console

z/VM V4.4 supports real and virtual integrated 3270 console devices. Real support enables this device, which is provided through a Hardware Management Console, to be used as a z/VM system operator console. This removes the requirement to have an external 3270 device to install and service z/VM V4.4. The z/VM Stand Alone Program Loader (SAPL) and stand-alone DASD Dump-Restore (DDR) program support the use of the integrated 3270 console as a system operator console. Virtual support enables testing of guest operating systems and utilities that support the integrated 3270 console device.

Note: There may be additional hardware requirements to use this support. See the section on server support in *z/VM: General Information*.

The following CP function has been updated:

- OPERATOR_CONSOLES system configuration statement

For more information, see:

- *z/VM: CP Planning and Administration*
- *z/VM: System Operation*

[V4.4] Installation Enhancements

The following improvements have been made to z/VM installation:

- There are fewer choices during installation, making z/VM easier to install.
- Language Environment is now a component of z/VM and has been upgraded to a new level.

Attention: Do not migrate Language Environment 1.8 or earlier to z/VM V4.4. The only level of Language Environment supported on z/VM V4.4 is the new Language Environment component. The files for the z/VM Language Environment component are installed on the MAINT 19E disk. Depending on how you plan to migrate other files from your old 19E disk to the new system, you may need to remove the old Language Environment files. See "Removing the Old Level of Language Environment" on page 7.

- A direct installation into SFS is now supported. Previously, the installation process installed all products, features, and components onto minidisks, then copied files to SFS.
- Service disks for all z/VM components (CP, CMS, VMSES/E, Dump Viewing Facility, REXX/VM, TSAF, AVS, GCS, and Language Environment) can now be either on minidisks or in SFS. Previously, only service disks for TSAF, AVS, GCS, and Language Environment could be in SFS.

  Note: The service disks for the z/VM components must be all on minidisks or all in SFS; you cannot have some components on minidisks and some in SFS.
- There are no longer separate source disks for the base components of z/VM; therefore, source code will be placed on the component's base disk.
- Source code, OSA/SF, and Tivoli Storage Manager are no longer optionally installed; they are now part of the installation base. (Tivoli Storage Manager is installed disabled.)
- The steps to move a component or product into SFS directories are more automated.
- The new HCD and HCM for z/VM facility is pre-installed.
- The new Performance Toolkit feature is pre-installed (disabled).
- A new level of OSA/SF, 440, is pre-installed.
- A new release of ICKDSF, 1.17.0, is pre-installed.

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

[V4.4] Service Enhancements

New service enhancements are as follows:

- Service processes and EXECs now support z/VM base component service disks either on minidisks or in SFS directories.
- TSAF and AVS now share SFS directories.
- Changes to the VMFREM command allow you to remove local modifications.
- New message log support for the SERVICE and PUT2PROD EXECs places all console messages into VMSES/E-formatted $MSGLOG files. This allows you to use VMFVIEW to view error and warning messages quicker and easier than in a console log.
- New functions in the VMFUPDAT command allow you to:
  - Change the manual status in the Service-Level Build Status table
  - Delete restart records in the System-Level Restart table or SERVICE $RESTART (from SERVICE EXEC) file.
- The VMFBDPMD part handler has been updated to allow building modules using **c89** and the CMS binder.

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

[V4.4] Product Parameter File Changes

The shipped ZVM $PPF file now includes the Language Environment component.

The ZVM $PPF file and all base component product parameter files ($PPFs) now have an SFS component name for each component.

TSAF and AVS now share SFS directories.

There are no longer separate source disks for the base components of z/VM; therefore, source code will be placed on the component's base disk.

[V5.1] Installation of z/VM from DVD

z/VM V5.1 includes the capability to install z/VM from a DVD to an ESS SCSI disk emulated as an FBA device or to a 3390 DASD. Installing from a DVD can significantly reduce the required installation media and allows you to install to a zSeries server using only SCSI disks. This capability requires the Hardware Management Console, Version 1.8 or later.

The following CP functions have been added for this support:
- DVDPRIME utility
- INSTDVD utility

The following CP functions have been updated:
- INSTPLAN utility
- INSTVM utility

For more information, see:
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: Installation Guide*

[V5.1] CP Is Now 64-Bit Only

z/VM V5.1 includes only a 64-bit CP module. The 32-bit and dual-mode CP modules included with previous releases are no longer provided. VMSES/E commands and files are affected as follows:

- The VMFHLASM and GENCPBLS commands now have NODUAL as the default.
- In the CP product parameter file, 5VMCPR10 $PPF:
  - The ALTCNTRL tag has been removed from the CNTRLOP section.
  - The HCPESAME, CPLOAD32, CPLOAD64, and HCPBLNUC build lists have been removed from the BLD section.
  - The CPLOAD build list has been added to the BLD section. This build list builds one CP nucleus, which is the 64-bit version.

    Note that the parts included in the build for the single 64-bit CP nucleus have the file type TEXT, not TXTESAME. Also, there is now only one CP control file, HCPVM CNTRL.

[V5.1] Installation and Service Enhancements

The automated service process has been updated to include:

- A new MAINT 500 minidisk, which may be used for loading certain service envelopes
- An easier query of RSU levels and individual PTF levels for a component

- Cataloging service memo files online and easily displaying them using the VMFUPDAT command
- An easier way to browse server restart records using VMFUPDAT SYSREST
- A new command, LOCALMOD, to automate the local modifications procedure

To help improve the ease-of-use for installing Linux with your z/VM system, spool and page space has been removed from the System Residence volume, and separate installation volumes for spooling and paging are now assigned by the user.

Installation of z/VM to 3390-1 or 3390-2 DASD is no longer supported and has been removed from the INSTPLAN utility.

For more information, see:

- *z/VM: Installation Guide*
- *z/VM: VMSES/E Introduction and Reference*

[V5.1] Relocation of the BFS /etc Directory

The BFS /etc directory, which contains the configuration files for the OpenExtensions Shell and Utilities, has been moved from the VMSYS file pool to the VMSYSU file pool. For more information, see *z/VM: CMS File Pool Planning, Administration, and Operation*.

[V5.2] Changes to Predefined User Directory Entries

Deleted User IDs

- LINUXSAV
- VMADMIN

Added User IDs

DTCVSW1
: System-default virtual switch (VSWITCH) controller virtual machine

DTCVSW2
: System-default virtual switch (VSWITCH) controller virtual machine (alternate, for provision of failover capability)

MIGMAINT
: User ID used with migration procedures

Renamed User IDs

- 5VMPTK10 renamed to 5VMPTK**20**.
- 5VMTCP10 renamed to 5VMTCP**20**.
- 4VMHCD40 renamed to **5**VMHCD**20**.

Changed User Specifications

- MAINT
  - Virtual storage increased:
    [V5.1] `USER MAINT MAINT 64M 1000M ABCDEFG`
    [V5.2] `USER MAINT MAINT 128M 1000M ABCDEFG`
  - Disk sizes for CF1, CF2, and CF3 increased to 120 cylinders.
- 5VMTCP20
  - Virtual storage increased:
    [V5.1] `USER 5VMTCP10 5VMTCP10 16M 32M BEG`

```
[V5.2] USER 5VMTCP20 5VMTCP20 100M 132M BEG
```

  - Links added:

```
LINK DTCVSW1 191 280 MR
LINK DTCVSW2 191 281 MR
```

- TCPIP
  - Privilege class changed:

```
[V5.1] USER TCPIP TCPIP 32M 128M ABCG
[V5.2] USER TCPIP TCPIP 32M 128M ABG
```

- CBDIODSP
  - Virtual storage increased:

```
[V5.1] USER CBDIODSP CBDIODSP 96M 128M BCEFG
[V5.2] USER CBDIODSP CBDIODSP 256M 999M BCEFG
```

[V5.2] Changes to Supplied SYSTEM CONFIG Definitions

SYSTEM_USERIDS statement:

- Removed the 'DISCONNECT' operand from the OPERATOR user ID.

FEATURES statement:

- Removed 'LOGMSG_FROM_FILE' from the ENABLE operand.
- Changed the RETRIEVE DEFAULT value from '7' to '20'.
- Changed the RETRIEVE MAXIMUM value from '7' to '255'.
- Added 'VDISK Userlim 144000 blocks'.

OPERATOR_CONSOLES statement:

- Moved 'SYSTEM_3270' ahead of 'SYSTEM_CONSOLE'.

EMERGENCY_MESSAGE_CONSOLES statement:

- Added 'SYSTEM_CONSOLE' at the end of the list.

CHARACTER_DEFAULTS statement:

- Changed 'CHAR_DELETE @' to 'CHAR_DELETE OFF'.
- Changed 'LINE_DELETE ¢' to 'LINE_DELETE OFF'.
- Changed 'TAB]' to 'TAB OFF'.

[V5.2] Migration and Service Enhancements

A simplified migration process and new migration tools allow the disks associated with the preinstalled z/VM components, facilities, and optional features (and the preinstalled ICKDSF and RSCS products) on your first-level z/VM V5.1 system to be made available to your second-level z/VM V5.2 system and then transfers the following types of files:

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

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

z/VM V5.1 provided a new command, LOCALMOD, to make the procedure for creating and applying local modifications easier. z/VM V5.2 provides further automation of the LOCALMOD process by adding a capability to rework local

modifications and provide support for local service. For more information, see *z/VM: VMSES/E Introduction and Reference*.

[V5.2] Support for IPv6 HiperSockets in TCP/IP Configuration

The basic TCP/IP configuration wizard (IPWIZARD) allows you to select IPv6 HiperSockets™ as your interface type.

[V5.3] Installation Changes

The following changes and enhancements have been made to the installation procedure:

- New 530 levels of Directory Maintenance Facility (DIRM), Performance Toolkit for VM, RACF Security Server for z/VM, RSCS Networking for z/VM, and TCP/IP are preinstalled.
- New support is provided that allows uploading the installation DVD media to an FTP server directory or VM minidisk and then installing from the FTP server or VM minidisk.
- Installation first level from DVD requires 576 MB of real storage.
- Installation might require additional volumes, depending on where products are installed.

[V5.3] Migration and Service Enhancements

Support is provided for simplified migration to z/VM V5.3 from either z/VM V5.2 or z/VM V5.1.

z/VM V5.3 includes the following service enhancements:

- The SERVICE, VMFSUFIN, and VMFUPDAT execs have been updated with new operands.
- The :PRODLEV column was added to VM SYSSUF.
- The :PPFVAR tag was added to VMSESE PROFILE.
- A new table, service-level production status (prodid SRVPROD) was added.
- You can now use the LOCALMOD command to apply local service to any component or product defined in the System-Level Service Update Facility (VM SYSSUF) software inventory table.

[V5.3] Changes to Predefined User Directory Entries

Deleted User IDs

- BOOTPD
- ROUTED

Added User IDs

LDAPSRV
LDAP server

RSCSAUTH
RSCS dynamic authorization server

SNMPSUBA
SNMP sub-agent server

VSMPROXY
Systems management API integrated systems management proxy server

VSMREQIN
Systems management API INET request server

VSMREQIU
Systems management API IUCV request server

VSMWORK1
Systems management API worker server

VSMWORK2
Systems management API worker server

VSMWORK3
Systems management API worker server

ZVMMAPLX
z/VM Managed Access Point server

Renamed User IDs

- 5VMDIR10 renamed to 5VMDIR**30**.
- 5VMPTK20 renamed to 5VMPTK**30**.
- 5VMTCP20 renamed to 5VMTCP**30**.
- 5767000P renamed to **5VMRAC30**.
- P684096K renamed to **5VMRSC30**.

Changed User Specifications

- MAINT
  - Minidisk 990 added.
  - Minidisk 51D increased to 26 cylinders.
- ADMSERV
  - Default storage size increased to 32M.
- BLDCMS
  - SPOOL 00C statement changed from class A to class *.
- BLDNUC
  - SPOOL 00C statement changed from class A to class *.
- BLDRACF
  - SPOOL 00C statement changed from class A to class *.
- BLDSEG
  - SPOOL 00C statement changed from class A to class *.
- DIRMAINT
  - DIAG88 added to the OPTION statement.
- IMAPAUTH
  - Default storage size increased to 32M.
- LPSERVE
  - Default storage size increased to 32M.
- MIGMAINT
  - Updated to more closely match MAINT.
- NAMESRV
  - Default storage size increased to 32M.
- NDBPMGR
  - Default storage size increased to 32M.
- NDBSRV01
  - Default storage size increased to 32M.
- PERFSVM

  - DIAG88 added to the OPTION statement.
- PORTMAP
  - Default storage size increased to 32M.
- RACFVM
  - Added the following statement:

    ```
    IUCV ALLOW MSGLIMIT 255
    ```
- RSCSDNS
  - Default storage size increased to 32M.
- RXAGENT1
  - Default storage size increased to 32M.
- SMTP
  - Default storage size increased to 32M.
- SNALNKA
  - Default storage size increased to 32M.
- SNMPD
  - Default storage size increased to 32M.
- SNMPQE
  - Default storage size increased to 32M.
- SSLSERV
  - Minidisk 203 increased to 2 cylinders.
- TCPIP
  - Added privilege class C.
- TCPMAINT
  - Default storage size increased to 32M.
  - Minidisk 591 increased to 61 cylinders.
  - Minidisk 592 increased to 70 cylinders.
- VMKERB
  - Default storage size increased to 32M.
- XCHANGE
  - Default storage size increased to 32M.
- X25IPI
  - Default storage size increased to 32M.
- 5VMRSC30
  - Minidisk 2B2 increased to 14 cylinders.
  - Minidisk 2B3 increased to 22 cylinders.
  - Minidisk 501 deleted.
  - Minidisk 505 deleted.
  - Some links to minidisks changed because of added or deleted user IDs. See the 5VMRSC30 PLANINFO file on this user ID's A disk.
- 5VMTCP30
  - Privilege class E removed.
  - Default storage size increased to 100M.
  - Defined as a file pool administrator to support installation and service of the LDAP function.
  - Minidisk 2B2 increased to 115 cylinders.
  - Minidisk 2B3 increased to 59 cylinders.
  - Minidisk 491 increased to 61 cylinders.
  - Minidisk 492 increased to 70 cylinders.
  - Minidisk 493 increased to 40 cylinders.
  - Links to MAINT 493 as 493C.
  - Some links to minidisks changed because of added or deleted user IDs. See the 5VMTCP30 PLANINFO file on this user ID's A disk.

[V5.3] Changes to Supplied SYSTEM CONFIG Definitions

TIMEZONE_BOUNDARY statement:

- These statements were updated to cover the years 2006 - 2011.

[V5.4] Installation Changes

The following changes and enhancements have been made to the installation procedure:

- New 540 levels of Directory Maintenance Facility (DIRM), HCD and HCM for z/VM (VMHCD), Language Environment (LE), Performance Toolkit for VM (PERFTK), RACF Security Server for z/VM (RACF), RSCS Networking for z/VM (RSCS), and TCP/IP (TCPIP) are preinstalled.
- New support is provided that allows:
  - Changing the default pack labels for DASD attached to a user at a virtual address.
  - Uploading the installation DVD media to a VM minidisk using a terminal emulator.
- Installation logging has been enhanced to ease problem determination and provide installation status.

[V5.4] Service Enhancements

z/VM V5.4 includes the following service enhancements:

- The LOCALMOD and SERVICE execs have been updated with new operands or options.

[V5.4] Changes to Predefined User Directory Entries

Added User IDs

GSKADMIN
SSL server certificate database manager

Renamed User IDs

- 5VMHCD20 renamed to 5VMHCD**40**.
- 5VMDIR30 renamed to 5VMDIR**40**.
- 5VMPTK30 renamed to 5VMPTK**40**.
- 5VMRAC30 renamed to 5VMRAC**40**.
- 5VMRSC30 renamed to 5VMRSC**40**.
- 5VMTCP30 renamed to 5VMTCP**40**.

Changed User Specifications

- All user IDs (except RACFSMF, RACFVM, RACMAINT, and 5VMRAC40)
  - MACHINE XA statement changed to MACHINE ESA (or MACH XA to MACH ESA).
- MAINT
  - Minidisk 49B increased to 400 cylinders 3390 or 576000 blocks FBA.
- MPROUTE
  - Added the following statement:

    ```
    SHARE RELATIVE 3000
    ```
- OSADMIN1
  - Privilege class changed to BG.
- OSADMIN2

  - Privilege class changed to BG.
- OSADMIN3
  - Privilege class changed to BG.
- RSCSAUTH
  - OPERATOR removed from CONSOLE statement.
- RXAGENT
  - Password changed from RXAGENT to AUTOONLY.
- SSLSERV
  - Default and maximum storage increased to 256M 2G.
  - Privilege class changed from BG to G.
  - APPLMON added to OPTION statement.
  - Added the following statement between the OPTION and LINK statements:

    ```
    SHARE RELATIVE 3000
    ```
  - Added the following statement between the INCLUDE and IUCV statements:

    ```
    POSIXINFO UID 7 GNAME security
    ```
  - MDISK 201 statement removed
  - MDISK 203 statement removed
  - LINK 5VMTCP30 493 493 RR statement removed.
- TCPIP
  - Default and maximum storage increased to 128M 256M.
- VSMPROXY
  - DIAG88 added to OPTION statement.
- VSMREQIU
  - IUCV ANY changed to IUCV ALLOW.
- VSMWORK1
  - DIAG88 removed from OPTION statement.
- VSMWORK2
  - DIAG88 removed from OPTION statement.
- VSMWORK3
  - DIAG88 removed from OPTION statement.
- 5VMHCD40
  - Minidisk 2B2 increased to 110 cylinders 3390 or 158400 blocks FBA.
  - Minidisk 2D2 increased to 200 cylinders 3390 or 288000 blocks FBA.
  - Minidisk 300 increased to 180 cylinders 3390 or 259200 blocks FBA.
  - Minidisk 400 increased to 180 cylinders 3390 or 259200 blocks FBA.
- 5VMTCP40
  - MDISK 493 statement removed.
  - LINK SSLSERV 201 274 MR statement removed.
  - LINK SSLSERV 203 276 MR statement removed.

[V5.4] Backing Off to a Previous z/VM Release

After running z/VM V5.4, if you need to back off z/VM V5.4 and IPL a previous z/VM release, it is suggested that you IPL with the CLEAR option on the HMC LOAD panel (that is, instead of using the NORMAL option on the HMC LOAD panel). This is necessary only in a small set of abnormal scenarios (for example, z/VM V5.4 enters a wait state, or if Standalone Dump is IPLed). If you intend to IPL either z/VM V5.3 or z/VM V5.2, and you have applied the PTF for APAR VM64413 to that system, you do not need to use the CLEAR option on the IPL.

[V6.1] Changes to Predefined User Directory Entries

Added User IDs

Note: These user IDs are not included in the supplied V6.1 source directory. However, they are included in the source directory for V6.2 and later.

DTCENS1
Primary virtual switch controller for intraensemble data network (IEDN) or intranode management network (INMN)

DTCENS2
Backup virtual switch controller for IEDN or INMN

VSMGUARD
z/VM systems management worker server and error recovery server for ensemble support

VSMREQI6
z/VM systems management IPv6 server for ensemble support

VSMREQIM
z/VM systems management network request server for ensemble support

ZVMLXAPP
Management guest (type MG) server for ensemble support

Renamed User IDs

- 5VMDIR40 renamed to **6**VMDIR**10**.
- 5VMPTK40 renamed to **6**VMPTK**10**.
- 5VMRAC40 renamed to **6**VMRAC**10**.
- 5VMRSC40 renamed to **6**VMRSC**10**.
- 5VMTCP40 renamed to **6**VMTCP**10**.

Changed User Specifications

- MAINT
  - MDISK 405 statement removed.
  - Minidisk 3B2 increased to 300 cylinders 3390 or 432000 blocks FBA.
  - Minidisk 400 increased to 130 cylinders 3390 or 187200 blocks FBA.
- ZVMMAPLX
  - Privilege classes B and E removed (leaving only G).
  - Added the following statement:

    ```
    CONSOLE 009 3215
    ```
- All z/VM systems management request servers and worker servers
  - Default and maximum storage increased to 128M 512M.
  - PARM AUTOCR added to the IPL CMS statement.

Note: These changes are not included in the supplied V6.1 source directory. However, they are included in the V6.2 source directory.

[V6.2] Changes to Predefined User Directory Entries

Deleted User IDs

- ADMSERV
- GCSXA
- NAMESRV
- NDBPMGR

- NDBSRV01
- SNALNKA
- SSLSERV
- TFTPD
- VMKERB
- VSMSERVE
- X25IPI

Renamed User IDs

- 5VMHCD40 renamed to **6**VMHCD**20**.
- 6VMDIR10 renamed to 6VMDIR**20**.
- 6VMPTK10 renamed to 6VMPTK**20**.
- 6VMRAC10 renamed to 6VMRAC**20**.
- 6VMRSC10 renamed to 6VMRSC**20**.
- 6VMTCP10 renamed to 6VMTCP**20**.

Changed Virtual Machine Definitions

The virtual machine definitions for the following user IDs have been changed from single-configuration virtual machine definitions (USER definitions) to multiconfiguration virtual machine definitions (IDENTITY definitions). For more information about this new type of virtual machine definition, see "[V6.2] SSI Cluster User Identity and Configuration" on page 201.

AUDITOR	LGLOPR	RACMAINT	VMSERVR
AUTOLOG1	LPSERVE	REXECD	VMSERVS
AUTOLOG2	MAINT	RSCS	VMSERVU
AVSVM	MIGMAINT	RSCSAUTH	VMUTIL
CBDIODSP	MONWRITE	RSCSDNS	VSMGUARD[1]
CMSBATCH	MPROUTE	RXAGENT1	VSMPROXY
DHCPD	OPERATNS	SMTP	VSMREQIM[1]
DISKACNT	OPERATOR	SNMPD	VSMREQIN
DTCENS1[1]	OPERSYMP	SNMPQE	VSMREQIU
DTCENS2[1]	OP1	SNMPSUBA	VSMREQI6[1]
DTCVSW1	OSADMIN1	SYSDUMP1	VSMWORK1
DTCVSW2	OSADMIN2	SYSMON	VSMWORK2
EREP	OSADMIN3	TCPIP	VSMWORK3
FTPSERVE	OSAMAINT	TCPMAINT	XCHANGE
GCS	OSASF	TSAFVM	ZVMLXAPP[1]
GSKADMIN	PERFSVM	UFTD	ZVMMAPLX
IMAP	PORTMAP	VMNFS	
IMAPAUTH	RACFSMF	VMRMADMN	
LDAPSRV	RACFVM	VMRMSVM	

Note on the table:

[1] This user ID was added (defined) in z/VM V6.1 but was not included in the supplied source directory. However, the IDENTITY definition for this user ID is included in the z/VM V6.2 source directory.

Added User IDs

- **USER** definitions:

  DATAMOV2
  DirMaint DATAMOVE service machine for SSI cluster member 2

  DATAMOV3
  DirMaint DATAMOVE service machine for SSI cluster member 3

  DATAMOV4
  DirMaint DATAMOVE service machine for SSI cluster member 4

  DIRMSAT2
  DirMaint DIRMSAT service machine for SSI cluster member 2

  DIRMSAT3
  DirMaint DIRMSAT service machine for SSI cluster member 3

  DIRMSAT4
  DirMaint DIRMSAT service machine for SSI cluster member 4

  MAINT620
  Owns the system resources that are specific to z/VM V6.2, and is used (rather than MAINT) to service all z/VM and preinstalled products, features, and components for z/VM V6.2

  PMAINT
  Owns the system resources that will be common in an SSI cluster, such as the parm disk and the minidisks that contain the source directory, cross release utilities (such as CPFMTXA and DIRECTXA), VMSES/E system-level files, and the new VMPSFS file pool

  VMSERVP
  Server for the new VMPSFS file pool

  6VMLEN20
  Owns the Language Environment disks

- **IDENTITY** definitions:

  DTCSMAPI
  SMAPI default TCP/IP stack

  LOHCOST
  SMAPI default database server

  PERSMAPI
  SMAPI default performance monitoring server

  SSL Base definition for the virtual machines in the SSL server pool

  SSLDCSSM
  SSL DCSS management agent server

  VSMEVSRV
  SMAPI default AF_EVNT server

  ZVMLXTS
  Provides an anchor for the Linux Terminal Server

Added Profiles

TCPSSLU
Defines the characteristics and configuration of the virtual machines in the SSL server pool

Changed User Specifications

- AUTOLOG1
  - SUBCONFIG entries:
    - Added links:

      ```
      LINK MAINT 193 193 RR
      LINK MAINT 19E 19E RR
      ```

    - Increased MDISK 191 from 001 to 005 cylinders.
- AUTOLOG2
  - IDENTITY entry:
    - Increased storage from 4M 4M to 32M 32M.
    - Changed class from BG to ABCDEG.
    - Changed ACCOUNT from SYSTEMS to 9 SYSTEM.
    - Changed IPL CMS PARM AUTOCR to IPL 190.
    - Added AUTOLOG OP1 MAINT.
  - SUBCONFIG entries:
    - Added link:

      ```
      LINK MAINT 0193 0193 RR
      ```

- AVSVM
  - Removed "64" from end of IDENTITY (formerly USER) statement.
- BLDCMS
  - Added link:

    ```
    LINK MAINT 990 990 RR
    ```

- BLDNUC
  - Added link:

    ```
    LINK MAINT 890 990 MW
    ```

- BLDRACF
  - Changed links from:

    ```
    LINK 6VMRAC20 590 490 MW
    LINK 6VMRAC20 505 305 RR
    ```

    to:

    ```
    LINK RACMAINT 490 490 MW
    LINK RACMAINT 305 305 RR
    ```

- BLDSEG
  - Added link:

    ```
    LINK MAINT 19D 19D MR
    ```

- CBDIODSP
  - IDENTITY entry:
    - Deleted link:

      ```
      LINK 6VMHCD20 400 400 RR
      ```

  - SUBCONFIG entries:
    - Added MDISK 400 (moved from 6VMHCD20) and increased from 180 to 360 cylinders.
- DATAMOVE
  - Changed password from NOLOG to AUTOONLY.
  - Changed links from:

    ```
    LINK 6VMDIR20 491 191 RR
    LINK 6VMDIR20 492 192 RR
    LINK 6VMDIR20 11F 11F RR
    LINK 6VMDIR20 41F 21F RR
    ```

    to:

```
LINK DIRMAINT 191 191 RR
LINK DIRMAINT 192 192 RR
LINK DIRMAINT 11F 11F RR
LINK DIRMAINT 21F 21F RR
```

  - Added link:

```
LINK PMAINT 551 551 RR
```

- DHCPD
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- DIRMAINT
  - Changed password from NOLOG to AUTOONLY.
  - Increased storage from 32M 64M to 128M 256M.
  - Changed IUCV ANY to IUCV ALLOW.
  - Added link:

```
LINK PMAINT 551 551 RR
```

  - Increased MDISK 155 from 9 to 12 cylinders.
  - Increased MDISK 1DB from 9 to 12 cylinders.
  - Increased MDISK 1DF from 9 to 12 cylinders.
  - Increased MDISK 1FA from 9 to 12 cylinders.
  - Increased MDISK 2DB from 9 to 12 cylinders.
  - Increased MDISK 2DF from 9 to 12 cylinders.
  - Added MDISK 15D.
- DIRMSAT
  - Changed password from NOLOG to AUTOONLY.
  - Increased storage from 16M 16M to 128M 256M.
  - Changed links from:

```
LINK 6VMDIR20 491 191 RR
LINK 6VMDIR20 492 192 RR
LINK 6VMDIR20 11F 11F RR
LINK 6VMDIR20 41F 21F RR
```

    to:

```
LINK DIRMAINT 191 191 RR
LINK DIRMAINT 192 192 RR
LINK DIRMAINT 11F 11F RR
LINK DIRMAINT 21F 21F RR
```

  - Added link:

```
LINK PMAINT 551 551 RR
```

  - Increased MDISK 1FA from 9 to 12 cylinders.
  - Added MDISK 1DE.
- DTCENS1
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- DTCENS2
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- DTCVSW1
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- DTCVSW2
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- FTPSERVE
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- GSKADMIN
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- IBMUSER
  - Changed links from:

```
LINK 6VMRAC20 29E 29E RR
LINK 6VMRAC20 505 305 RR
LINK 6VMRAC20 191 192 RR
```

    to:

```
LINK RACMAINT 29E 29E RR
LINK RACMAINT 305 505 RR
LINK RACMAINT 192 192 RR
LINK RACFVM   305 305 RR
```

- IMAP
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- IMAPAUTH
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- LDAPSRV
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- LGLOPR
  - SUBCONFIG entries:
    - Deleted link:

```
LINK MAINT 194 194 RR
```

- LPSERVE
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- MAINT
  - IDENTITY entry:
    - Removed MAINT from AUTOLOG statement.
    - Added links:

```
LINK PMAINT 2CC 2CC MR
LINK PMAINT 551 551 RR
```

    - Deleted MDISK 2CC. Minidisk has been split into PMAINT 2CC (10 cylinders) and MAINT620 4CC (10 cylinders).
    - Moved the following minidisks to MAINT620 (and increased some sizes as indicated).

194[1]	3A2	493[6]	5D2	6D2
201	3A4	500[7]	5E5[8]	7A2
2A2	3A6	51D	5E6	7A4
2A4	3B2[3]	5A2	6A2	7A6
2A6	3C2	5A4	6A4	7B2
2C2	3C4	5A6	6A6	7C2
2C4	3D2	5B2	6B2	7C4
2D2[2]	400[4]	5C2	6C2	7D2
319	490[5]	5C4	6C4	CF2

Notes on the table:

[1] Increased from 333 to 360 cylinders.
[2] Increased from 150 to 350 cylinders.
[3] Increased from 300 to 375 cylinders.
[4] Increased from 130 to 275 cylinders.
[5] Increased from 107 to 214 cylinders.
[6] Increased from 167 to 250 cylinders.
[7] Increased from 600 to 900 cylinders.
[8] Increased from 9 to 18 cylinders.

- Moved the following minidisks to 6VMLEN20:
  - 49B
  - 49E
  - 4A2
  - 4A4
  - 4A6
  - 4B2
  - 4C2
  - 4C4
  - 4D2

– SUBCONFIG entries:
- Added links:

```
LINK MAINT620 194 194 RR
LINK MAINT620 201 201 RR
LINK MAINT620 2A2 2A2 RR
LINK MAINT620 2A4 2A4 RR
LINK MAINT620 2A6 2A6 RR
LINK MAINT620 2C4 2C4 RR
LINK MAINT620 2D2 2D2 RR
LINK MAINT620 490 490 RR
LINK MAINT620 493 493 RR
LINK MAINT620 51D 51D RR
LINK MAINT620 5E5 5E5 RR
LINK MAINT620 890 890 RR
LINK 6VMLEN20 49E 49E RR
```

- Increased MDISK 190 from 107 to 214 cylinders.
- Increased MDISK 193 from 167 to 500 cylinders.
- Increased MDISK 19D from 146 to 292 cylinders.
- Increased MDISK 19E from 250 to 500 cylinders.
- Increased MDISK 401 from 146 to 292 cylinders.
- Increased MDISK 402 from 146 to 292 cylinders.
- Added MDISK CFD.

• MIGMAINT
– SUBCONFIG entries:
- Added minidisks:

    - MDISK 2222
    - MDISK 24CC
    - MDISK 2CF0
  - Increased MDISK 191 from 10 to 175 cylinders.
  - Added links:

```
LINK PMAINT 551 551 RR
LINK PMAINT 41D 41D RR
```

- MONWRITE
  - SUBCONFIG entries:
    - Increased MDISK 191 from 090 to 300 cylinders.
- MPROUTE
  - IDENTITY entry:
    - Added IPL ZCMS statement.
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- OPERATNS
  - IDENTITY entry:
    - Increased storage from 32M 32M to 128M 128M.
    - Added OPTION DIAG88 statement.
    - Added NAMESAVE VSMDCSS statement.
  - SUBCONFIG entries:
    - Increased MDISK 191 from 015 to 100 cylinders.
- OP1
  - IDENTITY entry:
    - Removed OP1 from AUTOLOG statement.
- OSADMIN1
  - Changed storage from 16M 32M to 128M 256M.
- OSADMIN2
  - Changed storage from 16M 32M to 128M 256M.
- OSADMIN3
  - Changed storage from 16M 32M to 128M 256M.
- OSASF
  - Changed storage from 64M 64M to 128M 256M.
- PERFSVM
  - IDENTITY entry:
    - Increased storage from 64M 512M to 128M 512M.
    - Changed ACCOUNT from xxxx to XXXXX.
    - Changed IUCV statements from:

```
IUCV *IDENT FCXRES00 GLOBAL
IUCV *IDENT FCXSYSTM GLOBAL
```

      to:

```
IUCV *IDENT FCXC1R01 GLOBAL
IUCV *IDENT FCXC1S01 GLOBAL
IUCV *IDENT FCXC1R02 GLOBAL
IUCV *IDENT FCXC1S02 GLOBAL
```

```
IUCV *IDENT FCXC1R03 GLOBAL
IUCV *IDENT FCXC1S03 GLOBAL
IUCV *IDENT FCXC1R04 GLOBAL
IUCV *IDENT FCXC1S04 GLOBAL
```

  - SUBCONFIG entries:
    - Deleted links:

```
LINK 6VMPTK20 1CC 1CC RR
LINK 6VMPTK20 201 201 RR
```

    - Added MDISK 1CC (moved from 6VMPTK20) and increased from 1 to 2 cylinders.
    - Added MDISK 201 (moved from 6VMPTK20) and increased from 10 to 20 cylinders.

- PORTMAP
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- RACFVM
  - SUBCONFIG entries:
    - Deleted links:

```
LINK 6VMRAC20 590 590 MR
LINK 6VMRAC20 505 505 MR
LINK RACMAINT 191 591 MR
```

    - Increased MDISK 305 from 68 to 136 cylinders.
    - Increased MDISK 490 from 38 to 70 cylinders.

- RACMAINT
  - SUBCONFIG entries:
    - Added links:

```
LINK 6VMRAC20 29E 29E RR
LINK 6VMRAC20 191 192 RR
```

- REXECD
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- RSCS
  - SUBCONFIG entries:
    - Changed link from:

```
LINK  6VMRSC20 401 191 RR
```

      to:

```
LINK RSCS 401 191 RR
```

    - Added MDISK 401 (moved from 6VMRSC20) and increased from 9 to 18 cylinders.
    - Added MDISK 403 (moved from 6VMRSC20) and increased from 3 to 6 cylinders.

- RSCSAUTH

  - SUBCONFIG entries:
    - Changed link from:

```
LINK  6VMRSC20 403 403 RR
```

      to:

```
LINK RSCS  403 403 RR
```

- RSCSDNS
  - SUBCONFIG entries:
    - Changed link from:

```
LINK  6VMRSC20 403 403 RR
```

      to:

```
LINK RSCS  403 403 RR
```

- RXAGENT1
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- SMTP
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- SNMPD
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- SNMPQE
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- SNMPSUBA
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- SYSADMIN
  - Changed links from:

```
LINK 6VMRAC20 29E 29E RR
LINK 6VMRAC20 505 505 RR
LINK 6VMRAC20 191 192 RR
```

    to:

```
LINK RACMAINT 29E 29E RR
LINK RACMAINT 305 505 RR
LINK RACMAINT 192 192 RR
```

- TCPIP
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- TCPMAINT
  - SUBCONFIG entries:
    - Added links:

```
LINK 6VMTCP20 491 491 RR
LINK 6VMTCP20 492 492 RR
```

    - Increased MDISK 591 from 61 to 122 cylinders.
    - Increased MDISK 592 from 70 to 140 cylinders.
- UFTD
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- VMNFS
  - SUBCONFIG entries:
    - Changed links from:

```
LINK  6VMTCP20 491 491 RR
LINK  6VMTCP20 492 492 RR
```

      to:

```
LINK  TCPMAINT 491 491 RR
LINK  TCPMAINT 492 492 RR
```

- VMSERVR
  - IDENTITY entry:
    - Changed IPL CMS to IPL 190.
- VMSERVS
  - IDENTITY entry:
    - Changed IPL CMS to IPL 190.
  - SUBCONFIG entries:
    - Deleted MDISKs 306, 307, 308, 309, 310, and 311.

- VMSERVU
  - IDENTITY entry:
    - Changed IPL CMS to IPL 190.
- VSMGUARD
  - SUBCONFIG entries:
    - Changed links from:

      ```
      LINK  MAINT  CF1 CF1 MD
      LINK  MAINT  CF2 CF2 MD
      ```

      to:

      ```
      LINK  PMAINT CF0 CF0 MD
      ```
    - Added link:

      ```
      LINK PMAINT  551 551 RR
      ```
- VSMPROXY
  - IDENTITY entry:
    - Increased storage from 32M 32M to 128M 512M.
    - Changed IPL CMS to IPL CMS PARM AUTOCR.
    - Removed IUCV *VMEVENT statement.
    - Added NAMESAVE VSMDCSS statement.
- VSMREQIM
  - IDENTITY entry:
    - Removed IUCV *VMEVENT statement.
- VSMREQIN
  - IDENTITY entry:
    - Increased storage from 32M 32M to 128M 512M.
    - Changed IPL CMS to IPL CMS PARM AUTOCR.
    - Added NAMESAVE VSMDCSS statement.
- VSMREQIU
  - IDENTITY entry:
    - Increased storage from 32M 32M to 128M 512M.
    - Changed IPL CMS to IPL CMS PARM AUTOCR.
    - Added NAMESAVE VSMDCSS statement.
- VSMREQI6
  - IDENTITY entry:
    - Removed IUCV *VMEVENT statement.
- VSMWORK1
  - IDENTITY entry:
    - Increased storage from 64M 64M to 128M 512M.
    - Changed IPL CMS to IPL CMS PARM AUTOCR.
    - Added DIAG88 to OPTION statement.
    - Added NAMESAVE VSMDCSS statement.
  - SUBCONFIG entries:
    - Changed links from:

      ```
      LINK  MAINT  CF1 CF1 MD
      LINK  MAINT  CF2 CF2 MD
      ```

      to:

      ```
      LINK  PMAINT CF0 CF0 MD
      ```
    - Added link:

      ```
      LINK PMAINT  551 551 RR
      ```
- VSMWORK2
  - IDENTITY entry:
    - Increased storage from 64M 64M to 128M 512M.

    - Changed IPL CMS to IPL CMS PARM AUTOCR.
    - Added DIAG88 to OPTION statement.
    - Added NAMESAVE VSMDCSS statement.
  – SUBCONFIG entries:
    - Changed links from:

```
LINK  MAINT  CF1 CF1 MD
LINK  MAINT  CF2 CF2 MD
```

      to:

```
LINK  PMAINT CF0 CF0 MD
```

    - Added link:

```
LINK PMAINT  551 551 RR
```

- VSMWORK3
  – IDENTITY entry:
    - Increased storage from 64M 64M to 128M 512M.
    - Changed IPL CMS to IPL CMS PARM AUTOCR.
    - Added DIAG88 to OPTION statement.
    - Added NAMESAVE VSMDCSS statement.
  – SUBCONFIG entries:
    - Changed links from:

```
LINK  MAINT  CF1 CF1 MD
LINK  MAINT  CF2 CF2 MD
```

      to:

```
LINK  PMAINT CF0 CF0 MD
```

    - Added link:

```
LINK PMAINT  551 551 RR
```

- XCHANGE
  – SUBCONFIG entries:
    - Changed link from:

```
LINK  6VMRSC20 403 403 RR
```

      to:

```
LINK RSCS  403 403 RR
```

    - Added MDISK 191.
- ZVMLXAPP
  – IDENTITY entry:
    - Changed password from AUTOONLY to ZVMLXAPP.
    - Changed LOADDEV statement from:

```
LOADDEV PORT 0 LUN 0 BOOT 0
```

      to:

```
LOADDEV PORT 0
LOADDEV LUN 0
LOADDEV BOOT 0
```

    - Removed statement:

```
SPOOL 000E 1403 A
```

- 40SASF40
  – Changed link from:

```
LINK 6VMTCP20 0492 0492 RR
```

    to:

```
LINK TCPMAINT 0492 0492 RR
```

- 6VMDIR20
  - Added link:

    ```
    LINK PMAINT 551 551 RR
    ```
  - Deleted MDISK 502.
  - Increased MDISK 41F from 8 to 16 cylinders.
  - Increased MDISK 491 from 15 to 30 cylinders.
- 6VMHCD20
  - Added link:

    ```
    LINK CBDIODSP 400 400 MR
    ```
  - Moved MDISK 400 to CBDIODSP.
- 6VMPTK20
  - Added links:

    ```
    LINK PERFSVM 1CC 1CC MR
    LINK PERFSVM 201 201 MR
    ```
  - Moved MDISK 1CC and MDISK 201 to PERFSVM.
- 6VMRAC20
  - Increased MDISK 29E from 2 to 10 cylinders.
  - Increased MDISK 590 from 38 to 70 cylinders.
- 6VMRSC20
  - Added links:

    ```
    LINK RSCS 401 401 MR
    LINK RSCS 403 403 MR
    ```
  - Deleted MDISK 502.
  - Moved MDISK 401 and MDISK 403 to RSCS.
  - Added MDISK 493.
- 6VMTCP20
  - Deleted links:

    ```
    LINK ADMSERV 191 263 MR
    LINK NAMESRV 191 269 MR
    LINK NDBPMGR 191 26A MR
    LINK NDBSRV01 191 26B MR
    LINK SNALNKA 191 270 MR
    LINK SSLSERV 191 273 MR
    LINK TFTPD 191 279 MR
    LINK VMKERB 191 27B MR
    LINK X25IPI 191 27D MR
    ```
  - Added links:

    ```
    LINK SSLDCSSM 191 263 MR
    LINK DTCENS1 191 269 MR
    LINK DTCENS2 191 26A MR
    ```
  - Increased MDISK 2B2 from 115 to 150 cylinders.
  - Increased MDISK 2D2 from 117 to 200 cylinders.

[V6.2] INSTALL Utility Renamed to INSTTAPE

The CP utility INSTALL has been renamed to INSTTAPE. The TAPE operand has been removed, but the RECOVER option is still provided.

The INSTALL EXEC used during the z/VM V6.2 installation process is a new function that is used only during z/VM installation. It does not invoke the old INSTALL utility (now INSTTAPE) and it has no operands or options.

[V6.2] CMS AUDITOR Utility Change

After you complete the customization of the CMS AUDITOR, you must uncomment the call to the AUDITOR EXEC in the AUDITOR user ID's PROFILE EXEC in order to start the AUDITOR utility.

[V6.2] High Level Assembler Not Required for RACF

The IBM High Level Assembler is no longer required to reassemble CP replacement parts for the RACF Security Server for z/VM.

[V6.2] Options Removed from the VMFINS Command

The RESOURCE, DFNAME, DFTYPE, and DFMODE options have been removed from the VMFINS DELETE, VMFINS INSTALL, and VMFINS MIGRATE commands. The option NORESOURCE remains on these VMFINS commands and is the default. Also, the VMFINS DEFAULTS file has been updated to remove the options NORESOURCE, DFNAME, DFTYPE, and DFMODE.

[V6.2] Change to the Linking Default of the VMFSETUP Command

The linking default of the VMFSETUP command has been changed from NOLINK to LINK.

[V6.2] Service Changes to Support the SSI Environment

z/VM V6.2 includes the following service changes to support the SSI environment:

- The MAINT*vrm* user ID, where *vrm* is the version, release, and modification level of the z/VM system, must now be used rather than the MAINT user ID to service all z/VM and preinstalled products, features, and components.
- When a product is placed into production, files that you placed on a test disk are not automatically moved to a production disk. If you placed files on a test disk, you must manually copy those files to the production disk.
- When a product is placed into production, the DDR of MAINT 490 to MAINT 190 will now DDR only the RECOMP area where the CMS nucleus resides. Because of this, the 490 disk and the 190 disk must be the same size, and the size of the RECOMP area on the 490 disk must be the same as the size of the RECOMP area on the 190 disk.
- It is recommended that you do not save backup copies of the CPLOAD MODULE on the MAINT CF3 parm disk. This disk will be erased when a product is placed into production. Instead, you should save backup copies of the CPLOAD MODULE on the MAINT CF1 or the MAINT*vrm* CF2 disk.
- VMPSFS is a new IBM supplied file pool that is the z/VM product service file pool for products loaded into SFS.
- The operation of key VMSES/E commands (SERVICE, PUT2PROD, and VMFSUFIN) is now based on using the content of the VMSES/E 5E6 Test Build minidisk, instead of the 5E5 Build disk. Because of this change, the VMSES/E 5E6 minidisk must be accessed (by convention, at file mode B) before these commands are used for service and maintenance activity. An ACCESS command for this purpose is included in the PROFILE EXEC for the MAINT*vrm* user ID that is supplied with the system. This change should be accounted for in any local changes or customization that pertain to the MAINT*vrm* user ID and its PROFILE EXEC.

[V6.2] Change to PPF File Used for Saved Segment Builds

The CMS and Language Environment saved segments supplied with z/VM are now built using the SERVP2P PPF file instead of the ZVM PPF file.

[V6.3] Changes to Predefined User Directory Entries

Deleted User IDs

- DHCPD (IDENTITY)
- LPSERVE (IDENTITY)

Added User IDs

- **IDENTITY** definitions:

  XCAT Supports the xCAT appliance packaged with z/VM.

  ZHCP Supports the xCAT appliance packaged with z/VM.

Renamed User IDs

- 6VMDIR20 changed to 6VMDIR30.
- 6VMPTK20 changed to 6VMPTK30.
- 6VMRAC20 changed to 6VMRAC30.
- 6VMRSC20 changed to 6VMRSC30.
- 6VMTCP20 changed to 6VMTCP30.
- MAINT620 changed to MAINT630.

Changed User Specifications

- All logon passwords that previously matched the user ID have been changed to WD5JU8QP.
- DIRMAINT
  - Added link:

    ```
    LINK TCPMAINT 592 592 RR
    ```
- GSKADMIN
  - IDENTITY entry:
    - Storage changed to 64M 128M.
- LOHCOST
  - IDENTITY entry:
    - Deleted IPL and LOADDEV statements.
  - SUBCONFIG entries:
    - Added statements:

      ```
      IPL 190 PARM AUTOCR
      LINK MAINT 0193 0193 RR
      LINK MAINT 0400 0400 RR
      ```
- MAINT
  - IDENTITY entry:
    - Storage changed to 256M 1000M.
  - SUBCONFIG entries:
    - Added statements:

      ```
      LINK MAINT630 400 400 RR
      MDISK 125 3390 000 END M01W02 MR
      MDISK 126 3390 000 END M01W03 MR
      ```
- MAINT630
  - Added links (if needed):

```
LINK MAINT 125 125 MR
LINK MAINT 126 126 MR
LINK PMAINT 141 141 MR
LINK PMAINT 142 142 MR
```

  - Increased minidisk sizes (if used):
    - MDISK 194 increased to 400 cylinders.
    - MDISK 3B2 increased to 580 cylinders.
    - MDISK 3D2 increased to 500 cylinders.
    - MDISK 400 increased to 450 cylinders.
    - MDISK 500 increased to 1200 cylinders.
    - MDISK 5B2 increased to 075 cylinders.
- MIGMAINT
  - IDENTITY entry:
    - Storage changed to 256M 1000M.
    - Added OPTION MAINTCCW statement.
- PMAINT
  - Privilege classes changed to G only.
- SYSADMIN
  - Added link:

```
LINK PMAINT 551 551 RR
```

- VMSERVP
  - MDISK 304 increased to 160 cylinders.
  - MDISK 305 increased to 400 cylinders.
- VSMGUARD
  - IDENTITY entry:
    - Logon password changed to AUTOONLY.
    - Added statements:

```
OPTION LNKNOPAS
NAMESAVE SMAPIOUT
```

  - SUBCONFIG entries:
    - Added minidisk:

```
MDISK A91 3390 0303 005 M01W01 MR ALL ALL ALL
```

- VSMWORK1
  - IDENTITY entry:
    - Logon password changed to AUTOONLY.
    - Added statements:

```
OPTION LNKNOPAS
NAMESAVE SMAPIOUT
```

  - SUBCONFIG entries:
    - Added minidisk:

```
MDISK A91 3390 3306 005 M01RES MR ALL ALL ALL
```

- VSMWORK2
  - IDENTITY entry:
    - Logon password changed to AUTOONLY.
    - Added statements:

```
OPTION LNKNOPAS
NAMESAVE SMAPIOUT
```

  - SUBCONFIG entries:
    - Added minidisk:

      ```
      MDISK A91 3390 0168 005 M01W01 MR ALL ALL ALL
      ```

- VSMWORK3
  - IDENTITY entry:
    - Logon password changed to AUTOONLY.
    - Added statements:

      ```
      OPTION LNKNOPAS
      NAMESAVE SMAPIOUT
      ```

  - SUBCONFIG entries:
    - Added minidisk:

      ```
      MDISK A91 3390 0198 005 M01W01 MR ALL ALL ALL
      ```

- 4OSASF40
  - MDISK 2D2 increased to 180 cylinders.
- 6VMDIR30
  - Links changed from:

    ```
    LINK MAINT 51D 51D MR * VMSES/E software inventory disk
    LINK MAINT 5E5 5E5 RR * VMSES/E code
    ```

    to:

    ```
    LINK MAINT630 51D 51D MR * VMSES/E software inventory disk
    LINK MAINT630 5E5 5E5 RR * VMSES/E code
    ```

- 6VMPTK30
  - MDISK 29D (if used) increased to 016 cylinders.
  - MDISK 2B2 (if used) increased to 012 cylinders.
- 6VMRAC30
  - Added minidisk:

    ```
    MDISK 651 3390 2487 001 630RL1 MR READ WRITE MULTIPLE
    ```

- 6VMTCP30
  - Removed links:

    ```
    LINK DHCPD 191 265 MR
    LINK LPSERVE 191 267 MR
    ```

[V6.3] ACIGROUP Keyword of RACF GLBLDSK Macro Changed to ACIGRP

The ACIGROUP keyword on the RACF GLBLDSK macro has been changed to ACIGRP to allow the use of Assembler XF. If you use the ACIGROUP keyword on the GLBLDSK macro you will need to change the keyword to ACIGRP. The GLBLDSK macro can be used in the HCPRWA and HCPRWAC CP parts.

[V6.3] Installation Changes

The following changes and enhancements have been made to the installation procedure:

- Starting with z/VM V6.3, a new installation technique called upgrade installation is introduced. In a traditional installation, the new release is installed on a separate set of volumes, after which the users and data from the current running system are migrated to the system running the new release. In an upgrade installation, a new release system is installed as a second-level guest of the current release system that you wish to upgrade. The new level of code from

the guest system is then moved to your current system with minimal impact to your current running system. This current running system can be a non-SSI system, the only member of a single-member SSI cluster, or any member of a multimember SSI cluster. In a multimember SSI cluster, you will upgrade one member at a time so that there is minimum impact to the other members.

- An upgrade installation is performed using a two-stage approach, with two separate sets of changes being defined and then made on you current system. The first set of changes, STAGE1, can be made to your current system without disrupting your normal system operation. After the STAGE1 changes are made, you must back up your current system before moving on to the disruptive STAGE2 changes. To perform the STAGE2 changes, you must stop all normal production work on your current system. In an SSI cluster environment, you can relocate production Linux workloads from your current system to other members of your cluster before performing STAGE2 activities.
- It is important to note that this procedure can only be used to upgrade from z/VM V6.2. If you wish to migrate from z/VM V5.4 or V6.1, you must use the traditional installation method to install a z/VM V6.3 system and then follow the migration procedures as documented both in this book and in the *z/VM: Installation Guide*, as necessary.
- Installation messages are contained in the new IUGMES repository rather than the HCPMES repository. This means that as of z/VM V6.3, both existing and new installation messages have the IUG prefix.

[V6.3] Relocation of Service-Level Production Status Table

The service-level production status table (*prodid* SRVPROD) has been moved from the system inventory disk (MAINT*vrm* 51D) to the production inventory disk (PMAINT 41D).

[V6.3] Service Enhancements

z/VM V6.3 includes the following service enhancements:

- The VMVFIEW EXEC has been updated to accommodate VMSES/E program-specific tracing output.
- Console data excluded from the SERVICE $CONS file by the SERVICE EXEC is spooled to the reader as a file called SERVICE $CONXCLD.

[V6.3] SSI Cluster Cross-System Highest Release Level Program Handling

In an SSI cluster in which members have different levels of z/VM installed, certain programs that manage shared resources, such as DIRECTXA and DISKMAP, are required to be at the highest release level that is running in the cluster. These programs must be on all members in the cluster regardless of the release level running on each member. Handling the highest release level program is accomplished by having a common test/production build disk set, where programs that must be at the highest release level are shipped and maintained. These disks are on the common volume and are owned by the PMAINT user ID. When a z/VM release which supersedes all other releases running on the members in an SSI cluster is installed on a member of the cluster, the z/VM installation processing places these programs from the superseding release on the SSI system common disk, replacing all programs from the superseded release.

[V6.3] ZVM $PPF Override File Removed

The override product parameter file (PPF), ZVM $PPF, which contained overrides to the base $PPF file for each component, has been removed. SERVP2P $PPF, which was introduced in z/VM V4.1, is the IBM-supplied override product parameter file which contains overrides to the base $PPF file for each component.

[V6.3] New LINKRr Option for the VMFSETUP Command

A new option for the VMFSETUP command, LINKRr, resolves any links specified in the :DCL section of the PPF using the CP LINK command, and acquires these disks with read-only status.

[V6.4] Changes to Predefined User Directory Entries

Deleted User IDs

- DTCENS1 (IDENTITY)
- DTCENS2 (IDENTITY)
- VSMREQIM (IDENTITY)
- VSMPROXY (IDENTITY)
- XCAT (IDENTITY)
- ZHCP (IDENTITY)
- ZVMLXAPP (IDENTITY)

Added User IDs

- IDENTITY definitions:

  DTCVSW3
  Additional TCP/IP stack for a virtual switch controller.

  DTCVSW4
  Additional TCP/IP stack for a virtual switch controller.

  OPNCLOUD
  Cloud Manager Appliance (CMA) virtual machine, in which the xCAT Management Node (MN) server and the ZHCP server run.

Replaced (Renamed) User IDs

- 6VMDIR30 changed to 6VMDIR40.
- 6VMHCD20 changed to 6VMHCD40
- 6VMPTK30 changed to 6VMPTK40.
- 6VMRAC30 changed to 6VMRAC40.
- 6VMRSC30 changed to 6VMRSC40.
- 6VMTCP30 changed to 6VMTCP40.
- MAINT630 changed to MAINT640.

Changed User Specifications

Note: Links to renamed user IDs have been updated to specify the new names.

- BLDSEG
  - Added statement:
    ```
    LINK MAINT640 5E6 5E6 RR
    ```
- CBDIODSP
  - Storage changed to `2G 2G`

- DATAMOVE
  - MDISK 1AA increased to 20 cylinders (28800 blocks).
  - MDISK 2AA increased to 20 cylinders (28800 blocks).
- DIRMAINT
  - MDISK 1AA increased to 20 cylinders (28800 blocks).
  - MDISK 2AA increased to 20 cylinders (28800 blocks).
- DIRMSAT
  - MDISK 1AA increased to 20 cylinders (28800 blocks).
  - MDISK 2AA increased to 20 cylinders (28800 blocks).
- DTCVSW1
  - IDENTITY entry:
    - Storage changed to `128M 256M`.
    - Added `SHARE RELATIVE 3000` statement.
- DTCVSW2
  - IDENTITY entry:
    - Storage changed to `128M 256M`.
    - Added `SHARE RELATIVE 3000` statement.
- IBMUSER
  - Storage changed to `128M 2G`.
- LOHCOST
  - IDENTITY entry:
    - Deleted `OPTION DEVINFO DEVMAINT LNKNOPAS DIAG88` statement.
- MAINT
  - SUBCONFIG entries:
    - Added statement:

      ```
      LINK MAINT640 5E6 5E6 RR
      ```
- MAINT640
  - MDISK 3B2 decreased to 620 cylinders (892800 blocks).
- MPROUTE
  - IDENTITY entry:
    - Storage changed to `64M 256M`
- OPERATNS
  - IDENTITY entry:
    - Storage changed to `256M 512M`.
- OPNCLOUD
  - IDENTITY entry:
    - Added statement:

      ```
      OPTION LNKNOPAS
      ```
- PERFSVM
  - IDENTITY entry:
    - Storage changed to `512M 3G`.
    - Changed `IPL CMS PARM AUTOCR` to `IPL ZCMS PARM AUTOCR`.
  - SUBCONFIG entries:
    - Added statement:

      ```
      COMMAND DEFINE STORAGE CONFIG 0.512M 2G.2G
      ```

    - MDISK 201 increased to 40 cylinders (57600 blocks).
- PERSMAPI
  - IDENTITY entry:
    - Changed `IPL CMS PARM AUTOCR` to `IPL ZCMS PARM AUTOCR`.
- RACFSMF
  - IDENTITY entry:
    - Storage changed to `128M 2G`.
  - SUBCONFIG entries:
    - Added statements:

      ```
      LINK RACFVM 301 301 MR
      LINK RACFVM 302 302 MR
      ```
- TCPMAINT
  - SUBCONFIG entries:
    - MDISK 591 increased to 160 cylinders (230400 blocks).
    - MDISK 592 increased to 240 cylinders (345600 blocks).
- VMSERVP
  - MDISK 305 increased to 500 cylinders (720000 blocks).
  - MDISK 306 increased to 500 cylinders (720000 blocks).
  - MDISK 307 increased to 500 cylinders (720000 blocks).
  - MDISK 308 increased to 500 cylinders (720000 blocks).
  - MDISK 309 increased to 500 cylinders (720000 blocks).
  - MDISK 310 increased to 500 cylinders (720000 blocks).
  - MDISK 311 increased to 500 cylinders (720000 blocks).
- VSMGUARD
  - IDENTITY entry:
    - Deleted `OPTION MAINTCCW DIAG88` statement.
  - SUBCONFIG entries:
    - Removed write and multiple-write passwords for MDISK A91.
- VSMWORK1
  - IDENTITY entry:
    - Deleted `OPTION MAINTCCW` statement.
  - SUBCONFIG entries:
    - Removed write and multiple-write passwords for MDISK A91.
- VSMWORK2
  - IDENTITY entry:
    - Deleted `OPTION MAINTCCW` statement.
  - SUBCONFIG entries:
    - Removed write and multiple-write passwords for MDISK A91.
- VSMWORK3
  - IDENTITY entry:
    - Deleted `OPTION MAINTCCW` statement.
  - SUBCONFIG entries:
    - Removed write and multiple-write passwords for MDISK A91.
- ZVMMAPLX
  - IDENTITY entry:

    - Changed password from `MAINT` to `WD5JU8QP`.
- 5684042J
  - MDISK 2D2 increased to 40 cylinders (57600 blocks).
- 6VMDIR40
  - MDISK 11F increased to 16 cylinders (23040 blocks)
  - MDISK 41F decreased to 8 cylinders (11520 blocks)
- 6VMHCD40
  - MDISK 2B2 increased to 200 cylinders (288000 blocks).
  - MDISK 2D2 increased to 750 cylinders (1080000 blocks).
  - MDISK 300 increased to 360 cylinders (518400 blocks).
- 6VMPTK40
  - MDISK 200 increased to 20 cylinders (28800 blocks).
- 6VMTCP40
  - Added links:

    ```
    LINK DTCVSW3 191 269 MR
    LINK DTCVSW4 191 26A MR
    ```

  - Deleted links:

    ```
    LINK DTCENS1 191 269 MR
    LINK DTCENS2 191 26A MR
    ```

  - MDISK 491 increased to 80 cylinders (115200 blocks).
  - MDISK 492 increased to 120 cylinders (172800 blocks).
  - MDISK 2D2 increased to 500 cylinders (720000 blocks).

[V6.4] 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.

The following CP monitor record has been added:
- Domain 1 Record 31 - MRMTRSRV - Service Configuration Sample Record

For more information, see *z/VM: CP Commands and Utilities Reference*.

[V6.4] Installation Changes

- Support is added for installation on the 3390 Model 27.
- If z/VM is installed to full pack minidisks, the minimum size is 3339 cylinders.

For more information, see:
- *z/VM: Installation Guide*.

[V6.4] 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 upgrades 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.

For more information, see:

- *z/VM: Installation Guide*.

[V6.4] 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 through via electronic delivery. In addition, z/VM service for all releases will no longer be orderable in 3590 or 3592 tape format.

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

[V6.4] Improved upgrade installation planning information

The instructions for planning an upgrade installation and completing the installation worksheets have been improved. For more information, see *z/VM: Installation Guide*.

[V6.4] Product Parameter File Names

Table 3 lists the base $PPF file for each z/VM component, facility, and optional feature and indicates whether the file name has changed for z/VM V6.4.

Table 3. Base $PPF File Names

Component, Facility, or Feature	File ID	Changed?
AVS	6VMAVS40 $PPF	Yes
CMS	6VMCMS40 $PPF	Yes
CP	6VMCPR40 $PPF	Yes
DirMaint	6VMDIR40 $PPF	Yes
Dump Viewing Facility	6VMDVF40 $PPF	Yes
GCS	6VMGCS40 $PPF	Yes
HCD/HCM	6VMHCD40 $PPF	Yes
Language Environment	6VMLEN20 $PPF	No
OSA/SF	4OSASF40 $PPF	No
Performance Toolkit	6VMPTK40 $PPF	Yes
RACF	6VMRAC40 $PPF	Yes
REXX	6VMREX40 $PPF	Yes
RSCS	6VMRSC40 $PPF	Yes
TCP/IP	6VMTCP40 $PPF	Yes
TSAF	6VMTSA40 $PPF	Yes
VMSES/E	6VMSES40 $PPF	Yes

Support and Exploitation of Hardware and Architectures

These topics describe z/VM support for, or exploitation of, new or enhanced hardware devices, facilities, and architectures.

Notes:

1. Some hardware z/VM support or exploitation might depend on hardware level or availability. See the sections on server support and device support in *z/VM: General Information*. For detailed information on hardware capabilities and requirements, refer to the appropriate hardware announcements and other hardware documentation.

2. This section also describes the z/VM simulation or emulation of certain hardware architectures and facilities. Other simulations and emulations might be described under "Connectivity and Networking" on page 125.

[V2.2] S/390 Open Systems Adapter Support Facility

The IBM S/390® Open Systems Adapter (OSA) is an integrated hardware feature that allows the S/390 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 VM that allows you to customize an OSA's modes of operation. You can access OSA/SF by a CMS user ID, by a REXX call to the OSA/SF application programming interface (API), or through the OSA/SF Windows or OS/2 graphical user interface (GUI).

[V2.2] Multi-Path Lock Facility

The IBM DASD Subsystem Multi-Path Lock Facility (MPLF) provides function to control processes and share data in a loosely coupled environment. VM now provides support to allow dedicated devices and full-pack minidisks to use this real hardware facility. The addition of this support allows Transaction Processing Facility (TPF) systems running as guests on VM to share data with native TPF systems.

MPLF controls locking through a set of channel commands implemented by the IBM DASD Subsystem. These commands result in a setting which indicates the requested operation is either compatible or incompatible with the current state of the lock. The control unit maintains the names and status of the locks in use and responds to requests to obtain or release a lock. The control unit also notifies a host when it permits lock ownership that was previously denied. A host can obtain a lock, release a lock, examine the status of active locks, and check the outcome of lock-related operations using the channel commands.

For more information, see the SET LKFACR and QUERY LKFACR commands in *z/VM: CP Commands and Utilities Reference*.

[V2.3] Guest Coupling Simulation

VM guest coupling simulation provides for the simulation of one or more complete Parallel Sysplex® systems within a single VM system image. This environment allows the testing and debugging of guest operating systems while running under VM. Guest coupling simulation is supported on the following processors (at the appropriate engineering-change levels):

- IBM S/390 Parallel Enterprise Server - Generation 3 and later
- IBM S/390 Multiprise 2000 and later

VM guest coupling simulation support simulates a sysplex environment by using software equivalents of the real hardware and software requirements. The hardware is simulated by the message facility support in CP. Other virtual machines, referred to as coupled guests, are set up to run the guest operating systems (that is, MVS™ or OS/390®). A special virtual machine called a Coupling Facilities (CF) service machine is defined and runs the coupling facility code. The coupled guest machines and the CF service machines are coupled together by a special message facility environment that passes information back and forth between the CF service machines and the coupled guests. The CF service machines manage data movement, scheduling, and locks, and maintain the status of the entire sysplex environment.

This facility also provides a CP command set to allow you to perform operations required to control the coupled guest environment:
- DEFINE MSGPROC
- DETACH MSGPROC
- QUERY VIRTUAL MSGDEVICES
- QUERY VIRTUAL MSGPROC
- RESTART MSGPROC
- SET MSGFACIL

These commands allow you to add and remove links to the CF service machines, request status from the CF service machines, and control the message facility environment. Also, a new CP SET VTOD command has been added to allow Year 2000 testing of a Parallel Sysplex environment on VM.

For more information, see:
- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: Running Guest Operating Systems*
- *z/VM: Diagnosis Guide*

[V2.4] Integrated Cluster Bus Channels

VM provides support for integrated cluster bus channels for S/390 CMOS processors that support this channel type. The existing dynamic I/O configuration support has been updated to allow VM, when running in an LPAR controlling the dynamic I/O configuration changes, to define integrated cluster bus channels for an OS/390 LPAR on the same Central Electronics Complex (CEC).

[V2.4] Fibre Connection Channels

VM provides support for fibre connection (FICON®) channels for processors that support this channel type. FICON channel technology can improve system performance and total aggregate system bandwidth. Each FICON channel provides the equivalent of eight ESCON channels. VM support includes:
- Dynamic I/O configuration support for the new FICON and FICON-converter channel types
- Support for the fibre-channel-to-ESCON converter function on the 9032-5 switch
- Support for the FICON 9042-1 switch
- Support for the S/390 architecture changes in the SCHIB, ORB, and IRB
- A new CP monitor record, Extended Channel Path Measurement Data

[V2.4] QDIO Facility and OSA-Express Gigabit Ethernet

VM provides guest support for the Queued Direct I/O (QDIO) facility on processors that support this new I/O architecture. The QDIO Facility allows a program to directly exchange data with an I/O device without performing traditional S/390 I/O instructions. To exchange data, both the I/O device and the program reference main storage directly through a set of data queues.

VM provides guest (dedicated device) support and dynamic I/O support for the new OSA-Express Gigabit Ethernet (OSA-Express GbE). A new channel path ID has been defined for this device.

The following CP commands have been updated:
- DEFINE CHPID / PATH

- QUERY VIRTUAL OSA
- SET CPTRACE

For more information, see *z/VM: CP Commands and Utilities Reference*.

[V2.4] Cryptographic Support

This support provides guest virtual machine access to the S/390 CMOS Cryptographic Coprocessor by extending the existing VM guest cryptographic support for Bipolar processors. The new cryptographic support is upwardly compatible with the existing Bipolar support.

This support is intended primarily for use by OS/390 Integrated Cryptographic Service Facility (ICSF) applications running in an OS/390 guest of VM. ISCF is currently the only IBM application program interface to the S/390 CMOS cryptographic hardware. The following commands and statements were updated:

- CP DEFINE CRYPTO command
- CP QUERY CRYPTO command
- CP QUERY VIRTUAL CRYPTO command
- CP SET CRYPTO command
- CPU directory control statement
- CRYPTO directory control statement

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Planning and Administration*

[V2.4] IEEE Floating Point

New function has been added to support IEEE Floating Point hardware on servers that provide this capability. CP has been updated to allow multiple levels of guest operating systems to use basic floating point extensions, floating point support extensions, hexadecimal floating point extensions, and binary floating point.

This support includes preservation and restoration of 16 floating point registers (Additional Floating Point registers 1,3,5,7,8-15 plus existing floating point registers 0,2,4,6) and the Floating Point Control (FPC) register which is provided by the IEEE Floating Point hardware. The following were updated:

- CP DISPLAY Registers command
- CP STORE (Registers) command
- CP STORE STATUS command
- CP TRACE command
- Messages HCP6153E and HCP6154E (new)

For more information, see *z/VM: CP Commands and Utilities Reference*.

Note: Applications that exploit the IEEE Floating Point hardware require the IBM High Level Assembler Version 1 Release 3.0.

[V2.4] Extended-TOD-Clock

The Extended-TOD-clock facility is a hardware facility available on certain processors which provides a 128-bit Time of Day (TOD) clock. VM supports the use of the Extended-TOD-clock facility from XA, ESA, and XC virtual machines.

[V3.1] Enterprise Storage Server FlashCopy

z/VM allows a native CP user to initiate a FlashCopy® function (an instant copy of a disk or data set) of a source device to a target device on an IBM Enterprise Storage Server® (ESS). Customers will find this feature especially beneficial for large databases, which normally require a long time to copy. FlashCopy support includes the new CP FLASHCOPY command. For more information, see *z/VM: CP Commands and Utilities Reference*.

[V3.1] ESS Parallel Access Volumes

z/VM provides guest support for the ESS Parallel Access Volumes feature. This feature allows the configuration of logical volumes (known as alias Parallel Access Volumes), where each logical volume (alias) has a unique device address but is actually an exposure of the existing real device (known as the base Parallel Access Volume). This allows the host to issue concurrent I/O requests to one real device, the base volume, through the different alias volumes.

Support for Parallel Access Volumes includes:

- The new CP QUERY PAV command, which displays information about the Parallel Access Volume devices on the system.
- Enhancements to the CP QUERY DASD DETAILS command to display additional information if the queried device is a Parallel Access Volume.
- A new CP Monitor Record, which has been added to Domain 6 (I/O) to record state change interrupts that indicate a change in the Parallel Access Volumes information:
  - Record 20 – MRIODSTC – State change

  Other Parallel Access Volumes information will be recorded in the existing Device Configuration Data Record (Domain 1, Record 6) and the Vary On Device - Event Data Record (Domain 6, Record 1).

[V3.1] Tape Support Enhancements

z/VM provides additional guest support for the IBM 3494 Virtual Tape Server (VTS):

- Peer-to-Peer VTS provides flexible component placement to meet 7x24 processing requirements and provides a solution for remote backup and recovery.
- Import/Export allows physical tapes to be removed from and inserted into the VTS logical tape library.

z/VM supports guest use of IBM 3590 A60 tape controllers attached with FICON channels when such guests themselves support the 3590 A60 on native FICON channels.

[V3.1] z/Architecture Support

z/Architecture (64-bit) capabilities are supported for guest operating systems.

[V4.1] Integrated Facility for Linux

z/VM supports the IBM S/390 Integrated Facility for Linux (IFL), a hardware feature available on zSeries, G6, and G5 servers which provides additional processing capacity for Linux workloads. CMS Level 17 or later, Linux for zSeries, Linux for S/390, z/VM V4.1 or later, and stand-alone utilities supplied with z/VM V4.1 or later are supported on IFL.

IFL is managed by Processor Resource/Systems Manager™ (PR/SM™) as a logical partition with dedicated CPs. Implementation of this facility requires an LPAR definition, following normal LPAR activation procedures. As with any change in the LPAR configuration of a processor, the introduction of additional resources to manage may have an impact on the capacity of the existing partitions and workloads running on the server. The size of the impact is dependent on the quantity of added resources and the type of applications being introduced. A tool is available to aid in assessing the impact to any server; contact your IBM representative for details.

[V4.2] Cryptographic Support

z/VM supports the IBM PCICA (PCI Cryptographic Accelerator) and the IBM PCICC (PCI Cryptographic Coprocessor) for Linux guest virtual machines. This support provides clear-key RSA support for Linux guests enabling hardware SSL acceleration on the zSeries and S/390 servers. A z/VM system can include Linux guests using the RSA-Assist support simultaneously with other VM guests using the CMOS Cryptographic support.

The following commands and statements were updated:
- CP QUERY CRYPTO command
- CP QUERY VIRTUAL CRYPTO command
- CP SET CRYPTO command
- CRYPTO directory control statement

For more information, see:
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Planning and Administration*

[V4.2] ESS Large Volumes

z/VM supports 3390 volumes greater than 10,017 cylinders on the ESS, up to the maximum supported by the DASD. This helps relieve device address constraints and improves disk resource utilization and storage administrator productivity by providing the ability to consolidate multiple disk volumes into a single address.

The following CP functions have been updated to reflect this new capacity:
- XLINK_DEVICE_DEFAULTS system configuration statement
- XLINK_VOLUME_INCLUDE system configuration statement
- MDISK directory control statement

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

[V4.2] FICON CTCA

z/VM supports FICON Channel-to-Channel communications between an IBM zSeries 900 (at the appropriate service level) and another z900 or an S/390 Parallel Enterprise Server G5 or G6. This enables more reliable and higher bandwidth host-to-host communication than is available with ESCON channels.

The CP DEFINE CU/CNTLUNIT command was updated.

Additional explanation was added to dynamic I/O return codes for control units.

For more information, see:
- *z/VM: CP Commands and Utilities Reference*

- *z/VM: CP Planning and Administration*

[V4.2] HiperSockets

z/VM supports the z/Architecture HiperSockets function for high-speed TCP/IP communication among virtual machines and logical partitions (LPARs) within the same zSeries server. The HiperSockets function uses an adaptation of the Queued Direct I/O (QDIO) high-speed I/O protocol. The HiperSockets function allows virtual machines and logical partitions to communicate internally over the memory bus using the internal-queued-direct (IQD) channel type in the z900.

Up to four IQD channels can be configured within a z900. Each IQD channel provides isolated communications among the logical partitions and virtual machines using that channel. There is no communication between different IQD channels, thereby providing communication security among separate groups of logical partitions and virtual machines connected to different IQD channels.

The z900 provides up to 1024 HiperSockets devices through which HiperSockets communication can be used, divided at your discretion among the four IQD channels.

z/VM supports HiperSockets for use by guest operating systems and by the TCP/IP server virtual machine. VM programs using traditional TCP/IP socket connections can communicate through HiperSockets with other VM programs, guest operating systems, and other logical partitions that are also using TCP/IP.

Note: z/VM support for HiperSockets requires a z900 server at EC level J10607 or J10608. For the most current information on the VM support for HiperSockets, see IBM: z/VM Operating System (www.ibm.com/vm/).

For customers who plan to run z/OS as a guest of VM, see IBM z Systems - Operating System - z/OS (www.ibm.com/systems/z/os/zos/) for the most current information on z/OS support for HiperSockets.

[V4.2] OSA-Express Token Ring

OSA-Express token ring support expands on previous OSA token ring support by adding 100 Mb/sec (megabit/second) support to the earlier 4 Mb/sec and 16 Mb/sec support.

The OSA-Express token ring feature supports the QDIO data transfer architecture. QDIO support requires z/VM V4.2.

The OSA-Express token ring feature continues to support non-QDIO environments, the traditional TCP/IP (LAN Channel Station (LCS)) and SNA (Link Services Architecture (LSA)) traffic. VM/ESA V2.3 or later is required.

z/VM V4.2 supports the OSA-Express token ring feature for use by guest operating systems and by the TCP/IP server virtual machine.

[V4.2] z/Architecture Guest Coupling Facility

VM guest coupling support has been enhanced to accommodate the z/Architecture guest coupling facility. These enhancements include the following:

- Larger Vector Support
- Coupling Facility Duplexing
- Multiple-Buffer Capability

- Message Architecture Enhancements
- New CP Commands:
  - DEFINE CFLINK
  - DETACH CFLINK
  - QUERY CFLINKS
  - QUERY VTOD
  - SET CFLINK
- Updated CP Commands and Directory Control Statements
  - DEFINE MSGPROC
  - SPECIAL Directory Control Statement

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: Running Guest Operating Systems*
- *z/VM: CP Planning and Administration*

[V4.3] Fibre Channel Protocol Guest Support

z/VM provides guest support for the new Fibre Channel Protocol (FCP) hardware channel. This channel allows a zSeries server to connect to a fibre-channel fabric. The z/VM support allows guest operating systems to access selected devices on Small Computer System Interface (SCSI) controllers connected to the fabric.

The programs operating in guest virtual machines on z/VM are responsible for providing the appropriate support to use and control the SCSI devices connected through the FCP channel. Your system administrator is responsible for ensuring the necessary level of access control and concurrent-access data integrity for the SCSI devices. z/VM itself neither uses any SCSI devices nor controls guest access to them. z/VM allows the dedication of zSeries subchannels on FCP channels to a guest. Each subchannel allows the guest to access any SCSI device for which access permission has been granted to the worldwide port name of the FCP channel within the fibre-channel infrastructure. Typical access controls within the infrastructure include zoning in the switches and Logical Unit Number-masking (LUN-masking) in the controllers.

When FCP-based SCSI support is available from Linux for zSeries, guest Linux operating systems can access data on the SCSI devices connected to the fibre-channel fabric. Access to SCSI devices by a Linux guest is governed by open standards for fibre-channel fabrics and controllers as well as by proprietary or commonly-available functions provided by individual vendors. Therefore, unlike the device isolation provided among virtual machines by z/VM and z/Architecture for devices connected to other types of channels, neither z/VM nor the FCP channels provide for the isolation of SCSI devices to a single virtual-machine configuration.

The following CP commands have been updated for this support:

- ATTACH
- DEFINE CHPID / PATH
- DETACH (Real Device)
- QUERY ALL
- QUERY CHPID / PATH
- QUERY (Real Device)
- QUERY VIRTUAL ALL
- QUERY (Virtual Device)

For more information, see *z/VM: CP Commands and Utilities Reference*.

[V4.3] Performance Monitor Enhancement

To correctly extract processor configuration data and performance data from each of the I/O processors on zSeries servers, CP has been updated to issue the new STSI (Store System Information) hardware instruction and record the appropriate monitor information. The Performance Reporting Facility (PRF) feature has been updated to process this new monitor data.

[V4.4] Cascaded FICON Directors

z/VM supports the zSeries capability to cascade two FICON directors within a fibre-channel fabric. The z/VM support for cascaded FICON directors is embodied in support for two-byte fibre-channel link addresses when defining and configuring control units within the fabric. Previously, only single-byte link addresses were needed in a single-switch fabric. Cascaded FICON directors require the use of two bytes of the standard three-byte fibre-channel link address; the third byte, for arbitrated-loop support, is not used in FICON fabrics. When using the new cascaded-FICON-director function, your hardware-I/O definitions need to be updated accordingly.

When CP is in control of the dynamic I/O, use the DEFINE CU and MODIFY CU commands to define a cascaded FICON director topology to the system. When HCD is in control of the dynamic I/O, use the ACTIVATE command from the HCD virtual machine to define the cascaded FICON director topology.

The following CP functions have been updated:
- DEFINE CU / CNTLUNIT command
- MODIFY CU / CNTLUNIT command

For more information, see *z/VM: CP Commands and Utilities Reference*.

[V4.4] Enhanced QDIO Performance

The QDIO architecture, originally introduced with the OSA-Express, was later extended to HiperSockets and the FCP channels. The architecture itself was extended in HiperSockets to include a new type of high-performance I/O interruption known as an adapter interruption. The use of adapter interruptions has been extended to the OSA-Express and FCP channels on the IBM eServer™ zSeries 990 (z990).

In addition to the use of adapter interruptions by the OSA-Express and FCP channels, the z990 server is designed to include a performance assist for the virtualization of adapter interruptions being given to operating systems running as guests of z/VM. This hardware performance assist is available to V=V guests (pageable guests) that support QDIO on z/VM V4.4.

This new IBM virtualization technology is designed to benefit all guest operating systems in environments where they can process adapter interruptions. This includes all users of HiperSockets, and guest operating systems that add adapter-interruption support for OSA-Express and FCP channels. With the enhancement of the TCP/IP stack in z/VM V4.4 to use adapter interruptions for OSA-Express, TCP/IP can benefit from this performance assist for both HiperSockets and OSA-Express.

The following CP functions have been added for this support:

- QUERY QIOASSIST command
- SET QIOASSIST command

The following CP functions have been updated for this support:
- QUERY VIRTUAL FCP command
- QUERY VIRTUAL OSA command

For more information, see *z/VM: CP Commands and Utilities Reference*.

[V4.4] ESS FlashCopy Version 2

z/VM supports the following ESS FlashCopy Version 2 enhancements:

- Data Set FlashCopy, which removes the restriction where source cylinders must be copied to the same physical target cylinders. The removal of this restriction is especially helpful when copying one VM minidisk to another. It allows a non-fullpack source minidisk to be copied to a non-fullpack target minidisk when the minidisks exist on different physical cylinder extents of the same physical volume or different physical volumes.
- Multiple Relationship FlashCopy, which allows a source to have multiple targets. This function allows one source volume to be copied to many target volumes.
- Elimination of Logical Storage System (LSS) constraint, which allows a source and target relationship to span across an LSS. This removes the restriction where both the source volume and the target volume must reside in the same logical control unit.

[V4.4] ESS Peer-to-Peer Remote Copy Enhancements

z/VM supports guest use of the ESS Peer-to-Peer Remote Copy Extended Distance (PPRC-XD) function. Guests who support PPRC-XD and have DATAMOVER authority in their user directory can copy full volumes of data in non-synchronous mode, which extends the distance between the primary and secondary ESS with a minimal effect on performance. PPRC-XD is suitable for data migration, backup, and disaster recovery procedures. Data can be copied at distances well beyond the 103 km supported for PPRC synchronous transmissions. Typically, the distance for PPRC-XD is limited only by the capabilities of the network and channel extension technologies.

Native z/VM support for PPRC-XD is provided by Device Support Facilities (ICKDSF), Version 1 Release 17, with ICKDSF running in a CMS virtual machine.

z/VM also provides guest support for PPRC Version 2, which is designed to offer an Asynchronous Cascading solution, providing a complete, consistent, and coherent copy of data at a remote site. Asynchronous Cascading provides a two-volume synchronous PPRC implementation with a non-synchronous third volume serving as a backup device that can provide a multi-site, long distance solution.

[V4.4] Extended Channel Measurement Data Support

z/VM supports the extended I/O-measurement facilities of the z990 server. This function, known as 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. Each extended I/O-measurement block has its own 64-bit address, allowing the measurement blocks to be stored in noncontiguous real memory. The extended measurement word provides channel measurement data for each I/O operation when it completes, streamlining measurement processing in the operating system.

[V4.4] Guest Coupling Enhancement

VM's virtual Coupling Facility (CF) support has been enhanced to allow z/VM systems to run as second-level (or higher) guests while simulating complete OS/390 and z/OS coupled sysplexes. This enhanced support loads the Coupling Facility Control Code (CFCC) from the Service Element of a zSeries or S/390 server into CF service virtual machines in a z/VM system running as a second-level (or higher) guest of z/VM V4.4. This allows you to test an OS/390 or z/OS Parallel Sysplex environment at any guest level.

Note: Each additional level of guest virtualization incurs a performance penalty that may make running a Parallel Sysplex impractical at higher guest levels.

[V4.4] Logical Channel Subsystems

A new logical channel subsystem (LCSS) structure is introduced with the z990 server. It is designed to allow the definition of more than one channel subsystem (CSS), 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.

Support for dynamic I/O configuration on z/VM V4.4 has been extended to allow channel paths, control units, and devices to be dynamically added, changed, and deleted in multiple logical channel subsystems when z/VM V4.4 is running on a z990 server with the applicable hardware support. When z/VM V4.4 is running on a z990 server that includes the hardware support, and z/VM is the controlling logical partition for dynamic-I/O configuration, z/VM is designed to handle all of the new elements of each CSS facility for changing your hardware I/O configuration. To dynamically change the I/O configuration, one of two methods can be employed:

- CP's suite of interactive dynamic-I/O-configuration commands
- HCD and HCM configuration-managment tools

The following CP functions have been changed for this support:

- DEFINE CHPID/PATH command
- DEFINE CU/CNTLUNIT command
- DEFINE DEVICE/IODEVICE command
- MODIFY CHPID/PATH command
- MODIFY CU/CNTLUNIT command
- MODIFY DEVICE/IODEVICE command
- DELETE CHPID/PATH command
- DELETE CU/CNTLUNIT command
- DELETE DEVICE/IODEVICE command
- QUERY CHPID command
- QUERY DYNAMIC_I/O command
- QUERY LPARS command
- DIAGNOSE code X'2AC' – HCD Dynamic I/O

For more information, see:

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

[V4.4] Support for 30 LPARs

Prior to the z990, there was an LPAR limit of 15. IBM plans to increase this limit for the z990 server. z/VM can handle I/O-configuration definition and dynamic-I/O configuration for up to 30 LPARs. In addition, the CP Monitor has been updated to allow performance data to be collected and recorded for z/VM systems running on servers with more than 15 LPARs.

[V4.4] IBM 3590 Model H

z/VM supports the IBM TotalStorage Enterprise Tape Drive 3590 Model H. This 384-track tape drive is faster and of higher capacity than the previous 256-track 3590 Model E. The 3590 Model H can be configured to emulate the device characteristics of either the 3590 Model B or the 3490 Model E. z/VM supports both emulation modes. The drive performs at 384-track speeds and capacity regardless of how it is configured.

The following external interfaces have been updated for this support:
- CP DIAGNOSE code X'210'
- CMS message DMS2139I

[V5.1] z/VM Operates Only on z/Architecture Servers

z/VM V5.1 operates only on IBM servers that support the IBM 64-bit z/Architecture.

z/VM V5.1 supports the enhanced z990 and the new z890, including supporting the enhanced LCSSs, spanned channels, and other server functions and features described below.

[V5.1] ASN-and-LX-Reuse Facility Support

z/VM adds support for guest use of the ASN-and-LX Reuse Facility when available on the processor. For example, z/OS Version 1 Release 6 will take advantage of this architectural extension to improve availability by reusing linkage indices in more circumstances than can currently be done.

The following CP commands are have been updated for this support:
- DISPLAY LKS
- DUMP LKS
- TRACE *mnemonic1*

For more information, see *z/VM: CP Commands and Utilities Reference*.

[V5.1] Enhanced LCSS Support

The scalability of LCSSs is being further expanded — you can now define up to four LCSSs on a single z990. The z890 can support up to two LCSSs. The channel subsystem structure now offers the following:

- Four LCSSs (z990), two LCSSs (z890)
  - Each LCSS can have up to 256 channels defined.
  - Each LCSS can be configured with 1 to 15 logical partitions (LPARs).
    - Cannot exceed 30 LPARs per system.

There is no change to the operating system maximums. One operating system image continues to support up to 256 Channel Path Identifiers (CHPIDs). z/VM V5.1 will support four LCSSs on a zSeries server with the capability to do dynamic-I/O configuration in any LCSS.

[V5.1] IBM ESS Model 750

z/VM V5.1 supports the IBM TotalStorage ESS Model 750. The Model 750 supports FlashCopy V1 and V2, as well as Peer-to-Peer Remote Copy (PPRC) V1 and V2. The Model 750 also includes support for the zSeries performance enhancers, and Parallel Access Volumes (PAV).

[V5.1] ESS PPRC over FCP Connections

z/VM V5.1 provides guest support for using the ESS PPRC function over FCP connections. Native support is provided by running Device Support Facilities (ICKDSF), Release 17, in a CMS virtual machine.

[V5.1] I/O Devices Not Supported

In addition to devices not supported by previous releases of z/VM, the following real I/O devices are not supported by z/VM V5.1:

- 3370 DASD
- 3375 DASD
- 3380 DASD on 3880 DASD Control Unit

  Note: Emulated 3380 devices on 3990 Model 2 or higher controllers will continue to be supported. This includes RAMAC emulated 3380 models and 3390 DASD running in 3380 track compatibility mode.
- Multiprise Internal Disk
- 3830 DASD Control Unit
- 3880 DASD Control Unit
- 3995 Optical Library Dataserver
- 9332 DASD
- 9335 DASD
- 9336 DASD, except simulated VDISKs and emulated SCSI LUNs
- 9340 DASD Subsystem, including all associated DASD and controllers
- 2440 Tape Unit
- 3420 Tape Unit
- 3422 Tape Unit and Control, except OMA/2 CD devices that emulate 3422s (supported for installation only)
- 3424 Tape Subsystem
- 3430 Tape Unit
- 9348 Tape Unit
- 3803 Tape Control Unit
- 9221 DASD/Tape Subsystem Control
- SDLC Integrated Communication Attachment
- BSC Integrated Communication Attachment
- CETI (9221) Integrated Communication Attachment

Omission of a device from this list does not mean that the device is supported. For a list of the devices that are supported by z/VM V5.1, see *z/VM: General Information*.

[V5.1] Improvements to Capacity Upgrade on Demand

z/VM V5.1 supports the On/Off Capacity on Demand (On/OffCoD) and the Capacity Backup Upgrade (CBU) functions on IBM zSeries servers and includes functional enhancements that can allow z/VM to:

- Recognize changed processor configuration settings on a zSeries system to:
  - Report the change in configuration to the system operator
  - Report the change in configuration to guests that support configuration change notification
- Recognize changed processor capacity settings on zSeries systems to:
  - Report the change in capacity to guests that support capacity measurement for billing purposes
  - Report the capacity change in z/VM monitor and accounting data

The following CP functions have been added for this support:

- QUERY CAPABILITY command determines the capability values of the processors in the configuration. A processor's capability value indicates its capability relative to other CPU models.
- Type D Accounting Record records the CPU capability of the processors in the configuration. This record is generated during system initialization and whenever the CPU capability changes.

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Planning and Administration*

[V5.1] OSA-Express Integrated Console Controller

With the new z890 and enhancements to the z990, IBM is introducing a new function for the OSA-Express 1000BASE-T Ethernet feature and a new Channel Path Identifier (CHPID) type, OSC. The new OSA-Express Integrated Console Controller (OSA-ICC) function supports TN3270E and non-SNA DFT 3270 emulation. Now, 3270 emulation for console session connections is integrated in the z990 and z890 through a port on the OSA-Express 1000BASE-T Ethernet feature. This can help eliminate the requirement for external console controllers (2074, 3174).

OSA-ICC support is available with the OSA-Express 1000BASE-T Ethernet feature, and can be configured on a port-by-port basis. A port on the 1000BASE-T Ethernet feature can be configured as an OSD, OSE, or OSC CHPID type. Use of the OSA-ICC and the OSC CHPID type is supported by z/VM V5.1.

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *Open Systems Adapter-Express Integrated Console Controller User's Guide*

[V5.1] PCIX Cryptographic Coprocessor

z/VM V5.1 provides z/OS and Linux for zSeries guest support for the PCIX Cryptographic Coprocessor (PCIXCC) feature available with the IBM z990 and z890. z/VM support includes:

- Dedicated-queue support for clear-key and secure-key cryptographic functions for z/OS guests
- Shared-queue and dedicated-queue support for clear-key cryptographic functions for Linux guests with up to 256 dedicated queues

The following CP functions have been changed for this support:
- CRYPTO directory control statement
- QUERY CRYPTO command
- QUERY VIRTUAL CRYPTO command

For more information, see:

- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*

[V5.1] Transparent Sharing of Additional Channel Resources Across LCSSs

When LCSSs were introduced, they were designed to provide transparent sharing of Internal Coupling Channels (ICs) and HiperSockets between separate LCSSs. This support has now been expanded to include the ICB-3, ICB-4, ISC-3, FICON Express, and OSA-Express features. They are now capable of being configured as Multiple Image Facility (MIF) spanning channels, allowing sharing of channel resources across LPARs.

Spanned channels can be shared among LPARs across LCSSs. ICB-3, ICB-4, ISC-3, FICON Express, and OSA-Express can be configured to multiple channel subsystems and are intended to be shared transparently by any or all of the configured LPARs without regard to the LCSS to which the LPAR is configured.

z/VM V5.1 supports all of the spanned channel types identified, and also supports internal spanned channels.

[V5.1] Up to 24 Processor Engines in a Single z/VM Image

z/VM supports up to 64 virtual processor engines in a single guest configuration. z/VM V5.1 allows up to 24 real processor engines in a single z/VM image on a z990. The specific workload will influence the efficiency with which a specific z/VM system can use larger numbers of engines. Generally, z/VM overhead is expected to be lower with fewer, more CPU-intensive guests than with many lightly loaded guests. Excessive overcommitment of storage could contribute to increased overhead as well.

[V5.1] IBM 3592

z/VM supports the IBM TotalStorage Enterprise Tape Drive 3592, which offers greater capacity and improved performance over previous 3590 tape drives. The 3592 supports standard read/write media as well as Write Once Read Many (WORM) media. z/VM provides both native and guest support for the 3592.

Message HCP2246E has been added for this support, and the following CP interface has been updated:
- DIAGNOSE code X'210'

[V5.1] Crypto Express2

Crypto Express2 is IBM's third-generation cryptographic hardware feature. Crypto Express2 combines the functions of the PCICA and PCXCC in one feature, providing improved secure-key operations. z/VM V5.1 provides z/OS and Linux for zSeries guest support for Crypto Express2 as described for the PCIXCC feature (see "[V5.1] PCIX Cryptographic Coprocessor" on page 63). Refer to the 2084/2086 PSP buckets to identify any required service for support of these features.

[V5.1] OSA-Express2

z/VM 5.1 supports the OSA-Express2, which is a new generation of the OSA-Express Gigabit Ethernet feature, as well as a 10-Gigabit Ethernet feature. This technology update can help to provide a balanced system with sufficient throughput to satisfy the bandwidth demands of your applications. The OSA-Express2 Gigabit Ethernet (GbE) and 10 Gigabit Ethernet (10 GbE) features support Queued Direct Input/Output (QDIO). Refer to the 2084/2086 PSP buckets to identify any required service for support of these features.

[V5.1] HyperSwap Support

z/VM is providing a new HyperSwap® function so that the virtual devices associated with one real disk can be swapped transparently to another. HyperSwap can be used to switch to secondary disk storage subsystems mirrored by Peer-to-Peer Remote Copy (PPRC). It may also be helpful in data migration scenarios to allow applications to use new disk volumes.

Geographically Dispersed Parallel Sysplex (GDPS®) 3.1 plans to exploit the new z/VM HyperSwap function to provide a disaster recovery solution for distributed applications, such as WebSphere®, that span z/OS images running natively and Linux guests running under z/VM. This disaster recovery solution requires GDPS, IBM Tivoli System Automation for Linux, Linux for zSeries, and z/VM V5.1.

To cover planned and unplanned outages, GDPS 3.1 is planning to provide the following recovery actions:

- In-place re-IPL of failing operating system images
- Site takeover/failover of a complete production site
- Coordinated planned and unplanned HyperSwap of storage subsystems transparently to the operating system images and applications using the storage

The following CP commands have been added:
- HYPERSWAP
- QUERY HYPERSWAP

For more information, see:

- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*

[V5.1] IBM TotalStorage DS8000 Series

z/VM V5.1 supports the IBM TotalStorage DS8000® series, which is designed to provide unmatched functionality, flexibility, and performance for enterprise disk storage. The DS8000 enables the creation of multiple storage LPARs (logical partitions), which can be used for completely separate production or test environments in a single physical DS8000. This may enable the use of one storage server where more than one was needed in the past. The DS8000 may also provide

up to six times the throughput of an ESS Model 800, which may enable workload consolidation into a single storage subsystem.

z/VM supports the definition of very large SCSI or extended count key data (ECKD™) disks on the DS8000. See "[V5.1] Large Disk Support on DS8000 and DS6000" on page 182.

The following CP functions have been updated for this support:
- EDEVICE system configuration statement
- SET EDEVICE command

For more information, see:
- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*

[V5.1] FICON Express2

z/VM V5.1 supports the FICON Express2 feature for the z890 and z990, which provides up to twice the number of channels in the same amount of physical space – up to 80 FICON Express2 channels on z890 and up to 240 FICON Express2 channels on z990. FICON Express2 is designed to deliver up to 270 megabyes per second (MBps) of throughput, depending upon application environment. By comparison, FICON Express delivers up to 170 MBps. A FICON Express2 channel can also process significantly more start I/Os per second than a FICON Express channel. FICON Express2 supports CHPID types FC (native FICON, including FICON CTC control units) and FCP (Fibre Channel Protocol for attachment to SCSI devices in z/VM and Linux on zSeries environments).

[V5.1] IBM TotalStorage DS6000 Series

z/VM V5.1 supports the IBM TotalStorage DS6000™ series, which is designed to deliver enterprise-class storage capabilities in a space-efficient, modular design. Depending on workload, the DS6000 can provide up to four times the maximum host I/O operations per second of an ESS Model 750.

z/VM V5.1 supports the definition of large SCSI FCP disks and 3390 Model 9 disks on the DS6000. See "[V5.1] Large Disk Support on DS8000 and DS6000" on page 182.

[V5.2] IBM System z9 Enterprise Class and System z9 Business Class

z/VM V5.2 supports the IBM System z9® Enterprise Class (z9 EC), formerly the IBM System z9 109 (z9-109), and the IBM System z9 Business Class (z9 BC), including the following new facilities and features.

Crypto Express2 Accelerator

z/VM V5.2 supports the Crypto Express2 Accelerator (CEX2A) in the same manner as prior cryptographic cards are supported for both dedicated and shared use by guests. The new cryptographic card support is compatible with existing support for other PCI cards.

The CEX2A is exclusive to the z9 EC and z9 BC. The PCICC and PCICA hardware are not supported on the z9 EC and z9 BC.

The following CP functions have been updated for this support:
- CPU directory control statement

- OPTION directory control statement
- QUERY CRYPTO command
- QUERY VIRTUAL CRYPTO command

For more information, see:

- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*

N_Port ID Virtualization

N_Port Identifier virtualization (NPIV) for FCP channels (CHPID type FCP) is designed to allow the sharing of a single physical FCP channel among multiple operating-system images, whether running in LPARs or as z/VM guests in virtual machines. z/VM V5.2 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. In particular, access controls based on the virtual port names may be applied in the fabric using standard mechanisms like zoning in the switches and logical unit number (LUN) masking in the storage controllers.

NPIV is the industry standard solution which supersedes the previously previewed FCP LUN Access Control.

NPIV is exclusive to the z9 EC and z9 BC servers and is applicable to all of the FICON features supported on them when configured as CHPID type FCP.

z/VM V5.2 supports guest use of NPIV when FCP subchannels are dedicated to a guest and supports CP use of NPIV for IBM SCSI FCP devices.

When NPIV is enabled, you can allow guest operating systems and z/VM users to query the hardware-defined virtual port names. The following CP command has been updated:

- QUERY FCP

Dynamic Addition or Deletion of a Logical Partition Name

z/VM V5.2 provides facilities to dynamically define and delete LPARs on z9 EC, z9 BC, z990, and z890 servers. This capability has been added to z/VM HCD/HCM support and to CP's dynamic I/O command interface.

The following CP functions have been added for this support:

- DEFINE LPAR command
- DELETE LPAR command

The following CP function has been updated:

- QUERY LPARS command

For more information, see *z/VM: CP Commands and Utilities Reference*.

Support for up to 60 LPARs

z/VM V5.2 supports the capability of the z9 EC server to define up to 60 LPARs, 15 per Logical Channel Subsystem.

Exploitation of Selected System z9 Instructions

z/VM V5.2 exploits selected instructions of the z9 EC and z9 BC, including:

- Guest support for the Signal Processor (SIGP) instruction Conditional-Emergency-Signal and Sense-Running-Status orders and for Program-Event-Recording 3
- CP exploitation of Program-Event-Recording 3, providing access to the guest breaking-event-address register to aid in debugging wild branches during virtual machine execution, and the Store-Clock-Fast Facility to reduce the overhead of Store Clock instructions and CP program tracing
- CP TRACE support for new instructions available with the z9 EC and z9 BC

The following CP commands have been updated:

- DISPLAY Registers
- TRACE TABLE

The following CP utility has been updated:

- TRACERED

The following new diagnose code has been added:

- DIAGNOSE code X'9C' Voluntary Time Slice Yield

[V5.2] Device Support Withdrawn

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

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

For a list of the real I/O devices that are supported by z/VM V5.2, see *z/VM: General Information*.

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

[V5.2] Enhanced IBM DS6000, DS8000, and ESS Support

z/VM V5.2 recognizes the DS6000 and DS8000 in their native control unit modes. That is, control unit mode 1750 will be recognized for the DS6000, and control unit mode 2107 will be recognized for the DS8000. In addition, the Enterprise Storage Server Models 750 and 800 will be recognized in their native control unit mode of 2105. z/VM will automatically virtualize the controller mode that a guest supports.

[V5.2] Enhanced IBM Parallel Access Volume Support

z/VM V5.2 supports IBM Parallel Access Volumes (PAVs) as linkable minidisks for guest operating systems such as z/OS that exploit the PAV architecture, and provides the potential benefit of PAVs for I/O issued to minidisks owned or shared by guests that do not exploit PAVs, such as CMS.

The following commands have been added for this support:

- DEFINE PAVALIAS
- QUERY VIRTUAL PAV

The following commands have been updated for this support:

- ATTACH

- DETACH (Real Device)
- DETACH (Virtual Device)
- QUERY VIRTUAL (Device) DETAILS
- QUERY VIRTUAL DASD DETAILS
- SET MDCACHE

[V5.2] Enhanced Performance Assists for z/VM Guests

z/VM V4.4 added virtualization technology that allows adapter interruptions to be passed directly to z/VM guests for HiperSockets, FCP, and OSA on System z9, z990, and z890 servers.

A complementary virtualization technology is being introduced for System z9, z990, and z890 servers:

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

These new hardware assists can allow a cooperating guest operating system to initiate QDIO operations directly to the applicable channel, without interception by z/VM, thereby helping to provide additional performance improvements.

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

z990 and z890 servers will require microcode level (MCL) updates for these new assists. Support is integrated in the licensed internal code (LIC) of System z9 servers. These assists are provided on System z9, z990, and z890 servers for the following channel path types: HiperSockets (CHPID type IQD), FICON (CHPID type FCP), and all OSA features (CHPID type OSD).

[V5.2] FICON Express4

The FICON Express4 feature, introduced on the z9 EC and z9 BC servers, provides improved capacity and performance of up to 4 gigabits per second (Gbps). The FICON Express4 feature also provides:

- Reduced cost of storage operations and shorter backup windows with faster channel link data rates
- Investment protection for storage infrastructure with manageable migration to higher performance provided by 1/2/4 Gbps auto-negotiating links
- Channel and link consolidation to help simplify management and reduce the cost of the storage infrastructure
- Cost-effective FICON exploitation for midrange and small enterprises with additional price granularity with 2-port or 4-port cards for the z9 BC

[V5.2] GARP VLAN Registration Protocol Support

Generic Attribute Registration Protocol (GARP) Virtual Local Area Network (VLAN) Registration Protocol, also referred to as GARP VLAN Registration Protocol (GVRP), is defined in the IEEE 802.1p standard and allows for the control of IEEE 802.1q VLANs. An Open Systems Adapter (OSA) communicates with

GVRP-capable (IEEE 802.1p) switches, registering the active VLANs on behalf of the z/VM TCP/IP stack and the z/VM virtual switch if so configured. This provides dynamic creation and management of VLANs, including pruning of unnecessary broadcasts and unicasts on switches connected through 802.1q trunk ports.

z/VM enables the virtual switch to support GVRP in a guest LAN environment to help reduce LAN traffic between the physical LAN and the z/VM guest. GVRP support is exclusive to System z9 servers and is applicable to all OSA-Express2 features when operating in QDIO mode.

The following CP functions have been updated:
- DEFINE VSWITCH system configuration statement
- DEFINE VSWITCH command
- QUERY CONTROLLER command
- QUERY VSWITCH command
- SET VSWITCH command

[V5.2] IPv6 HiperSockets Support

IPv6 support is offered for HiperSockets (CHPID type IQD), exclusive to System z9 servers. z/VM V5.2 supports IPv6 HiperSockets for TCP/IP and virtualizes IPv6 HiperSockets in a guest LAN environment.

[V5.2] Open Systems Adapter for the Network Control Program

The Open Systems Adapter for the Network Control Program (OSA NCP, CHPID type OSN) is a new function of the OSA-Express2 Gigabit Ethernet and 1000BASE-T Ethernet features and is available only on System z9 servers. OSA-Express2 OSN supports the Channel Data Link Control (CDLC) protocol and provides channel connectivity from System z9 operating systems to the IBM Communication Controller for Linux (CCL). OSA-Express2 OSN is exclusively for internal communication, LPAR-to-LPAR.

If you continue to need SNA solutions that require NCP functions, you can now consider CCL as a migration strategy to replace your IBM Communications Controllers (374x).

z/VM V5.2 provides guest support for the OSA-Express2 OSN through HCD and HCM and CP's dynamic I/O command interface.

The following CP functions have been updated for this support:
- DEFINE CHPID/PATH command
- DEFINE DEVICE/IODEVICE command
- QUERY CHPID command
- QUERY OSA command
- QUERY VIRTUAL FCP command
- QUERY VIRTUAL OSA command

For more information, see *z/VM: CP Commands and Utilities Reference*.

[V5.2] SCSI Disk I/O Performance Improvements

z/VM V5.2 provides performance improvements for SCSI disk I/O, including:
- QDIO efficiency improvements

The z/VM SCSI I/O subsystem (SCSI-driver stack) will now support and exploit the QDIO architecture's ability to chain address lists, allowing more data to be moved with a single I/O request. Chaining the address lists helps to reduce fragmentation of large channel programs, decreases CP overhead, and significantly lowers the number of I/O operations driven through the SCSI stack, thereby improving the performance of SCSI I/O.

- Paging and spooling optimization

  CP's paging and spooling operations on SCSI disks currently use FBA emulation to transform standard I/O requests (that is, channel programs) into SCSI commands. With this new optimization, CP's paging and spooling function will build and send SCSI commands directly to the z/VM SCSI stack, thereby improving efficiency and performance.
- Improved FBA emulation efficiency
  - CP's FBA emulation function will now read CMS I/O buffers more efficiently, reducing the number of I/O requests to the z/VM SCSI stack. This enhancement will help to significantly improve SCSI I/O performance for CMS guests.
  - FBA emulation will now handle LOCATE CCWs more efficiently, reducing the number of I/O requests to the z/VM SCSI stack. This enhancement helps to significantly improve SCSI I/O performance for Linux guests, including the length of time to IPL a Linux guest.
  - FBA emulation will now handle the FBA padding function (initiated using a CCW data count less than the associated LOCATE data count) more efficiently, reducing the number of I/O requests to the z/VM SCSI stack. This enhancement helps to significantly improve SCSI I/O performance for format functions.

[V5.3] FlashCopy Version 2 Support Enhancements

z/VM V5.3 support for the FlashCopy V2 feature of IBM System Storage® disk storage devices has been enhanced to simplify the tasks required to automate backups. This includes the capabilities to:

- Specify multiple target minidisks

  The CP FLASHCOPY command can now accept up to 12 target minidisks to be copied.
- Determine the status of FlashCopy requests

  The new CP QUERY VIRTUAL FLASHCOPY command allows the user to query the number of FlashCopy relationships active for one or more of their virtual DASD.
- Exploit hardware asynchronous cache destage and discard

  This is designed to eliminate delayed hardware response messages and provides quicker responses to the CP FLASHCOPY command. This makes the FlashCopy appear synchronous to the virtual machine and may simplify automating processes that exploit this technology.

The following commands have also been updated for this support:

- CPHX now has the ability to terminate a wait loop for FLASHCOPY.
- FLASHCOPY added a new SYNChronous parameter, can accept multiple targets, and performs I/O retry and more front-end checking.

In addition, z/VM has reduced the number of FlashCopy hardware-related error conditions that can be reflected to the guest for the z/VM FLASHCOPY command. z/VM will attempt to redrive the I/O on some error conditions before reflecting the command response back to the guest.

[V5.3] Guest ASCII Console Support

Guest ASCII console support allows the real integrated ASCII console presented by the Hardware Management Console (HMC) to be dedicated to a guest, to which it appears as an integrated ASCII console.

This support includes new commands to

- Attach and detach the ASCII console to a guest
- Query the status of the guest ASCII console

For more information, see *z/VM: CP Commands and Utilities Reference*.

[V5.3] HyperParallel Access Volume Support

z/VM V5.3 supports the Hyper Parallel Access Volume (HyperPAV) function optionally provided by the IBM System Storage DS8000 disk storage systems. The HyperPAV function 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. z/VM provides support of HyperPAV volumes as linkable minidisks for guest operating systems, such as z/OS, that exploit the HyperPAV architecture. z/VM also transparently provides the potential benefits of HyperPAV volumes for minidisks owned or shared by guests that do not specifically exploit HyperPAV volumes, such as Linux and CMS.

The following CP functions have been added for this support:

- CU system configuration statement
- SYSTEM_ALIAS system configuration statement
- DEFINE HYPERPAVALIAS command
- QUERY CU command
- SET CU command

The following CP functions have been updated for this support:

- ATTACH command
- DEFINE command
- DETACH (Real Device) command
- QUERY CACHE command
- QUERY DASD DETAILS command
- QUERY PAV command
- QUERY VIRTUAL (Device) DETAILS command
- QUERY VIRTUAL DASD DETAILS command
- QUERY VIRTUAL PAV command
- SET MDCACHE command
- DIAGNOSE code X'A4' Synchronous I/O (Standard CMS Blocksize)
- DIAGNOSE code X'18' Standard DASD I/O
- DIAGNOSE code X'240'
- DIAGNOSE code X'250' Block I/O (Standard Blocksize)

- DASD Block I/O System Service (*BLOCKIO)

The following CP monitor records have been added or updated:

- New records:
  - Domain 1 Record 20 - MRMTRHPP - HyperPAV Pool Definition
  - Domain 6 Record 28 - MRIODHPP - HyperPAV Pool Activity
  - Domain 6 Record 29 - MRIODHPC - HyperPAV Pool Creation
  - Domain 6 Record 30 - MRIODLPT - LSS PAV Transition
  - Domain 8, Record 1 - MRVNDSES - Virtual QDIO Session Activity
  - Domain 8, Record 2 - MRVNDLSU - Virtual Network Guest Link State Change - Link Up
- Updated records:
  - Domain 1 Record 6 - MRMTRDEV - Device Configuration Data

    Fields have been added to contain a HyperPAV Base/Alias indicator and a HyperPAV Pool Number.
  - Domain 6 Record 1 - MRIODVON - Vary On Device - Event Data

    Fields have been added to contain a HyperPAV Base/Alias indicator and a HyperPAV Pool Number.
  - Domain 6 Record 3 - MRIODDEV - Device Activity

    Fields have been added to contain a HyperPAV Base/Alias indicator and a HyperPAV Pool Number and SCM type measurements for PAV/HyperPAV Alias I/O done on behalf of a PAV/HyperPAV Base device.
  - Domain 6 Record 20 - MRIODSTC - State change

    Fields have been added to contain a HyperPAV Base/Alias indicator and a HyperPAV Pool Number.

Note: If the HyperPAV feature is enabled on an IBM System Storage DS8000 Series DASD Subsystem, z/VM 5.3.0 by default will initialize each associated logical control unit (LCU) or subsystem so that it executes in “HyperPAV mode”. If the legacy PAV mode is required on one or more LCUs, the CU configuration file statement or CP SET CU command provides the capability to revert the required LCUs into “PAV mode”.

[V5.3] IBM System Storage SAN Volume Controller (2145) Device Support

The 2145 permits you to combine several heterogeneous storage subsystems behind a single piece of hardware that can be presented to host operating systems as a single cohesive pool of available SAN storage. This allows the host operating system (z/VM) to use a generic SCSI driver customized for the 2145 and rely on the 2145 to handle any of the device-specific requirements. As a result, subsystems can be added, removed, or reconfigured without disrupting the availability of the devices seen by the host, as well as making certain hardware features (such as Copy Services) available that might not otherwise be found because of the hardware involved.

The new 2145 attribute has been added to the EDEVICE configuration statement.

The new 2145 attribute has been added to the SET EDEVICE command. In addition, the QUERY EDEVICE command has been updated to display the 2145 information.

[V5.3] Improved Memory Management for Linux Guests

z/VM V5.3 supports the Collaborative Memory Management Assist (CMMA) on System z9 servers. This z/VM support, in conjunction with CMMA exploitation in 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, in the following ways:

- CP knows when a Linux application releases storage, allowing CP to select those pages for removal at a higher priority and to reclaim the page frames without the overhead of paging-out their data contents to expanded storage or disk.
- CP recognizes clean disk cache pages, the contents of which Linux is able to reconstruct, allowing CP to bypass paging-out the data contents when reclaiming the backing frames for these pages. If Linux or its application subsequently tries to refer to the discarded page, Linux is notified that the page has been discarded and can reread the contents from disk or otherwise reconstruct them.
- The guest further benefits from the Host Page-Management Assist (HPMA). In conjunction with CMMA, HPMA allows the machine to supply fresh backing page frames for guest memory when the guest reuses a previously discarded page, eliminating the need for z/VM hypervisor interception and resolution of these host page faults.

The following changes have been made to CP commands:

- The response to the K option on the DISPLAY (Guest Storage) commands for ESA/390, ESA/XC, and z/Architecture virtual machines includes the block-usage state and the block-content state.
- The CP QUERY command has a new operand, MEMASSIST, that displays the status of the assist.
- The CP SET command has a new operand, MEMASSIST, that controls the use of the assist.

A new option has been added to the OPTION directory control statement: NOMEMASSIST. The DirMaint feature has been updated to support this directory option.

For more information about the assist, see "Collaborative Memory Management Assist", in *z/VM: CP Programming Services*.

[V5.3] MIDAW Support

z/VM support for MIDAWs (modified-indirect-data-address words) allows guests to exploit MIDAWs when running on servers that provide MIDAW support. MIDAWs allow a program to transfer large amounts of data that span noncontiguous blocks in main storage without the overhead of data chaining and without the strict boundary and count restrictions imposed by IDAWs. The following CP external interfaces have been updated for this support:

- TRACE I/O command response
- TRACERED utility formatted I/O data
- DIAGNOSE code X'210' output

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Programming Services*

[V5.3] SCSI Device Support Enhancements

z/VM V5.3 provides additional enhancements for SCSI disk support for Linux users, including:

- Point-to-Point Fibre Channel links, as an alternative to the current requirement for a Fibre Channel switched fabric
- Dynamically determined preferred paths for emulated FBA devices (EDEVICEs) on SCSI disks in an IBM System Storage DS6000, instead of the current need to specify which paths are preferred in a SET EDEVICE command or an EDEVICE configuration file statement
- Faster formatting of emulated FBA devices (EDEVICEs) on SCSI disks in an IBM Enterprise Storage Server (ESS) or IBM System Storage DS8000
- Display of additional SCSI device characteristics when using the QUERY EDEVICE DETAILS command
- Checking for erroneous mapping of multiple emulated-device (EDEVICE) definitions onto the same SCSI disk when bringing emulated disks online

[V5.3] Specialty Processors Support

z/VM V5.3 provides guest support for virtual CPU types of zAAP (IBM zEnterprise® Application Assist Processor), zIIP (IBM z Integrated Information Processor), and IFL (IBM Integrated Facility for Linux), in addition to general-purpose CPs (central processors). A z/VM user can issue the DEFINE CPU command to define these types of virtual CPUs. The system administrator can issue the new SET CPUAFFINITY command to specify whether z/VM should create virtual specialty CPUs for a user by dispatching them on real CPUs that match their types (if available) or simulate them by dispatching them on real CPs. z/VM guests can create virtual specialty processors on processor models that support the same types of specialty processors but do not necessarily have them installed. zIIPs can be simulated only on System z10 and System z9 servers. zAAPs can be simulated only on System z10, System z9, and zSeries 990 and 890 servers.

The following CP commands have been added for this support:

- QUERY CPUAFFINITY
- SET CPUAFFINITY

The following CP commands have been updated for this support:

- BEGIN
- DEDICATE
- DEFINE
- DEFINE CPU
- DEFINE CRYPTO
- DETACH
- INDICATE LOAD
- INDICATE USER
- IPL
- LOGON
- QUERY PROCESSORS
- QUERY VIRTUAL CPUS
- SYSTEM
- UNDEDICATE

- VARY PROCESSOR

The CPU directory control statement has been updated for this support.

The accounting records for Virtual Machine Resource Usage (Record Type 1) and for CPU Capability (Record Type D) have been updated for this support.

The following CP monitor records are new for this support:

- Domain 0 (System)
  - Record 24 - MRSYTSPT - Scheduler Activity (per processor type)
- Domain 2 (Scheduler)
  - Record 12 - MRSCLSCA - SET CPUAFFINITY Changes

The following CP monitor records are changed for this support:

- Domain 0 (System)
  - Record 1 - MRSYTSYP - System Data (per processor)
  - Record 2 - MRSYTPRP - Processor Data (per processor)
  - Record 4 - MRSYTRSP - Real Storage Data (per processor)
  - Record 5 - MRSYTXSP - Expanded Storage Data (per processor)
  - Record 10 - MRSYTSCG - Scheduler Activity (global)
  - Record 11 - MRSYTCOM - Processor Communication Activities (per processor)
  - Record 12 - MRSYTUWT - User wait states
  - Record 13 - MRSYTSCP - Scheduler Activity (per processor)
  - Record 15 - MRSYTCUG - Logical CPU Utilization Data (global)
  - Record 16 - MRSYTCUP - CPU Utilization in a Logical Partition
  - Record 17 - MRSYTCUM - Physical CPU Utilization Data for LPAR Management
  - Record 22 - MRSYTSXP - System Execution Space (per processor)
- Domain 1 (Monitor)
  - Record 4 - MRMTRSYS - System Configuration Data
  - Record 5 - MRMTRPRP - Processor Configuration (per processor)
  - Record 7 - MRMTRMEM - Memory Configuration Data
  - Record 15 - MRMTRUSR - Logged on Users
- Domain 2 (Scheduler)
  - Record 4 - MRSCLADL - Add User to Dispatch List
  - Record 5 - MRSCLDDL - Drop User from Dispatch List
  - Record 6 - MRSCLAEL - Add User to Eligible List
  - Record 9 - MRSCLSHR - SET SHARE Changes
- Domain 3 (Storage)
  - Record 2 - MRSTORSP - Real Storage Activity (per processor)
  - Record 20 -MRSTOSXP - System Execution Space (per processor)
- Domain 4 (User)
  - Record 2 - MRUSELOF - User Logoff Data
  - Record 3 - MRUSEACT - User Activity Data
  - Record 4 - MRUSEINT - User Interaction Data
  - Record 5 - MRUSEDFC - DEFINE CPU

  - Record 6 - MRUSEDTC - DETACH CPU
  - Record 7 - MRUSERDC - DEFINE CPU n AS
  - Record 8 - MRUSETRE - User Transaction End
  - Record 9 - MRUSEATE - User Activity data at Transaction End
  - Record 10 - MRUSEITE - User Interaction data at Transaction End
- Domain 5 (Processor)
  - Record 1 - MRPRCVON - Vary On Processor
  - Record 3 - MRPRCPRP - Processor Data (per processor)

[V5.3] Tape Data Encryption Support

z/VM V5.3 supports drive-based data encryption with the IBM System Storage TS1120 Tape Drive (machine type 3592, model E05).

Encryption of tapes by z/VM itself requires that the IBM Encryption Key Manager be running on another operating system, using an out-of-band (such as TCP/IP) connection to the tape control unit. z/VM native support includes encryption for DDR and SPXTAPE, as well as transparent support for guests that do not provide for their own encryption (for example, CMS and Linux on z Systems).

z/VM also enables encryption of tapes by guests (such as z/OS) that have the ability to control the tape-encryption facilities themselves and to optionally run the Encryption Key Manager. Key management for such guests can use either an out-of-band or an in-band (such as an ESCON or FICON channel) connection between the Encryption Key Manager and the tape control unit. z/VM V5.3 allows any key label to be used, with key labels being accessible through a key alias that is defined to z/VM.

With the PTF for APAR VM64062, DFSMS/VM FL221 supports locating encryption-capable 3592 tape drives in an Enterprise Automated Tape Library. This DFSMS/VM support provides tape-encryption capabilities for a z/VSE® guest running on z/VM.

[V5.3] IBM System z10 Enterprise Class Support

z/VM V5.3 supports guest exploitation of the IBM System z10 Enterprise Class (z10 EC) at the level of System z9 functionality.

z/VM V5.3 also exploits selected functions of the z10 EC, including:

- Using dynamic I/O configuration to define, modify, and query a coupling over InfiniBand (CIB) CHPID when z/VM V5.3 is the controlling LPAR for I/O.
- Dynamically adding processors to or removing processors from an LPAR in reserve without preplanning.

Notes:

1. EREP support for the z10 EC requires the PTF for APAR VM64367.
2. CMS IOCP support for the z10 EC requires the PTF for APAR VM64302.
3. HCD support for the z10 EC requires the PTF for APAR VM64020.
4. OSA/SF support for the z10 EC requires the PTF for APAR OA23824.

[V5.3] DS8000 Dynamic Volume Expansion Support

z/VM V5.3 supports IBM System Storage DS8000 Dynamic Volume Expansion, which provides the ability to increase the size of a logical volume while it is online to a host system. The expanded size is not available on z/VM until the device is varied offline and online.

The following commands have been updated for this support:

- QUERY DASD
- QUERY EDEVICE

[V5.3] InfiniBand Based Coupling CHPID Support

z/VM V5.3 supports the coupling over InfiniBand (CIB) channel-path type for coupling channels that are built on InfiniBand protocols and connect to InfiniBand links.

The following CP commands have been added for this support:

- QUERY LSYSTEM
- SET LSYSTEM

The following CP commands have been updated for this support:

- DEFINE
- DEFINE CHPID / PATH
- MODIFY CHPID / PATH
- QUERY CHPID
- QUERY CONFIGMODE
- SET
- SET CONFIGMODE

The following CP utility has been updated for this support:

- IOCP

[V5.3] OSA-Express3 10 Gigabit Ethernet Support

z/VM V5.3 supports the OSA-Express3 10 Gigabit Ethernet (10 GbE) feature on the z10 EC, which can provide performance benefits for TCP/IP and z/VM virtual switches. OSA-Express3 10 GbE adapters can be included in link aggregation groups.

[V5.3] Tape Data Encryption Rekey Support

z/VM V5.3 provides drive-based encryption rekey support with the IBM System Storage TS1120 Tape Drive (machine type 3592, model E05). This support builds on the Tape Data Encryption Support that was previously available by providing z/VM with a method to change the key labels stored on a tape cartridge. Only the key labels change in this case, which eliminates the requirement to rewrite the entire data portion of the tape when changing the access rights.

[V5.4] Additional IBM System z10 EC Support

z/VM V5.4 supports guest exploitation of the following z10 EC facilities:

- Execute-Extensions Facility
- General-Instruction-Extension Facility
- Parsing-Enhancement Facility

The VM Dump Tool supports these new hardware instructions.

The following CP commands have been updated for this support:
- DISPLAY (Guest Storage - ESA/390)
- DISPLAY (Guest Storage - ESA/XC)
- DISPLAY (Guest Storage - z/Architecture)
- DISPLAY (Host Storage)
- DUMP (Guest Storage - ESA/390)
- DUMP (Guest Storage - ESA/XC)
- DUMP (Guest Storage - z/Architecture)
- DUMP (Host Storage)
- TRACE

For more information, see "CP Commands" on page 289.

[V5.4] z/VM Mode LPAR and Specialty Processors Support

z/VM V5.4 supports the z/VM mode logical partition on the IBM System z10™. Multiple types of IBM Z® processors can be defined in the same z/VM mode LPAR, including:
- Central Processor (CP)
- Integrated Facility for Linux (IFL)
- IBM z Integrated Information Processor (zIIP)
- IBM zEnterprise Application Assist Processor (zAAP)
- Internal Coupling Facility (ICF)

Two new commands, SET VCONFIG MODE and QUERY VCONFIG, are provided so that users can set and query the virtual configuration mode of their virtual machines to allow certain sets of virtual specialty processors. This builds upon the guest support for virtual specialty processors that was introduced in z/VM V5.3, and now includes support for a virtual CPU type of ICF. The DEFINE CPU command can be used to define ICF and other types of virtual CPUs. The QUERY PROCESSORS command has been updated to include a new EXPANDED operand that displays the logical partition mode.

The SET SHARE command has also been enhanced to provide the ability to set a user's system-resource-access priority for each processor type. The QUERY SHARE command now displays a user's share setting by processor type.

The following CMS utility has been updated for this support:
- ACCOUNT

The following CP commands have been added for this support:
- QUERY VCONFIG
- SET VCONFIG

The following CP commands have been updated for this support:
- AUTOLOG
- DEFINE CPU
- INDICATE LOAD
- INDICATE USER
- QUERY PROCESSORS
- QUERY SHARE
- QUERY VIRTUAL CPUS
- SET SHARE

- XAUTOLOG

The following directory control statements have been updated for this support:
- CPU statement
- OPTION statement

The accounting records for Virtual Machine Resource Usage (Record Type 1) have been updated for this support.

The following DIAGNOSE code has been updated for this support:
- DIAGNOSE Code X'2FC'

The following monitor records have been updated to indicate that type ICF is now a possible value:
- Domain 0 Record 1 - MRSYTSYP - System Data (per processor)
- Domain 0 Record 2 - MRSYTPRP - Processor Data (per processor)
- Domain 0 Record 4 - MRSYTRSP - Real Storage Data (per processor)
- Domain 0 Record 5 - MRSYTXSP - Expanded Storage Data (per processor)
- Domain 0 Record 11 - MRSYTCOM - Processor Communication Activities (per processor)
- Domain 0 Record 13 - MRSYTSCP - Scheduler Activity (per processor)
- Domain 0 Record 15 - MRSYTCUG - Logical CPU Utilization Data (global)
- Domain 0 Record 22 - MRSYTSXP - System Execution Space (per processor)
- Domain 0 Record 24 - MRSYTSPT - Scheduler Activity (per processor type)
- Domain 1 Record 5 - MRMTRPRP - Processor Configuration (per processor)
- Domain 3 Record 2 - MRSTORSP - Real Storage Activity (per processor)
- Domain 3 Record 20 - MRSTOSXP - System Execution Space (per processor)
- Domain 4 Record 2 - MRUSELOF - User Logoff Data - Event Record
- Domain 4 Record 3 - MRUSEACT - User Activity Data
- Domain 4 Record 4 - MRUSEINT - User Interaction Data
- Domain 4 Record 5 - MRUSEDFC - DEFINE CPU - Event Record
- Domain 4 Record 6 - MRUSEDTC - DETACH CPU - Event Record
- Domain 4 Record 7 - MRUSERDC - DEFINE CPU n AS - Event Record
- Domain 4 Record 8 - MRUSETRE - User Transaction End - Event Record
- Domain 4 Record 9 - MRUSEATE - User Activity data at Transaction End - Event Record
- Domain 4 Record 10 - MRUSEITE - User Interaction data at Transaction End - Event Record
- Domain 5 Record 1 - MRPRCVON - Vary On Processor - Event Data
- Domain 5 Record 3 - MRPRCPRP - Processor Data (per processor)

The following monitor records have been updated to add new fields:
- Domain 0 Record 12 - MRSYTUWT - User Wait State

  The following fields were added for ICF processors:
  - SYTUWT_CALLLICF
  - SYTUWT_CALCFICF
  - SYTUWT_CALSWICF
  - SYTUWT_CALCWICF
  - SYTUWT_CALCRICF
- Domain 1 Record 4 - MRMTRSYS - System Configuration Data

  The following field was added to indicate the processor configuration mode:
  - MTRSYS_SYSCMODE
- Domain 1 Record 15 - MRMTRUSR - Logged on User

  The following field was added to indicate the number of guest-defined ICF CPUs:

  - MTRUSR_CALICFCT

  The following fields were added to display the user's system-resource-access-priority for CP processors:
  - MTRUSR_CP_SSHRELSH
  - MTRUSR_CP_SSHABSSH
  - MTRUSR_CP_SSHMXSHR
  - MTRUSR_CP_SSHFLG1

  The following fields were added to display the user's system-resource-access-priority for zAAP processors:
  - MTRUSR_ZAP_SSHRELSH
  - MTRUSR_ZAP_SSHABSSH
  - MTRUSR_ZAP_SSHMXSHR
  - MTRUSR_ZAP_SSHFLG1

  The following fields were added to display the user's system-resource-access-priority for IFL processors:
  - MTRUSR_IFL_SSHRELSH
  - MTRUSR_IFL_SSHABSSH
  - MTRUSR_IFL_SSHMXSHR
  - MTRUSR_IFL_SSHFLG1
  -

  The following fields were added to display the user's system-resource-access-priority for ICF processors:
  - MTRUSR_ICF_SSHRELSH
  - MTRUSR_ICF_SSHABSSH
  - MTRUSR_ICF_SSHMXSHR
  - MTRUSR_ICF_SSHFLG1

  The following fields were added to display the user's system-resource-access-priority for zIIP processors:
  - MTRUSR_ZIP_SSHRELSH
  - MTRUSR_ZIP_SSHABSSH
  - MTRUSR_ZIP_SSHMXSHR
  - MTRUSR_ZIP_SSHFLG1
- Domain 2 Record 9 - MRSCLSHR - SET SHARE Changes - Event Record

  The following fields were added to display the user's system-resource-access-priority for CP processors:
  - SCLSHR_CP_SSHRELSH
  - SCLSHR_CP_SSHABSSH
  - SCLSHR_CP_SSHMXSHR
  - SCLSHR_CP_SSHFLG1

  The following fields were added to display the user's system-resource-access-priority for zAAP processors:
  - SCLSHR_ZAP_SSHRELSH
  - SCLSHR_ZAP_SSHABSSH
  - SCLSHR_ZAP_SSHMXSHR
  - SCLSHR_ZAP_SSHFLG1

  The following fields were added to display the user's system-resource-access-priority for IFL processors:
  - SCLSHR_IFL_SSHRELSH
  - SCLSHR_IFL_SSHABSSH
  - SCLSHR_IFL_SSHMXSHR
  - SCLSHR_IFL_SSHFLG1

  The following fields were added to display the user's system-resource-access-priority for ICF processors:
  - SCLSHR_ICF_SSHRELSH

  - SCLSHR_ICF_SSHABSSH
  - SCLSHR_ICF_SSHMXSHR
  - SCLSHR_ICF_SSHFLG1

  The following fields were added to display the user's system-resource-access-priority for zIIP processors:
  - SCLSHR_ZIP_SSHRELSH
  - SCLSHR_ZIP_SSHABSSH
  - SCLSHR_ZIP_SSHMXSHR
  - SCLSHR_ZIP_SSHFLG1
- Domain 4 Record 1 - MRUSELON - User Logon - Event Record

  The following fields were added to display the user's system-resource-access-priority for CP processors:
  - USELON_CP_SSHRELSH
  - USELON_CP_SSHABSSH
  - USELON_CP_SSHMXSHR
  - USELON_CP_SSHFLG1

  The following fields were added to display the user's system-resource-access-priority for zAAP processors:
  - USELON_ZAP_SSHRELSH
  - USELON_ZAP_SSHABSSH
  - USELON_ZAP_SSHMXSHR
  - USELON_ZAP_SSHFLG1

  The following fields were added to display the user's system-resource-access-priority for IFL processors:
  - USELON_IFL_SSHRELSH
  - USELON_IFL_SSHABSSH
  - USELON_IFL_SSHMXSHR
  - USELON_IFL_SSHFLG1

  The following fields were added to display the user's system-resource-access-priority for ICF processors:
  - USELON_ICF_SSHRELSH
  - USELON_ICF_SSHABSSH
  - USELON_ICF_SSHMXSHR
  - USELON_ICF_SSHFLG1

  The following fields were added to display the user's system-resource-access-priority for zIIP processors:
  - USELON_ZIP_SSHRELSH
  - USELON_ZIP_SSHABSSH
  - USELON_ZIP_SSHMXSHR
  - USELON_ZIP_SSHFLG1
- Domain 4 Record 2 - MRUSELOF - User Logoff - Event Record

  The following fields were added to display the user's system-resource-access-priority for CP processors:
  - USELOF_CP_SSHRELSH
  - USELOF_CP_SSHABSSH
  - USELOF_CP_SSHMXSHR
  - USELOF_CP_SSHFLG1

  The following fields were added to display the user's system-resource-access-priority for zAAP processors:
  - USELOF_ZAP_SSHRELSH
  - USELOF_ZAP_SSHABSSH
  - USELOF_ZAP_SSHMXSHR
  - USELOF_ZAP_SSHFLG1

The following fields were added to display the user's system-resource-access-priority for IFL processors:
- USELOF_IFL_SSHRELSH
- USELOF_IFL_SSHABSSH
- USELOF_IFL_SSHMXSHR
- USELOF_IFL_SSHFLG1

The following fields were added to display the user's system-resource-access-priority for ICF processors:
- USELOF_ICF_SSHRELSH
- USELOF_ICF_SSHABSSH
- USELOF_ICF_SSHMXSHR
- USELOF_ICF_SSHFLG1

The following fields were added to display the user's system-resource-access-priority for zIIP processors:
- USELOF_ZIP_SSHRELSH
- USELOF_ZIP_SSHABSSH
- USELOF_ZIP_SSHMXSHR
- USELOF_ZIP_SSHFLG1

The following monitor records have had comments updated to explain that the system-resource-access-priority values are given for the CPU and are not global values for the user:
- Domain 2 Record 4 - MRSCLADL - Add User to Dispatch List - Event Record
- Domain 2 Record 5 - MRSCLDDL - Drop user from Dispatch List - Event Record
- Domain 2 Record 6 - MRSCLAEL - Add User to Eligible List - Event Record
- Domain 4 Record 3 - MRUSEACT - User Activity Data
- Domain 4 Record 9 - MRUSEATE - User Activity data at Transaction End - Event Record

[V5.4] Installation of Linux on z Systems from the HMC

Through this support, you can install Linux through a Hardware Management Console (HMC) DVD or flash drive. The support allows you to avoid connection to an intranet, which is a possible security issue. The z/VM FTP server supports file transfer to and from the HMC removable media. Through FTP, you can retrieve Linux installation files from the HMC DVD or flash drive and place the files onto a z/VM minidisk or SFS directory.

Note: This support is intended for customers who have no alternative, such as a LAN-based server, for serving the DVD contents for Linux installations. The elapsed time for installation using the HMC DVD drive can be an order of magnitude, or more, longer than the elapsed time for LAN-based alternatives.

The z/VM FTP server supports the following FTP functions with an HMC directory:
- Change directory to HMC directory
- File transfer using type ASCII and IMAGE (GET, PUT, Store Unique, APPEND)
- Directory/file listing (DIR and NLST)
- File deletion (DELETE)
- File sizing (SIZE)
- Authorization of HMC directory access
- Buffer changes

For more information, see *z/VM: Getting Started with Linux on z Systems*.

[V5.4] Exploitation of OSA-Express3 Ports

z/VM V5.4 recognizes all four ports on an OSA-Express3 Gigabit Ethernet (GbE) machine feature. There are two PCI Express (PCI-E) adapters per feature, with two ports per adapter, for a total of four ports per feature. For more information, see *z/VM: Connectivity*.

The following CP commands have been updated for this support:
- DEFINE VSWITCH
- QUERY LAN
- QUERY PORT
- QUERY VIRTUAL NIC
- QUERY VSWITCH
- SET PORT GROUP
- SET VSWITCH

The following CP system configuration file statements have been updated for this support:
- DEFINE VSWITCH
- MODIFY PORT

The following TCP/IP commands have been updated for this support:
- IFCONFIG
- IPWIZARD
- NETSTAT DEVLINKS

The following TCP/IP configuration statement has been updated for this support:
- DEVICE (OSD)

The following DIAGNOSE code has been updated for this support:
- DIAGNOSE Code X'26C'

The following monitor records have been updated for this support:
- MRIODVSW - VSWITCH Activity Sample Record
- MRIODVSF - VSWITCH Failure Event Record
- MRIODVSR - VSWITCH Recovery Event Record

[V5.4] Virtual CPU SHARE Redistribution Support

A guest's share is distributed as equally as possible among its started virtual processors by dividing its share by the number of nondedicated virtual processors in the configuration. This redistribution occurs when a SET SHARE command is issued or when a virtual configuration change alters the number of virtual processors. Virtual CPU SHARE Redistribution also performs share redistribution when a virtual processor is started or stopped.

The following command has been updated for this support:
- SET SHARE

The following CP monitor records have been updated:
- Domain 1 Record 15 - MRMTRUSR - Logged on Users
  - MTRUSR_CALSHARF - This field was updated to indicate that a virtual processor is stopped or in a hard stopped state.
- Domain 4 Record 3 - MRUSEACT - User Activity Data:
  - USEACT_CALSHARF - This field was updated to indicate that a virtual processor is stopped or in a hard stopped state.

  - USEACT_VMDCTSTA - This field accumulates the number of times the virtual CPU has been started.
  - USEACT_VMDCTSTO - This field accumulates the number of times the virtual CPU has been stopped.
- Domain 4 Record 4 - MRUSEINT - User Interaction Data
  - USEINT_CALCWAIT - This field was updated to indicate whether or not the virtual processor is in a hard stopped state.
  - USEINT_VMDCTSTA - This field accumulates the number of times the virtual CPU has been started.
  - USEINT_VMDCTSTO - This field accumulates the number of times the virtual CPU has been stopped.
- Domain 4 Record 9 - MRUSEATE - User Activity data at Transaction
  - USEATE_CALSHARF - This field was updated to indicate that a virtual processor is stopped or in a hard stopped state.

[V5.4] System z10 BC and Enhanced System z10 EC

z/VM V5.4 supports the IBM System z10 Business Class (z10 BC) server and the enhanced (GA2) System z10 Enterprise Class (z10 EC) server.

The following additional support is required:

- EREP/VM requires APAR VM64475 to support the z10 BC.
- HCD for z/VM requires the PTF for APAR VM64410.
- OSA/SF requires the PTF for APAR OA26286.
- IOCP requires the PTF for APAR VM64474.

[V5.4] Additional OSA-Express3 Support

z/VM V5.4 recognizes all four ports on the OSA-Express3 1000BASE-T Ethernet feature on the z10 EC. There are two PCI Express (PCI-E) adapters per feature, with two ports per adapter, for a total of four ports per feature. The OSA-Express3-2P Gigabit Ethernet and 1000BASE-T Ethernet features on the z10 BC provide two ports per feature.

[V5.4] IBM FlashCopy SE

z/VM supports the IBM FlashCopy SE feature on the IBM DS8000, providing an instantaneous space-efficient snapshot capability that can greatly reduce the storage capacity needed for point-in-time copies. With a FlashCopy SE relationship, disk space will be consumed for the target copy only when data is written to the source volume or when a write is directed to the target. For a source volume without much write activity, the target volume can consume significantly less physical space than the source. In addition to the benefit of more efficient storage utilization, less capacity can mean fewer disk drives and lower power and cooling requirements. FlashCopy SE might be especially useful in the creation of temporary copies for tape backup, online application checkpoints, or copies for preproduction or disaster recovery testing.

The FLASHCOPY command has been enhanced to establish and manage persistent FlashCopy relationships, including those that specify a space-efficient volume as the target. This command also provides:

- The ability to utilize the FlashCopy NOCOPY feature, in order to prevent the DASD subsystem from copying unmodified source tracks to the corresponding target.
- Options to preserve or change the volume label directly from the FLASHCOPY command.

A new QUERY FLASHCOPY command has been created to allow the interrogation of persistent FlashCopy relationships that have been established on specified real devices. This includes information about relationships established from other LPARs.

The QUERY VIRTUAL FLASHCOPY command has been enhanced to provide detailed information about persistent FlashCopy relationships that affect the specified virtual machine.

The following CP commands were added for this support:
- FLASHCOPY BACKGNDCOPY
- FLASHCOPY ESTABLISH
- FLASHCOPY RESYNC
- FLASHCOPY TGTWRITE
- FLASHCOPY WITHDRAW
- QUERY FLASHCOPY

The following CP commands have been updated for this support:
- ATTACH
- DEFINE
- DEFINE MDISK
- FLASHCOPY
- LINK
- QUERY CU DASD
- QUERY DASD
- QUERY VIRTUAL DASD
- QUERY VIRTUAL FLASHCOPY

For more information, see *z/VM: CP Commands and Utilities Reference*.

For information about using persistent FlashCopy and space-efficient volumes, see *z/VM: CP Planning and Administration*

[V5.4] Additional Tape Encryption Support

This z/VM enhancement is designed to support drive-based data encryption with the IBM System Storage TS1130 Tape Drive (machine type 3592, model E06) to help protect data on tape in a cost-effective way. Guest operating systems running under z/VM can take advantage of these encryption features through z/VM, even if the guest does not exploit them itself for reading and writing tapes. In addition, this tape drive can be used for native z/VM tape functions, such as SPXTAPE, DDR, and CMS TAPE.

DIAGNOSE code X'210' has been updated to support a new VRDCBLOK CLASTAPE value for the 3592 E06. For more information, see *z/VM: CP Programming Services*.

With the PTF for APAR VM64458, DFSMS/VM FL221 supports locating encryption-capable 3592 tape drives in an Enterprise Automated Tape Library. This DFSMS/VM support provides an environment for a z/VSE guest running on z/VM to exploit tape encryption.

[V5.4] Coupling Facility Enhancement

z/VM's Coupling Facility simulation has been enhanced to support Coupling Facility Control Code (CFCC) Release 16.

[V5.4] OpenSolaris Supported on IFLs

IBM has expanded the authorized use of an IBM Integrated Facility for Linux (IFL) feature. An IFL may now be used for execution of the OpenSolaris operating system in addition to or instead of the Linux operating system, respective OpenSolaris or Linux workloads, z/VM, and z/VM applications in support of either operating system.

Previously, z/VM was authorized to run on IFL processors only if, on the IFL processors, z/VM was being used exclusively to run Linux workloads and, if required, z/VM applications in support of those Linux workloads. z/VM is now also authorized to run on IFL processors if, on the IFL processors, z/VM is being used to run OpenSolaris workloads and, if required, z/VM applications in support of those OpenSolaris workloads.

This modification applies to all IFLs previously shipped on any IBM Z product, and to all newly ordered IFLs. Unless IBM specifies otherwise, non-IBM operating systems are offered by other parties alone; IBM does not warrant, and is not responsible for support of, a non-IBM operating system.

z/VM guest support for OpenSolaris requires DIAGNOSE code X'2A8'. See "[V5.4] DIAGNOSE Code X'2A8' – Network Diagnose" on page 258.

[V5.4] OSA-ICC 3215

The OSA-Express Integrated Console Controller (OSA-ICC) function has a built-in NIC/Ethernet port. Previously, the OSA-ICC supported only a 3270 (block mode) data stream, and not a 3215 (line mode) data stream. The 3215 support enhancement to the OSA-ICC function will enable it to receive channel programs with 3215 data streams and transmit this data to remote TCP/IP clients using the Telnet protocol.

Refer to the Preventive Service Planning (PSP) bucket for the minimum MCL level that supports the OSA-ICC 3215 function on the OSA-Express2 (System z9 or z10) or OSA-Express3 (System z10 only):

- DEVICE2094 for the z9 EC
- DEVICE2096 for the z9 BC
- DEVICE2097 for the z10 EC
- DEVICE2098 for the z10 BC

To support the OSA-ICC 3215 for z/TPF guests, the following CP functions have been added:

- RDEVICE (Terminals) system configuration statement
- SET RDEVICE (Terminals) command

The following CP command has been updated:

- DEFINE CHPID / PATH

[V6.1] Architecture Level Set

z/VM version 6 requires a new architecture level set (ALS) and supports only the IBM System z10 Enterprise Class (z10 EC) and z10 Business Class (z10 BC) and new generations of IBM Z servers, such as the IBM zEnterprise 196 (z196). Refer to the appropriate Preventive Service Planning (PSP) bucket for the minimum microcode level (MCL) and any required updates:

- DEVICE2097 for the z10 EC
- DEVICE2098 for the z10 BC
- DEVICE2817 for the z196

For specific server facilities required by z/VM V6.1, see IBM: z/VM V6.1 - Architecture Level Set (www.ibm.com/vm/zvm610/architecture/).

Documentation has been removed for the following CP commands, which no longer have any function:

- DEFINE CRYPTO
- DETACH CRYPTO
- DISPLAY CDX
- DUMP CDX
- SET CRYPTO
- STORE CDX

In the following CP interfaces, documentation has been removed for operands that no longer have any function:

- RDEVICE (Special Devices) system configuration statement
- CPU user directory statement
- CRYPTO user directory statement
- QUERY CRYPTO command
- QUERY VIRTUAL CRYPTO command
- SET RDEVICE (Special Devices) command

Documentation has been removed for the following TCP/IP server configuration statements, which no longer have any function:

- ATMARPSERVER statement
- ATMLIS statement
- ATMPVC statement
- DEVICE and LINK statements for ATM devices
- LINK statement for Token Ring Network or PC Network LCS
- LINK statement for FDDI LCS
- LINK statement for QDIOATM
- LINK statement for OSD Token Ring Network

In the following TCP/IP server configuration interfaces, documentation has been removed for operands that no longer have any function:

- DEVICE statement for LCS devices
- TRANSLATE statement
- IFCONFIG command

In the following systems management APIs, some parameters no longer have any function:

- Image_CPU_Define_DM
- Image_CPU_Query_DM

In the following DirMaint commands, documentation has been removed for operands that no longer have any function:

- CPU
- CRYPTO
- SETCPU

[V6.1] IBM System Storage Support

z/VM 6.1 provides support for IBM Full Disk Encryption and Solid State Disk features of the IBM System Storage DS8000. The QUERY DASD DETAILS command now indicates when a DASD volume is an encrypted volume or when it is comprised of Solid State Drives.

With the PTF for APAR VM64657, DFSMS/VM provides services that enable a z/VSE guest to manage resources of the IBM Virtualization Engine for the TS7720 configured without a physical tape library.

[V6.1] Worldwide Port Name Prediction Tool

A worldwide port name (WWPN) prediction tool is now available from IBM Resource Link to assist you with pre-planning of your Storage Area Network (SAN) environment prior to the installation of your System z10 server. This stand alone tool is designed to allow you to set up your SAN in advance, so that you can be up and running much faster once the server is installed. The tool assigns WWPNs to each virtual Fibre Channel Protocol (FCP) channel/port using the same WWPN assignment algorithms a system uses when assigning WWPNs for channels utilizing N_Port Identifier Virtualization (NPIV).

With the PTF for APAR VM64579, the Hardware Configuration Definition (HCD) facility of z/VM provides I/O device information from the input/output definition file (IODF) for input into the WWPN prediction tool. This PTF has been pre-applied to the HCD level that is included and preinstalled with z/VM V6.1.

For more information on setting up a SAN using the WWPN prediction tool, see the IBM Z Hardware Announcement dated April 28, 2009.

[V6.1] IBM zEnterprise 196

z/VM V6.1 includes support that enables guests to exploit selected new features provided by the IBM zEnterprise 196 (z196) server. In addition, z/VM will recognize and report processing capability reduction due to customer-initiated power-saving mode or an autonomic response to environmental conditions and will reflect the change and its cause in monitor and accounting data streams. Supporting guests will also be notified of capability changes.

Note: The TRACE MC command will not provide output when the Enhanced-Monitor Facility is used and a monitor-event-counting operation is performed.

The following CP functions have been updated for this support:

- QUERY CAPABILITY command response includes information about machine capacity changes.
- Accounting record format for CPU capability (Record Type D) includes new fields.
- Monitor records include new fields:
  - Domain 0 (System) Record 19 - MRSYTSYG - Primary and Secondary CPU Capability (global):
    - Nominal CPU capability in the configuration

  - Domain 1 (Monitor) Record 4 - MRMTRSYS - System Configuration Data:
    - Capacity-Change Reason
    - Capacity-Adjustment Indication
  - Domain 1 (Monitor) Record 18 - MRMTRCCC - CPU Capability Change:
    - Nominal CPU capability in the configuration
    - Capacity-Change Reason
    - Capacity-Adjustment Indication

Guest support for the following z196 facilities is available transparently without any changes required. Trace, DISPLAY I, and VMDUMPTL INSTR support is not provided for these facilities.

- Floating-Point Extension Facility
- New general instructions
- Fast-BCR-Serialization Facility
- CMPSC Enhancement Facility
- Non-quiescing Key-Setting Facility
- Message-Security-Assist Extension 4

[V6.1] Crypto Express3

z/VM V6.1 provides guest support for the new Crypto Express3 feature for the IBM System z10 and the IBM zEnterprise 196. The Crypto Express3 adapter, like the Crypto Express2, can be defined as either a coprocessor or as an accelerator and supports both encrypted key and clear key applications. The z/VM support is compatible with the guest support provided for the Crypto Express2 feature as described for the PCIXCC feature (see "[V5.1] PCIX Cryptographic Coprocessor" on page 63). In addition to the guest compatibility support, the QUERY CRYPTO APQS command has been enhanced to provide information about active users who have virtual crypto queues defined for sharing.

Responses for the following CP commands have been updated to include new AP types:

- QUERY CRYPTO
- QUERY VIRTUAL CRYPTO

The following CP monitor record has been updated to include the new types 8 and 9 when Crypto Express3 cards are installed:

- Domain 5 Record 10 - MRPRCAPM - Crypto Performance Measurement Data

[V6.1] Dynamic Emulated Device Path Control

z/VM V6.1 provides the ability to add and remove paths from an emulated device (representing a real SCSI device) while the device is online and possibly in use. This is especially useful for making configuration changes for z/VM system volumes that cannot be varied offline while z/VM is running.

[V6.1] HyperSwap Improvements

z/VM V6.1 provides finer control over z/VM's missing-interrupt automatic quiesce trigger. This support is enabled with new INCLUDE and EXCLUDE options on the HYPERSWAP command. For example, devices can be excluded from the trigger, or the trigger for a device can be made a multiple of the current z/VM MITIME setting.

[V6.1] IBM Extended Address Volumes Support

The Extended Address Volumes (EAV) function on the IBM DS8000 adds support for ECKD volumes up to 262,668 cylinders. z/VM V6.1 provides support for ECKD volumes up to 262,688 cylinders for:

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

CMS is enhanced to support volumes up to 65,520 ECKD cylinders for its own use, or about 45 GB of data. CP continues to support volumes up to 65,520 cylinders for its own use. The maximum size for FBA SCSI disks supported for use by CMS or GCS remains at 381 GB. IBM suggests that customers defining ECKD disks for use by CMS should set a practical limit of about 22 GB. If larger disks are defined, they should be limited to contain very few files, or the CMS file system might not be able to obtain enough virtual storage below 16 MB to format or access those disks. For more information, see the ACCESS command in *z/VM: CMS Commands and Utilities Reference*.

[V6.1] IBM XIV Storage Systems Direct Attachment Support

z/VM V6.1 supports the direct attachment of IBM XIV® Storage Systems for system use (such as paging, spooling, and IPL). This support eliminates the need for z/VM to access XIV devices through the IBM SAN Volume Controller (SVC) and provides the ability to define system and guest volumes as emulated devices on XIV devices.

The following CP functions have been updated:

- EDEVICE system configuration statement
- SET EDEVICE command

[V6.1] Protected Key CPACF Support

z/VM 6.1 provides support to enable guest virtual machines to use the new encrypted key functions of the CP Assist for Cryptographic Function (CPACF) available on the IBM System z10 and IBM zEnterprise 196.

[V6.1] OSX and OSM Channel Path Identifiers

z/VM V6.1 supports two new channel path ID (CHPID) types that have been added to support the private networks that interconnect the parts of the IBM zEnterprise System. These new CHPID types are based on the existing OSA-Express Direct (OSD) Ethernet CHPID:

- CHPID type OSX provides connectivity between the z196 and the intraensemble data network (IEDN) for application workloads.
- CHPID type OSM provides connectivity between the z196 and the intranode management network (INMN) for management related functions.

The following CP commands have been updated for this support:

- DEFINE
- DEFINE CHPID / PATH
- DEFINE CU
- QUERY CHPID
- QUERY FCP
- QUERY OSA

- VARY

The following CP monitor records have been updated for this support:
- Monitor Record Domain 0, Record 20 - Extended Channel Measurement Data supports new values for the OSX and OSM CHPIDs.
- Monitor Record Domain 1, Record 4 - System Configuration Data added a new field (MTRSYS_SYSENSBL) to indicate whether the system is in an ensemble.

The following dynamic I/O return codes have been modified for this support:
- Return code X'0119' - Adding a Channel Path
- Return code X'0105' - Adding an I/O Device

[V6.1] XRC Time Stamping Support

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).

The following CP command has been added:
- QUERY STP

The following CP commands have been updated:
- QUERY TIMEZONES
- SET TIMEZONE

[V6.2] IBM zEnterprise 196 (Enhanced) and IBM zEnterprise 114

z/VM V6.2 supports the IBM zEnterprise 196 (z196) and the IBM zEnterprise 114 (z114) and enables guests to exploit selected new features.

System topology information (Store System Information (STSI) instruction 15.1.x) for the z196 and z114 (and the z10) is collected in the CP monitor records. If the STSI 15.1.x support is available (not available on second-level systems), a configuration monitor record is used to collect the initial system topology and an event monitor record is cut whenever periodic polling by CP sees the system topology as changed.

CP will also include the first 180 bytes of the Basic-machine configuration (STSI 1.1.1), which includes the hardware model number, in a configuration monitor record (D1R4) and an event monitor record (D1R18).

z196 and z114 Driver D93 provides a new physical-system-management time (SysMgmtime) value, which CP will provide in a monitor record (D0R17).

The following CP monitor records have been added:
- Domain 1 Record 26 - MRMTRTOP - System Topology Configuration
- Domain 5 Record 14 - MRPRCTOP - System Topology

The following CP monitor records have been updated:
- Domain 0 Record 17 - MRSYTCUM - Physical CPU Utilization Data for LPAR Management
- Domain 1 Record 4 - MRMTRSYS - System Configuration Data
- Domain 1 Record 18 - MRMTRCCC - CPU Capability Change

[V6.2] ACCOUNT Utility Enhancement

The ACCOUNT utility has been enhanced with a new CPUCAP option that can create a CPU capability record based on CPU capability type 0D accounting records.

[V6.2] CPU-Measurement Counter Facility Host Support

CP can collect counter data generated by the CPU-Measurement Counter Facility and include the counter data in the CP MONITOR data stream. Support for the CPU-Measurement Sampling Facility and virtualization of the CPU-Measurement Facility interfaces for guest use are not provided.

The MONITOR SAMPLE ENABLE PROCESSOR command is used to enable collection of CPU-Measurement Facility counter data. If the logical partition is authorized for the basic counter set and collection is enabled, whichever basic, problem-state, crypto-activity, or extended counter sets that are authorized are enabled and activated when MONITOR START is issued. When a monitor sample is collected, the counters will be extracted on each processor and recorded in a new Processor domain monitor record (Domain 5 Record 13). There are also updates to Domain 1 Record 14. If you do not want to collect CPU-Measurement Facility counter data, use MONITOR SAMPLE DISABLE PROCESSOR CPUMFC to disable it, or unauthorize basic counters in the partition's activation profile.

Counters can also be enabled (through MONITOR SAMPLE ENABLE) or disabled (through MONITOR SAMPLE DISABLE) while the monitor is active. Additionally, authorization for counter sets can be changed dynamically. z/VM recognizes and responds to these dynamic authorization changes.

When a processor is brought online by VARY PROCESSOR, its counters are enabled and activated if collection is enabled.

[V6.2] Crypto Terminology Change

Crypto terminology in z/VM has changed to eliminate confusion over the use of the terms "queue" and "domain". Previously, both terms were used to refer to the same component of a crypto device. The responses for the QUERY CRYPTO and QUERY VIRTUAL CRYPTO commands that include the term "queue" have been changed to use the term "domain".

A new DOMAIN operand will be supported on the QUERY CRYPTO command. The existing APQS operand will continue to be supported as a synonym for domain to ease compatibility.

The response to QUERY CRYPTO DOMAIN (formerly QUERY CRYPTO APQS) now reports the hardware status, z/VM's use of the device, and the planned use of the device more clearly.

[V6.2] Cryptographic Coprocessor Facility (CCF) Support Removed

The IBM Z servers supported by z/VM provide the following cryptographic hardware features: CP Assist for Cryptographic Function (CPACF), Crypto Express2 feature, and Crypto Express3 feature. Because the old Cryptographic Coprocessor Facility (CCF) and its predecessors are no longer available on these processors, CP support for old cryptographic hardware has been removed.

The following monitor records are no longer available:

- Domain 5 Record 6 - MRPRCCFN - Vary On Crypto Facility Event Data
- Domain 5 Record 7 - MRPRCCFF - Vary Off Crypto Facility Event Data

[V6.2] Dynamic Discovery of FCP Disks

Dynamic discovery of FCP disks formalizes the SCSIDISC sample utility to be a fully supported z/VM tool. SCSIDISC provides CMS users with the ability to recognize any worldwide port numbers (WWPNs) and associated logical unit numbers (LUNs) accessible from some specified FCP subchannel. The use of this DASD in both emulated devices and directly attached FCP devices requires specific knowledge of the storage area networking topology. The SCSIDISC utility provides a means to analyze this topology with no prior knowledge of it.

The SCSIDISC utility has been added for this support.

[V6.2] ISFC Infrastructure Enhancements

The Inter-System Facility for Communications (ISFC) SSI infrastructure provides tools that will be used for cross-system communication. This enhances the ISFC subsystem to improve the transport mechanism and provide convenient interfaces for exploitation by other subsystems with the CP nucleus.

The following command has been added for this support:
- QUERY ISFC

The following commands have been updated for this support:
- ACTIVATE ISLINK
- DEACTIVE CONV
- DEACTIVE ISLINK
- MONITOR EVENT
- MONITOR SAMPLE
- QUERY COLLECT
- QUERY ISLINK
- QUERY MONITOR
- SET CPTRACE

[V6.2] Real Device Mapping

Real device mapping provides a means of identifying a device either by a customer-generated equivalency ID (EQID) or by a CP-generated EQID. This mapping is used to ensure virtual machines relocated by live guest relocation continue to use the same or equivalent devices following a relocation.

The following CP command has been added for this support:
- QUERY EQID

The following CP commands have been updated for this support:
- ATTACH
- DETACH
- QUERY EDEVICE
- QUERY (Real Device)
- SET EDEVICE
- SET RDEVICE (Advanced Function Printers)
- SET RDEVICE (Card Punches)
- SET RDEVICE (Card Readers

- SET RDEVICE (Clear)
- SET RDEVICE (Communication Controllers)
- SET RDEVICE (DASD)
- SET RDEVICE (Graphic Display Devices)
- SET RDEVICE (Impact Printers)
- SET RDEVICE (Integrated Communication Adapters)
- SET RDEVICE (Special Devices)
- SET RDEVICE (Tape Units)
- SET RDEVICE (Terminals)
- SET RDEVICE (Unsupported Devices)
- SET RDEVICE (3800 Printers)
- VARY (Real Device)

The following CP configuration statements have been updated for this support:

- EDEVICE Statement
- RDEVICE Statement (Advanced Function Printers)
- RDEVICE Statement (Card Punches)
- RDEVICE Statement (Card Readers)
- RDEVICE Statement (Communication Controllers)
- RDEVICE Statement (DASD)
- RDEVICE Statement (Graphic Display Devices)
- RDEVICE Statement (Impact Printers)
- RDEVICE Statement (Special Devices)
- RDEVICE Statement (Tape Units)
- RDEVICE Statement (Terminals)
- RDEVICE Statement (Unsupported Devices)
- RDEVICE Statement (3800 Printers)

The following CP monitor records have been updated to add the device EQID:

- Domain 1 Record 6 - MRMTRDEV - Device Configuration Data
- Domain 6 Record 1 - MRIODVON - Vary On Device - Event Data

[V6.2] Shared Disk Enhancements

The shared disk enhancements improve the support for sharing real DASD among z/VM images and simplifies the management of minidisk links and minidisk cache for minidisks shared by multiple images.

The following commands have been updated for this support:

- ATTACH
- CPACCESS
- DEFINE (General)
- DEFINE CPOWNED
- LINK
- QUERY CPOWNED
- QUERY LINKS
- QUERY MDCACHE
- QUERY MDISK

- QUERY SYSTEM
- RESET
- SET MDCACHE
- SET WRKALLEG

The following utility has been updated for this support:

- CPFMTXA

The following statements have been updated for this support:

- CP_OWNED configuration statement
- DASDOPT directory statement
- LINK directory statement
- MDISK directory statement
- REDEVICE configuration statement (DASD)

The following monitor records have been added:

- Domain 6 (I/O) Record 31 - MRIODMDE - Minidisk Activity
  - This record was added to indicate when the minidisk cache (MDC) setting has been changed automatically by CP.
- Domain 11 (SSI) Record 6 - MRSSIXLK - XDISK Serialization Sample
  - This record was added to indicate usage information for XDISK cross system serialization routines.
- Domain 11 (SSI) Record 7 - MRSSIXDI - XDISK Activity
  - This record was added to provide information about XDISK activity.

The following monitor records have been updated:

- Domain 0 Record 14 - MRSYTXSG - Expanded Storage Data (global)
  - SYTXSG_TCMCPURG - This field accumulates the number of times cache is purged due to a write link obtained in the SSI environment.
- Domain 4 Record 2 - MRUSELOF - User Logoff Data - Event Record

  The following fields contain LINK, DETACH and WRKALLEG usage information in an SSI environment:
  - USELOF_VMDLKTOT
  - USELOF_VMDLKPLX
  - USELOF_VMDLKTOD
  - USELOF_VMDDTTOT
  - USELOF_VMDDTPLX
  - USELOF_VMDDTTOD
  - USELOF_VMDWKTOT
  - USELOF_VMDWKPLX
  - USELOF_VMDWKTOD
- Domain 4 Record 3 - MRUSEACT - User Activity Data

  The following fields contain LINK, DETACH and WRKALLEG usage information in an SSI environment:
  - USEACT_VMDLKTOT
  - USEACT_VMDLKPLX
  - USEACT_VMDLKTOD
  - USEACT_VMDDTTOT
  - USEACT_VMDDTPLX
  - USEACT_VMDDTTOD
  - USEACT_VMDWKTOT
  - USEACT_VMDWKPLX

– USEACT_VMDWKTOD

[V6.2] FICON Express8S and OSA-Express4S

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

- For the FICON/zHPF/FCP storage area network:

  FICON Express8S for single mode fiber (LX) and multimode fiber (SX) environments, with two channels per feature and two channel path identifiers (CHPIDs).
- For the 10 Gigabit Ethernet and Gigabit Ethernet local area network:

  OSA-Express4S for single mode fiber (LX, LR) and multimode fiber (SX, LR) environments. The 10 GbE features have one port per feature and one CHPID. The GbE features have two ports per feature and one CHPID shared by the two ports.

[V6.2] Support for Gen 4 IBM System Storage TS1140 Tape Drive (3592 Model E07)

z/VM supports the Gen 4 IBM System Storage TS1140 Tape Drive (machine type 3592, model E07).

The following CP commands have been updated:

- QUERY TAPES
- QUERY (Virtual Device)

In addition, DIAGNOSE code X'210' has been updated to support a new VRDCUNDV value for the 3592 E07. For more information, see *z/VM: CP Programming Services*.

[V6.2] Guest Support for High Performance FICON for z Systems

With the PTF for APAR VM65041, z/VM V6.2 provides guest support for High Performance FICON for z Systems™ (zHPF). zHPF is a performance and reliability, availability, and serviceability (RAS) enhancement of the z/Architecture and the FICON channel architecture implemented in the zEnterprise 196, zEnterprise 114, and System z10® servers, and the System Storage DS8000 series.

Exploitation of zHPF by the FICON channel, the operating system, and the control unit is designed to help reduce the FICON channel overhead. When both the FICON channel and the control unit indicate support for zHPF, transport mode I/O (multiple commands per frame) is supported in addition to command mode I/O (one command per frame). The FICON Express8S, FICON Express8, FICON Express4, and FICON Express2 features support transport mode, and these features support the FICON architecture, FICON channel-to-channel (CTC), and the zHPF architecture simultaneously.

Use of zHPF requires the function to be installed on the machine (processor complex and channels) and the storage controller, and requires support in z/VM and the exploiting guest. For information about the specific hardware requirements for zHPF support, see *z/VM: General Information*.

The following CP commands have been updated with new and changed responses:

- QUERY CHPID
- QUERY MDCACHE
- QUERY PATHS
- TRACE IO
- TRSOURCE
- VARY PATH

The following CP monitor record has been added for this support:

- Domain 6 Record 32 - MRIODHPF - zHPF Feature Change

The following CP monitor records have been updated for this support:

- Domain 0 Record 19 - MRSYTSYG - System Data
- Domain 0 Record 20 - MRSYTEPM - Extended Channel Path Measurement Data
- Domain 1 Record 4 - MRMTRSYS - System Configuration Data
- Domain 6 Record 3 - MRIODDEV - Device Activity

With the PTF for APAR VM65044, the z/VM V6.2 Performance Toolkit feature recognizes the new monitor record for High Performance FICON.

[V6.3] IBM zEnterprise EC12 and IBM zEnterprise BC12

z/VM supports the IBM zEnterprise EC12 (zEC12) and the IBM zEnterprise BC12 (zBC12) and enables guests to exploit selected new features.

z/VM support for the zEC12 and zBC12 includes support for:

- **Local-TLB-Clearing Facility**

  The Local-TLB-Clearing Facility provides guests running on z/VM the ability to use the IPTE or IDTE instructions with the Local-Clearing Control (LC) bit on. Guests that are able to meet the architecture requirements and change their IPTE/IDTE instructions might see performance improvements.

- **Crypto Express4S**

  z/VM provides guest support for the Crypto Express4S feature, which is exclusive to the zEC12 and zBC12, for z/Architecture guests.

  The Crypto Express4S can be configured in one of three ways using the Hardware Management Console (HMC) panels:

  – 2877 IBM Common Cryptographic Architecture (CCA) coprocessor
  – 2878 IBM Enterprise Public-Key Cryptography Standards (PKCS) #11 (EP11) coprocessor
  – 2879 Accelerator

  The new EP11 coprocessor is designed to provide open industry standard cryptographic services. It provides enhanced firmware capabilities that are designed to meet the rigorous FIPS 140-2 Level 4 and Common Criteria EAL 4+ certifications.

  EP11 provides enhanced security qualities when using PKCS #11 functions. EP11 supports secure PKCS #11 keys that never leave the secure boundary of the coprocessor unencrypted. With EP11, keys can now be generated and securely wrapped under the EP11 Master Key, all within the bounds of the coprocessor.

  z/Architecture guests can be authorized in the z/VM directory for shared or dedicated access to the Crypto Express4S. When Crypto Express4S is configured as an accelerator or a CCA coprocessor, z/VM supports the device for shared or

dedicated use. When Crypto Express4S is configured as an EP11 coprocessor, z/VM supports the device for dedicated use only.

- **CPU-Measurement Counter Facility enhancement**

  The CPU-Measurement Counter Facility has been expanded. Code was added to z/VM to support current variations of the new version numbers when collecting the counters and when putting them into the monitor record (Domain 5 Record 13).

- **New subchannel measurement block fields**

  The Interrupt-Delay Time and I/O-Priority-Delay Time fields have been added to the subchannel measurement block.

- **OSA-Express5S**

  OSA-Express5S is the next generation of devices in the OSA family, and is exclusive to the zEC12 and zBC12. It contains a technology update compared with the OSA-Express4S, with similar performance characteristics.

The following CP commands have been updated for this support:
- QUERY CRYPTO
- QUERY VIRTUAL CRYPTO

The following CP monitor records have been updated:
- Domain 5 Record 10 - MRPRCAPM - Crypto Performance Measurement Data
- Domain 5 Record 13 - MRPRCMFC - CPU-Measurement Facility Counters
- Domain 6 Record 3 - MRIODDEV - Device Activity
- Domain 6 Record 14 - MRIODMOF - Set subchannel measurement off - Event

The TCP/IP NETSTAT OSAINFO command has been updated to display information for the OSA-Express5S.

[V6.3] Fibre Channel Protocol Data Router Support

z/VM FCP QEBSM support has been updated to support guest exploitation of the Data Router facility for direct memory access between an FCP device SCSI interface card and real memory.

The QUERY VIRTUAL FCP command has been updated for this support.

The following CP monitor record has been updated:
- Domain 1 Record 19 - MRMTRQDC - QDIO Device Configuration Record

[V6.3] FICON DS8000 Series New Functions

z/VM supports the following functions on IBM System Storage DS8700 and later devices:

- Storage Controller Health Message

  This new attention message will be generated from the hardware, one per corresponding Logical Storage System (LSS), to alert the operating system of a condition that in the past would have surfaced as a general Equipment Check. This message will give more details and is intended to reduce the number of false HyperSwap events that have occurred with the less descriptive Equipment Check.

- Peer-to-Peer Remote Copy (PPRC) Summary Unit Check

  When a PPRC suspend condition existed in the past, a corresponding state change interrupt would be presented to each affected DASD in the LSS. When there are a large number of DASD involved, the amount of processing, time, and

memory needed by the operating system to process each state change interrupt can accumulate. For customers with large numbers of DASD, this state change processing can cause timeouts to occur in Geographically Dispersed Parallel Sysplex (GDPS) HyperSwap scenarios, because a PPRC suspend occurs on every DASD under HyperSwap control. To prevent these timeouts, a new Summary Unit Check has been created that gives only one interrupt per affected LSS.

This support is intended to reduce the amount of false HyperSwap events by GDPS and reduce the amount of processing required to handle PPRC suspend events, which previously came via state change interrupts. This is particularly useful for customers with large sets of DASD under GDPS PPRC control.

The following CP functions have been updated:
- CU system configuration statement
- QUERY CU command
- QUERY HYPERSWAP command

[V6.3] HiperSocket Completion Queue Guest Exploitation Support

z/VM provides guest support for asynchronous HiperSockets data transmission provided by completion queues when deploying a HiperSockets network in a virtualized environment. Completion queues perform requests asynchronously to improve performance for peak workload intervals. IBM is working with its Linux distribution partners to include support in future Linux for z Systems distribution releases.

The following CP monitor records have been updated:
- Domain 0 Record 20 - MRSYTEPM - Extended Channel Path Measurement Data (per channel)
- Domain 1 Record 19 - MRMTRQDC - QDIO Device Configuration
- Domain 6 Record 25 - MRIODQDA - QDIO Device Activation Event
- Domain 6 Record 26 - MRIODQDS - QDIO Device Activity Sample
- Domain 6 Record 27 - MRIODQDD - QDIO Device Deactivation Event

[V6.3] Multiple Subchannel Set Support for Mirrored DASD

The multiple subchannel set support for mirrored DASD provides a subset of host support for the Multiple Subchannel Set (MSS) facility to allow using an alternate subchannel set for Peer-to-Peer Remote Copy (PPRC) secondary volumes.

The following CP command has been added for this support:
- QUERY MSS

The following CP commands have been updated for this support:
- HYPERSWAP
- QUERY DASD
- QUERY HYPERSWAP
- QUERY PATHS
- SET IPLPARMS
- SHUTDOWN
- VARY PATH
- VARY (Real Device)
- VARY SUBCHANNEL

The following CP utility has been updated for this support:
- SALIPL

The Stand-Alone Program Loader (SAPL) provides an expanded field on its main panel where a 5-digit device number for the minidisk on which the CP nucleus resides is displayed and can be specified. For more information about SAPL, see *z/VM: System Operation*.

You should not use the Stand-Alone Program Loader Creation Utility (SALIPL) to install SAPL if you specify the DEVICE option on the SALIPL utility because you will lose the information from any DEVICE option that you specified using SALIPL. For this function, you must use the SALIPL utility under CMS.

Different values will now appear in monitor record fields for devices in subchannel sets other than zero.

The following monitor records all use RDEVSID to set the subchannel ID values that they contain. With MSS support, the first halfword of this value does not have to be X'0001' because other subchannel sets can be used:

Domain 1 (Monitor)
- Record 6 - MTRDEV - Device Configuration Data
- Record 8 - MTRPAG - Paging Configuration Data

Domain 2 (Scheduler)
- Record 1 - SCLRDB - Begin Read
- Record 2 - SCLRDC - Read Complete
- Record 3 - SCLWRR - Write Response

Domain 3 (Storage)
- Record 4 - STOASP - Auxiliary Storage Management (per exposure)
- Record 7 - STOATC - Page/Spool Area of a CP Volume
- Record 11 - STOASS - Auxiliary Shared Storage Management (per explosure)

Domain 6 (I/O)
- Record 1 - IODVON - Vary Device On
- Record 2 - IODVOF - Vary Device Off
- Record 3 - IODDEV - Device Activity
- Record 4 - IODCAD - Cache Activity Data
- Record 5 - IODATD - Attach Device
- Record 6 - IODDTD - Detach Device
- Record 7 - IODENB - Enable Terminal
- Record 8 - IODDSB - Disable Terminal
- Record 9 - IODATS - Attach Shared Device
- Record 11 - IODSON - Vary On Subchannel
- Record 12 - IODSOF - Vary Off Subchannel
- Record 13 - IODMON - Set Subchannel Measurement On
- Record 14 - IODMOF - Set Subchannel measurement Off
- Record 15 - IODDDV - Delete Device
- Record 16 - IODMDV - Modify Device
- Record 18 - IODTON - Set Throttle On
- Record 19 - IODTOF - Set Throttle Off
- Record 20 - IODSTC - State Change
- Record 21 - IODVSW - Virtual Switch Activity

Domain 7 (Seek)
- Record 1 - SEKSEK - Seek Data

[V6.3] z/VM HiperDispatch

z/VM HiperDispatch is intended to improve efficiency in the use of CPU resources. This enhancement can help to improve performance and allow you to get more business value from your existing IBM Z servers, perhaps by consolidating z/VM workloads into fewer (larger) z/VM LPARs.

HiperDispatch improves CPU efficiency by causing the Control Program to run work in a manner that recognizes and exploits IBM Z machine topology to increase the effectiveness of physical machine memory cache. This includes:

- Requesting the LPAR hypervisor 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

HiperDispatch can also improve CPU efficiency by automatically tuning the LPAR's use of its logical CPUs to try to reduce multiprocessor effects. This includes:

- Sensing and forecasting key indicators of workload intensity and of elevated multiprocessor effect
- Autonomically tuning the z/VM system to reduce multiprocessor effect when it is determined that HiperDispatch can help to improve CPU efficiency

The supported processors limit for z/VM V6.3 remains 32, but with improved multiprocessor efficiency.

Externals for HiperDispatch support include:

- Changes to the following CP commands:
  - DEDICATE
  - INDICATE LOAD
  - INDICATE QUEUES
  - QUERY PROCESSORS
  - QUERY SRM
  - SET CPTRACE
  - SET SRM
- A new configuration statement: SRM.
- A change to the meaning of NOLIMIT on the FEATURES MAXUSERS configuration statement.
- New or changed monitor records:
  - D0R2 - MRSYTPRP - Processor Data (Per Processor)
  - D0R16 - MRSYTCUP - CPU Utilization in a Logical Partitition
  - D0R23 - MRSYTLCK - Formal Spin Lock Data
  - D0R24 - MRSYTSPT - Scheduler Activity (per processor type)
  - D1R4 - MRMTRSYS - System Configuration Data
  - D1R5 - MRMTRPRP - Processor Configuration
  - D1R16 - MRMTRSCH - Scheduler Settings
  - D2R7 - MRSCLSRM - SET SRM Changes
  - D4R2 - MRUSELOF - User Logoff event record
  - D4R3 - MRUSEACT - User Activity Data
  - D5R2 - MRPRCVOF - Vary Off Processor
  - D5R3 - MRPRCPRP - Processor Data

  - D5R15 - MRPRCDSV - Dispatch Vector Assignments (new event record)
  - D5R16 - MRPRCPUP - Park/Unpark Decision (new event record)
  - D5R17 - MRPRCRCD - Real CPU Data (per CPU) (new sample record)
  - D5R18 - MRPRCDHF - Dispatch Vector High Frequency Data (new sample record).
- A change to the VM Dump Tool CPUUSE macro.
- New trace table entries:
  - Input unpark mask - 3610
  - Switch master - 3611
  - Park processor - 3612
  - Unpark processor - 3613
  - SIGP instruction AE01 replaces AE00
  - Time stamp trace entry FFFD replaces FFFE.

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

[V6.3 APAR] Soft Fence and Query Host Access GDPS Enhancements

With the PTF for APAR VM65322, z/VM supports the new Soft Fence and Query Host Access DS8000 functions.

Soft Fence is designed to protect the integrity of a volume in the fence state by limiting most I/O operations directed to that volume. This function is exploited in GDPS 3.10 and is used to fence the old primary volumes in the GDPS-managed consistency group during a HyperSwap scenario. GDPS automatically resets the fence state upon completion of a HyperSwap scenario. z/VM offers support to manually reset and query the soft fence state of a volume or the volumes on a given controller. Additionally, the Query Host Access function reports all LPARs that have established a path group ID on a volume, whether or not in the grouped state. z/VM now offers support to query the CPU serial number and LPAR ID associated with all LPARs that have established a path group ID to a volume.

The following CP command has been added for this support:
- SET DASD

The following CP commands have been updated for this support:
- QUERY CU
- QUERY DASD
- QUERY FENCES
- SET CU

[V6.3 APAR] PCIe Guest Direct Attach Support

With the PTFs for APARs VM65417 (CP), VM65572 (CP TRACERED support), and PI20509 (TCP/IP), z/VM V6.3 guests will be allowed to dedicate PCIe adapters to their virtual machines.

z/OS V2.1 guests (with required PTFs) running under z/VM V6.3 can exploit the zEDC Express and 10GbE RoCE features. See the PSP bucket for specific details.

- zEDC Express feature

  IBM zEnterprise Data Compression (zEDC) for z/OS V2.1 and the zEDC Express feature are designed to support a new data compression function to help provide high-performance, low-latency compression without significant CPU overhead. This can help to reduce disk usage, provide optimized cross-platform exchange of data, and provide higher write rates for SMF data.

- 10GbE RoCE Express feature

  The term RoCE refers to Remote Direct Memory Access over Converged Ethernet. The 10 Gigabit Ethernet (10GbE) RoCE Express feature is designed to help 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 is designed to help reduce network latency with memory-to-memory transfers utilizing Shared Memory Communications-Remote Direct Memory Access (SMC-R) in z/OS V2.1. 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.

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.

For more information about guest support for PCIe functions, see *z/VM: CP Planning and Administration*.

The following system configuration statements have been updated for this support:
- FEATURES
- STORAGE

The following CP commands have been added for this support:
- DEFINE PCIFUNCTION
- DELETE PCIFUNCTION
- DETACH PCIFUNCTION
- LOCATE RPCI
- LOCATE VPCI
- MODIFY PCIFUNCTION
- QUERY IOASSIST (The function of this command is similar to but different from the QUERY IOASSIST command that was removed in V5.1.)
- QUERY PCIFUNCTION
- QUERY VIRTUAL PCIFUNCTION
- SET IOASSIST (The function of this command is similar to but different from the SET IOASSIST command that was removed in V5.1.)
- VARY PCIFUNCTION

The following CP commands have been updated for this support:
- ATTACH
- DEFINE CHPID / PATH
- MONITOR
- QUERY CHPID
- QUERY FRAMES
- QUERY MONITOR
- QUERY TRFILES
- QUERY TRSOURCE
- RESET
- SET CPTRACE
- TRSOURCE ENABLE
- TRSOURCE ID
- VMRELOCATE

PCIFUNCT has been added to the list of restricted user IDs.

The following CP monitor records have been added for this support:

- Domain 1 Record 27 - MRMTRPCI - PCI function Configuration Data
- Domain 6 Record 36 - MRIODPAT - Attach PCI Function
- Domain 6 Record 37 - MRIODPDT - Detach PCI Function
- Domain 6 Record 38 - MRIODPEN - Guest Enables a PCI Function
- Domain 6 Record 39 - MRIODPAC - PCI Activity
- Domain 6 Record 40 - MRIODPDS - Guest Disables a PCI Function
- Domain 6 Record 41 - MRIODPER - PCI function error
- Domain 6 Record 42 - MRIODPAD - PCI function added to the system
- Domain 6 Record 43 - MRIODPDL - PCI function deleted from the system
- Domain 6 Record 44 - MRIODPMD - PCI function program controls modified
- Domain 6 Record 45 - MRIODPON - Real PCI function varied on
- Domain 6 Record 46 - MRIODPOF - Real PCI function varied offline

The following CP monitor records have been updated:
- Domain 0 Record 3 - MRSYTRSG - Real Storage Data (global)
- Domain 1 Record 4 - MRMTRSYS - System Configuration Data
- Domain 1 Record 7 - MRMTRMEM - Memory Configuration Data

The following DIAGNOSE code has been updated for this support:
- DIAGNOSE Code X'2AC'

Trace codes in the range 1200 - 123F have been added for this support. For more information about these trace codes, see *z/VM: Diagnosis Guide*.

The following dynamic I/O return codes have been modified for adding, deleting, or changing a PCI function:
- 0106
- 0107
- 010A
- 010D
- 010F
- 0110
- 0112
- 011C
- 011D

[V6.3 APAR] IBM z13 Support

With the PTF for APAR VM65577, z/VM provides support that will enable guests to exploit zEC12 function supported by z/VM on the IBM z13™ (z13™). z/VM support for the z13 also includes support for:

- New hardware facilities

  z/VM supports the following new hardware facilities transparently.

  – Load/Store-on-condition Facility 2
  – Load-and-Zero-Rightmost-Byte Facility
  – Decimal-Floating Point Packed Conversion Facility
  – Delay Facility

  New facility bits have been defined that will be passed to a guest and can be tested to determine if a hardware facility is available. A guest can use the STFLE instruction to obtain the facilities list. These new facility bits are also incorporated into the live guest relocation domains support.

  TRACE, DISPLAY I, and VMDUMPTL support for interpreting the new instructions associated with these facilities for display purposes is not included.
- Updates to monitor records for CPU-Measurement Counter Facility

The new Counter Second Version Number value (4) is supported, and appropriate counters are collected and stored in the CPU-Measurement Facility Counters monitor record (Domain 5 Record 13).

- CPU-Measurement Counter Facility enhancement

  Host exploitation is provided for a new instruction to allow collection of multiple counters simultaneously, when the store-CPU-counter-multiple facility is installed. This is expected to reduce the overhead for collecting CPUMF data for z/VM monitor records. Guest support is not provided for the new store-CPU-counter-multiple facility or the new MT-CPU-timer-enhancement facility.
- New I/O related architectures

  See "[V6.3 APAR] z13 Compatibility I/O Support" on page 106.
- Crypto Express5S and enhanced domain support

  See "[V6.3 APAR] Crypto Express5S and Enhanced Domain Support" on page 106.
- FICON Express16S

  See "[V6.3] FICON Express16S Support" on page 107.

Note: The IBM zEnterprise Application Assist Processor (zAAP) CPU type is not supported on z13, either in real hardware or in the LPAR configuration. Similarly, z/VM will not allow guests to define this CPU type if the machine does not support it.

The following CP commands have been updated:
- DEFINE CPU
- QUERY CAPABILITY

A new CPU Capability continuation data (Type E) accounting record has been added.

The following monitor records have been updated:
- Domain 0 Record 19 - MRSYTSYG - System Data - will now report two values for capability.
- Domain 1 Record 4 - MRMTRSYS - System Configuration Data - will now report two values for capability.
- Domain 1 Record 18 - MRMTRCCC - CPU Capability Change - will now report two values for capability.
- Domain 5 Record 13 - MRPRCMFC - CPU-Measurement Facility Counters - will now be able to report a Counter Second Version Number value of 4.

For more information, see:
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Planning and Administration*

See the following additional z13 support:
- "[V6.3 APAR] Simultaneous Multithreading (SMT) Support" on page 108
- "[V6.3 APAR] Increased CPU Scalability" on page 111
- "[V6.3 APAR] Multi-VSwitch Link Aggregation Support" on page 155

Support for the z13 also requires the PTFs for the following APARs. For more information, see the documentation for the product, feature, or function.
- APAR VM65495 provides EREP/VM support.

- APAR VM65568 provides IOCP support.
- APAR VM65489 provides HCD support.
- APAR VM64437 provides HCM support.
- APAR PM79901 provides HLASM support.
- APAR VM65527 provides Performance Toolkit support.

[V6.3 APAR] z13 Compatibility I/O Support

With the PTF for APAR VM65577, z/VM supports new I/O related architectures and features of the z13, including:

- PCI function measurement block enhancements for the RDMA over Converged Ethernet (RoCE) adapter
- Dynamic I/O support for new channel path type CS5 (Coupling over PCIe)
- Dynamic I/O support for specifying virtual channel IDs (VCHIDs) for HiperSockets (IQD) channels

The following commands have been updated for this support:

- DEFINE CHPID/PATH
- DEFINE PCIFUNCTION
- QUERY CPUID
- VARY PCIFUNCTION

The following CP monitor records have been updated:

- Domain 0 Record 20 - MRSYTEPM - Extended Channel Path Measurement Data (per channel)
- Domain 1 Record 27 - MRMTRPCI - PCI Function Configuration Data
- Domain 6 Record 39 - MRIODPAC - PCI Activity
- Domain 6 Record 40 - MRIODPDS - Guest Disables a PCI Function
- Domain 6 Record 42 - MRIODPAD - PCI Function Added to the System
- Domain 6 Record 45 - MRIODPON - Real PCI Function Varied On

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Messages and Codes*
- *z/VM: I/O Configuration*

[V6.3 APAR] Crypto Express5S and Enhanced Domain Support

With the PTF for APAR VM65577, z/VM provides z/Architecture guest support for the Crypto Express5S feature, which is exclusive to the z13, and enhanced domain support for Crypto Express5S.

Crypto Express5S can be configured in one of three ways using the Hardware Management Console (HMC) panels:

- IBM Common Cryptographic Architecture (CCA) coprocessor
- IBM Enterprise Public Key Cryptography Standards (PKCS) #11 (EP11) coprocessor
- Accelerator

z/Architecture guests can be authorized in the z/VM directory for shared or dedicated access to the Crypto Express5S. When Crypto Express5S is configured as an accelerator or a CCA coprocessor, z/VM supports the device for shared or dedicated use. When Crypto Express5S is configured as an EP11 coprocessor, z/VM supports the device for dedicated use only.

In addition, z/VM supports enhanced domain support for the Crypto Express5S feature on the z13:

- z/VM supports an architected increase in the maximum number of crypto features, referred to as Adjunct Processors (APs), from 64 to 256, and an architected increase in the maximum number of domains per AP from 16 to 256.
- The z13 supports up to 16 APs, and up to 85 domains on each AP, with the Crypto Express5S feature.

Attention: This support provides a new version of DIRECTXA, which requires more directory space on the DASD volume that contains the object directory. An additional 60 bytes per user plus 2 pages for system use is required. The count of users should include each user defined with a USER, IDENTITY, or POOL statement.

With APAR VM65577 applied, z/VM supports Message-Security-Assist Extension 5 on zEC12 (Driver D15F) and later.

With the PTF for APAR VM65588, DirMaint supports enhanced domain support for the Crypto Express5S feature on z13.

The following configuration statement is new:
- CRYPTO APVIRTUAL

The following directory statement was updated:
- CRYPTO

The following commands have been updated for this support:
- QUERY CRYPTO
- QUERY VIRTUAL CRYPTO
- DIRMAINT CRYPTO

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: Directory Maintenance Facility Commands Reference*
- *z/VM: CP Planning and Administration*

[V6.3] FICON Express16S Support

z/VM supports the FICON Express16S feature on z13. With the introduction of FICON Express16S on the z13, you now have additional growth opportunities for your storage area network (SAN). FICON Express16S supports a link data rate of 16 gigabits per second (Gbps) and autonegotiation to 4 or 8 Gbps for synergy with existing switches, directors, and storage devices. With support for native FICON, High Performance FICON for z Systems (zHPF), and Fibre Channel Protocol (FCP), the z13 server enables you to position your SAN for even higher performance, helping you to prepare for an end-to-end 16 Gbps infrastructure to meet the increased bandwidth demands of your applications.

The new FICON Express16S channel will work with your existing fiber optic cabling environment, both single mode and multimode optical cables. The FICON Express16S feature running at end-to-end 16 Gbps link speeds will provide reduced latency for large read/write operations and increased bandwidth compared with the FICON Express8S feature.

[V6.3 APAR] Simultaneous Multithreading (SMT) Support

Simultaneous multithreading (SMT) technology is available on z13 and is supported by z/VM.

z13 SMT support

Incremental throughput improvements are achieved on z13 partly because the new processor chip offers intelligently implemented 2-way simultaneous multithreading (SMT). SMT allows two active instruction streams per core, each dynamically sharing the core's execution resources. SMT will be available on z13 for workloads running on the IBM Integrated Facility for Linux (IFL) and the IBM z Integrated Information Processor (zIIP).

Each software operating system or hypervisor has the ability to intelligently drive SMT in a way that is best for its unique requirements. z/OS SMT management consistently drives the cores to high thread density, in an effort to reduce SMT variability and deliver repeatable performance across varying CPU utilization, thus providing more predictable SMT capacity. z/VM SMT management optimizes throughput by spreading a workload over the available cores until it demands the additional SMT capacity.

z13 z/VM SMT support

With the PTFs for APARs VM65586 and VM65696, z/VM provides host exploitation support for SMT on z13, which will enable 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 z13 to dispatch work on an individual thread of a 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. For a performance analysis on how SMT benefited z/VM workloads, see IBM: VM Performance Resources (www.ibm.com/vm/perf/).

Simultaneous multithreading support is available on a z/VM system only if the facility is installed 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.

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.

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. Existing distributions of Linux for z Systems can benefit from SMT exploitation in a z/VM host.

Stand-alone dump support for SMT requires APARs VM65676 and VM65677.

Performance Toolkit support for SMT requires APAR VM65529. See "[V6.3 APAR] Additional Performance Toolkit Enhancements" on page 219.

Note: z/VM CPU pools provide a mechanism for limiting the CPU resources consumed by a group of virtual machines to a specific capacity. In an environment without SMT, these capacities are enforced in terms of a number of cores. In an SMT environment, these capacities are enforced in terms of a number of threads. Consequently, it might be necessary to increase the capacities of CPU pools, in order to provide adequate resource to CPU pool members.

Three Measures of CPU Time When Multithreading Is Enabled

z/VM host exploitation support for SMT provides three measures of CPU time when multithreading is enabled, because the hardware CPU timer is no longer an indication of core utilization. These three measures of CPU time are described below and are reported by accounting records and monitor records.

Raw Time
: This is a measure of the CPU time each virtual CPU spent dispatched on a thread, and is the CPU timer information provided directly by the hardware. When all cores have only one thread, this is an accurate measure of CPU time used by the task running on the single-threaded core. When multithreading is enabled, and some cores are running with more than one thread, the CPU Timer is no longer a direct indication of physical core consumption, so you might want one of the other times.

MT-1 Equivalent Time
: This is a measure of effective capacity, taking into account the multithreading benefit. The CPU time charged approximates the time that would have been spent if the workload had been run with multithreading disabled; that is, with all core resources available to one thread. The effect is to "discount" the time charged to compensate for the slowdown induced by the activity on other threads in the core.

Prorated Core Time
: This is a measure of core utilization regardless of the multithreading benefit. Time is charged by dividing the time the core was dispatched evenly among the threads dispatched in that interval. Under this method, the total time charged to all guests equals the total time the logical cores of the z/VM partition were dispatched. This method is consistent with cost recovery for core-based software licensing.

Notes:

1. When a user is running on a system where multithreading is not installed or not enabled, MT-1 equivalent time and prorated core time consumed will be identical to raw time.
2. The PTF for APAR VM65680 is required for prorated core time fields to be populated.

Changes to z/VM Interfaces

The following CP interfaces have been added:

- MULTITHREADING system configuration statement
- INDICATE MULTITHREAD command
- QUERY MULTITHREAD command
- VARY CORE command
- Virtual Machine Resource Usage 2 (Type F) accounting record

- Trace ID 3615 (Results of CPU Quiesce)
- Monitor record Domain 5 Record 20 - MRPRCMFM - MT CPUMF Counters

The following CP interfaces have been updated:
- SRM system configuration statement
- INDICATE LOAD command
- INDICATE USER command
- LOGOFF command
- MONITOR SAMPLE command
- QUERY TIME command
- QUERY PROCESSORS command
- SET CPTRACE command
- SET SRM command
- VARY PROCESSOR command
- Virtual Machine Resource Usage (Type 1) accounting record
- Store Hypervisor Information (STHYI) instruction
- Trace ID 3612 (Enter Parked Wait State)
- VM Dump Tool CPUUSE macro
- VM Dump Tool HCQGDSPL function
- Monitor records:
  - Domain 0 Record 2 - MRSYTPRP - Processor Data (per processor)
  - Domain 0 Record 15 - MRSYTCUG - Logical CPU Utilization Data (global)
  - Domain 0 Record 16 - MRSYTCUP - CPU Utilization in a Logical Partition
  - Domain 0 Record 17 - MRSYTCUM - Physical CPU Utilization Data for LPAR Management
  - Domain 0 Record 19 - MRSYTSYG - System Data (global)
  - Domain 0 Record 23 - MRSYTLCK - Formal Spin Lock Data (global)
  - Domain 1 Record 4 - MRMTRSYS - System Configuration Data
  - Domain 1 Record 5 - MRMTRPRP - Processor Configuration Data (per processor)
  - Domain 1 Record 16 - MRMTRSCH - Scheduler Settings - Sample Record
  - Domain 1 Record 18 - MRMTRCCC - CPU Capability Change
  - Domain 2 Record 4 - MRSCLADL - Add User To Dispatch List - Event Record
  - Domain 2 Record 5 - MRSCLDDL - Drop User From Dispatch List - Event Record
  - Domain 2 Record 7 - MRSCLSRM - SET SRM Changes - Event Record
  - Domain 2 Record 13 - MRSCLALL - Add VMDBK to the limit list - Event Record
  - Domain 2 Record 14 - MRSCLDLL - Drop VMDBK from the limit list - Event Record
  - Domain 4 Record 2 - MRUSELOF - User Logoff Data - Event Record
  - Domain 4 Record 3 - MRUSEACT - User Activity Data
  - Domain 4 Record 9 - MRUSEATE - User Activity data at Transaction End - Event Record
  - Domain 5 Record 1 - MRPRCVON - Vary On Processor - Event Data
  - Domain 5 Record 2 - MRPRCVOF - Vary Off Processor - Event Data
  - Domain 5 Record 11 - MRPRCINS - Instruction Counts (per processor)
  - Domain 5 Record 13 - MRPRCMFC - CPU-Measurement Facility Counters
  - Domain 5 Record 16 - MRPRCPUP - Park/Unpark Decision (Event)
  - Domain 5 Record 17 - MRPRCRCD - Real CPU Data (per CPU) (Sample)
  - Domain 5 Record 19 - MRPRCCPU - CPU Pool Utilization (Sample)

For more information, see:
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Planning and Administration*

- *z/VM: CP Programming Services*

[V6.3 APAR] Increased CPU Scalability

With the PTFs for APARs VM65586 and VM65696, z/VM will support up to 64 logical processors on z13:

- 64 cores with multithreading disabled
- 32 cores (up to 2 threads per core) with multithreading enabled

z/VM continues to support up to 32 logical processors on prior machines.

The following CP monitor records have been updated:
- Domain 0 Record 23 - MRSYTLCK - Formal Spin Lock Data (global)
- Domain 3 Record 1 - MRSTORSG - Real Storage Management (global)
- Domain 3 Record 2 - MRSTORSP - Real Storage Activity (per processor)

[V6.3 APAR] Multi-Target PPRC Support

With the PTFs for APARs VM65544 and VM65674, z/VM provides support for the Multi-Target Peer-to-Peer Remote Copy (PPRC) storage feature. This support exploits multiple PPRC secondary volumes (targets) for an associated primary volume when configured in default subchannel set 0. z/VM's alternate subchannel set support is tolerated with this new function, but z/VM does not support a multi-target secondary in the alternate subchannel set at this time.

The following CP commands have been updated for this support:
- HYPERSWAP
- QUERY DASD
- QUERY HYPERSWAP

Device Support Facilities (ICKDSF) APAR PM99490 is required to support Multi-Target PPRC DASD.

[V6.3 APAR] Multithreading Prorated Core Time Support

When the PTF for APAR VM65680 is applied and multithreading is enabled, CP provides support for multithreading prorated core time. Prorated core time is a method that provides reports and caps based on core utilization regardless of the multithreading benefit, which might be used as a means of distributing charges for subcapacity based products that are based on core capacity. Time is charged by dividing the time the core was dispatched evenly among the threads dispatched in that interval. Under this method, the total time charged to all guests equals the total time the logical cores of the z/VM partition were dispatched. Prorated core time is calculated for all users only when multithreading is enabled, and used in the limiting calculations of CPU pools and guests with individual LIMITHARD shares.

With this support, the prorated core time fields that were added to accounting records and monitor records for multithreading support will now be populated with data when multithreading is enabled.

The following CP interfaces have been updated:
- DEFINE CPUPOOL command
- QUERY CPUPOOL command
- SET CPUPOOL command
- SET SHARE command
- Virtual Machine Resource Usage 2 (Type F) accounting record

- Store hypervisor information (STHYI) instruction
- Monitor records:
  - Domain 2 Record 13 - MRSCLALL - Add VMDBK to the limit list - Event Record
  - Domain 2 Record 14 - MRSCLDLL - Drop VMDBK from the limit list - Event Record
  - Domain 5 Record 19 - MRPRCCPU - CPU Pool Utilization (Sample)

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Planning and Administration*
- *z/VM: CP Programming Services*

[V6.3 APAR] Support for IBM LinuxONE Systems

With the PTF for APAR VM65716, z/VM V6.3 can be selected as the hypervisor for IBM LinuxONE™ systems:

- IBM LinuxONE Emperor™ (based on the IBM z13 server)
- IBM LinuxONE Rockhopper™ (based on the IBM z13s™ server)

Additional z/VM and other APARs might be required to support these servers. For more information, see the server support table in *z/VM: General Information*.

[V6.3 APAR] IBM z13 (Driver D27) Compatibility and IBM z13s Support

With the PTF for APAR VM65716, z/VM provides support for z/VM V6.3 and V6.2 guests to exploit z/VM supported z13 functions on the z13 (Driver D27) and the IBM z13s.

This APAR also provides support for the following additional functions on z13 (Driver D27) and z13s™:

- **LPAR enhancement to provide group physical capacity limit enforcement**

  Processor Resource/System Manager (PR/SM) and the Hardware Management tool have been enhanced to support an option to limit the amount of physical processor capacity consumed by a group of logical partitions (LPARs) when a processor unit (PU) is defined as a general purpose processor (CP) or an Integrated Facility for Linux (IFL) shared across a set of LPARs.

  This enhancement is designed to provide a group physical capacity limit enforced as an absolute (versus relative) limit; it is not affected by changes to the logical or physical configuration of the system. This group physical capacity limit can be specified in units of CPs or IFLs.

  The z/VM support includes use of this information in functions that exploit or report processor capacity information for the purpose of adapting to its use. In z/VM publications this function might also be referred to as LPAR group absolute capacity capping.

- **LPAR enhancement for dynamic memory management**

  Processor Resource/Systems Manager (PR/SM) has been enhanced to support more flexibility as to how additional physical memory is dynamically added to a logical partition. Rather than attempting to fully populate a logical partition's reserved storage element when it is initially configured online, the operating system in the partition can request a single storage increment be attached (and subsequently can request additional increments if desired). This allows a more gradual, flexible addition of memory to the partition as needed over time.

- **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. Supported for z/VM V6.3 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 usage of the new Internal Shared Memory (ISM) PCI function type.

  See the z/OS subset within the 2964DEVICE or 2965DEVICE preventive service planning (PSP) bucket for z/OS service required in support of SMC-D connectivity.
- **Regional Crypto Enablement (RCE)**

  z/VM supports z/OS guest exploitation of IBM Regional Crypto Enablement (RCE) adapters. IBM will enable geo-specific cryptographic support that will be supplied by approved vendors. The RCE support will reserve the I/O slots for the vendor-supplied cryptographic cards. Clients will need to directly contact the approved vendor for purchasing information.

  Note: RCE support also requires the PTF for APAR VM65577.

The following dynamic I/O return code has been added for adding, deleting, and changing a PCI function:

- 0321

The following CP interfaces have been updated:

- DEFINE PCIFUNCTION command
- Store hypervisor information (STHYI) instruction

  The partition section of the function code X'0000' response buffer is updated to include information about the LPAR group absolute capacity caps if defined.
- Monitor records:
  - Domain 0 Record 16 - MRSYTCUP - CPU Utilization in a Logical Partition, updated to report the LPAR group name and associated per CPU type LPAR Group Absolute Capacity Value.
  - Domain 0 Record 23 - MRSYTLCK - Formal Spin Lock Data (global), updated to include DSVBK lock statistics.
  - Domain 1 Record 4 - MRMTRSYS - System Configuration Data, updated to include the STSI 1.2.2 CPU speeds.
  - Domain 1 Record 18 - MRMTRCCC - CPU Capability Change, updated to include the STSI 1.2.2 CPU speeds.
  - Domain 1 Record 27 - MRMTRPCI - PCI function Configuration Data.
  - Domain 5 Record 9 - MRPRCAPC - Crypto Performance Counters, updated to include the crypto delay times and the minimum and maximum sleep times.
  - Domain 5 Record 16 - MRPRCPUP - Park/Unpark Decision (Event), updated to report the per CPU type LPAR Group Absolute Capacity Value.
  - Domain 6 Record 39 - MRIODPAC - PCI Activity.
  - Domain 6 Record 40 - MRIODPDS - Guest Disables a PCI Function.

  - Domain 6 Record 42 - MRIODPAD - PCI function added to the system.
  - Domain 6 Record 45 - MRIODPON - Real PCI function varied on.

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Programming Services*
- *z/VM: I/O Configuration*

Support for z13 (Driver D27) and z13s also requires the PTFs for the following APARs (in addition to the base z13 support APARs listed under "[V6.3 APAR] IBM z13 Support" on page 104). For more information, see the documentation for the product, feature, or function.

- APAR VM65704 provides EREP/VM support.
- APAR VM65736 provides IOCP support, including support for LinuxONE.
- APAR VM65729 provides VMHCD support, including support for LinuxONE.
- APAR VM64844 provides VMHCM support, including support for LinuxONE..
- APAR VM65698 provides Performance Toolkit support for SMC-D.

[V6.3 APAR] IBM Z Vector Facility (SIMD) Support

With the PTF for APAR VM65733, z/VM enables guests to exploit the Vector Facility for z/Architecture. The Vector Facility instructions provide a powerful framework for development of new Business Analytics workloads, porting numerically intensive workloads from other platforms, and accelerating Business Analytics workloads on IBM z13.

This support enables guest use of Vector Facility instructions and the 128-bit vector registers used by these instructions. This support also enables command access to the additional floating-point (AFP) registers without requiring prior guest program use. As a result, a request to display or dump all floating-point registers (DISPLAY Y or DUMP Y) will always output 16 registers.

The documentation for the following commands has been updated for this support:

- DISPLAY
- DISPLAY (Registers)
- DUMP (Registers)
- SAVESYS
- STORE (Registers)
- STORE STATUS
- VMDUMP

The following CP monitor records have been updated:

- Domain 1 Record 15 - MRMTRUSR - Logged on User
- Domain 4 Record 2 - MRUSELOF - User Logoff Data - Event Record
- Domain 4 Record 3 - MRUSEACT - User Activity Data
- Domain 4 Record 9 - MRUSEATE - User Activity data at Transaction End - Event Record

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Messages and Codes*

[V6.4] 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 the server support information in *z/VM: General Information*.

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

[V6.4] ESA/390 Removal

z/VM V6.4 enhancements enable hypervisor initialization and termination and the Stand-Alone Program Loader (SAPL), DASD Dump Restore (DDR), and Stand-Alone Dump utilities to run entirely in z/Architecture mode. The IBM z13 and z13s are planned to be the last IBM Z 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. On those future systems, ESA/390 logical partition mode will be renamed to General logical partition 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.

You should reinstall the SAPL, DDR, and Stand-Alone Dump utilities so your system is ready to run on a z/Architecture-only machine. Also note that IPLable installation DVDs from previous releases will not IPL on z/Architecture-only machines.

The following directory statements have been updated:
- GLOBALOPTS
- MACHINE

The following CP commands have been updated:
- DEFINE CPU
- DEFSYS
- INDICATE USER
- IPL
- LOGON
- QUERY NSS
- QUERY PROCESSORS
- QUERY SET
- QUERY VCONFIG
- SAVESYS
- SET MACHINE
- SET VCONFIG
- SET 370ACCOM
- XAUTOLOG

The following CP utilities have been updated:
- DDR
- SALIPL
- SDINST

The following DIAGNOSE code has been updated:
- X'84'

The following monitor records have been updated:
- Domain 1 Record 1 - MRMTREPR - Event Profile
- Domain 4 Record 1 - MRUSELON - User Logon - Event Record
- Domain 4 Record 2 - MRUSELOF - User Logoff Data - Event Record
- Domain 4 Record 3 - MRUSEACT - User Activity Data
- Domain 4 Record 9 - MRUSEATE - User Activity Data at Transaction End - Event Record

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

ICKDSF support for z/Architecture-only servers requires APAR PI46151.

For more information, see:
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Planning and Administration*
- *z/VM: CP Programming Services*

[V6.3 APAR] ESA/390 Removal Support Available for z/VM V6.3

With the PTF for APAR VM65856, the V6.3 CP hypervisor (CPLOAD), Stand-Alone Program Loader (SAPL) utility, and DASD Dump Restore (DDR) utility are enhanced to run entirely in z/Architecture mode.

With the PTFs for APARs VM65921 (CP) and VM65922 (CMS), the V6.3 Stand-Alone Dump utility is enhanced to run entirely in z/Architecture mode.

Note: These V6.3 APARs do not include support for defining Z mode virtual machines.

You should regenerate and install the updated CPLOAD module and reinstall SAPL, DDR, and Stand-Alone Dump so your system is ready to run on a z/Architecture-only machine. Also note that IPLable installation DVDs from V6.3 and previous releases will not IPL on z/Architecture-only machines.

[V6.4] 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 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.

The CP TRACE *mnemonic1* command has been updated with the new PFMF mnemonic.

The following monitor record has been updated:
- Domain 5 Record 11 - MRPRCINS - Instruction Counts (per processor)

[V6.4] 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.

The following CP command has been added for this support:
- EXPLORE FCP

The following CP commands have been updated:
- QUERY DASD
- QUERY EDEVICE
- VARY (Real Device)

The following CP utility has been added for this support:
- IOEXPLOR

For more information, see:
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Messages and Codes*

[V6.4] SCSI Reliability, Availability, and Serviceability (RAS) 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 is now supported without quiescing EDEVICE I/O. This was previously restricted, and continues to be restricted on releases prior to z/VM V6.4.

The following CP functions have been updated for this support:
- The following CP commands have been updated:
  - SET CPTRACE
  - SET EDEVICE
  - SET MITIME
- The EDEVICE configuration statement has been updated.

- The following new trace codes have been added to the CP trace table: 2891, 28A1, 28A2, 2C91, 2CA1, 2CB1, and 6026–6030.
- The following trace codes have been deleted in the CP trace table: 2890, 28A0, 2C90, 2CA0 and 2CB0.

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Planning and Administration*
- *z/VM: CP Messages and Codes*
- *z/VM: Diagnosis Guide*

[V6.4] Guest Transactional Execution (TX) 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.

Note: The following instructions are restricted from use within transactions:
- Store Hypervisor Information (STHYI)
- Inter-User Communications Vehicle (IUCV)
- DIAGNOSE

The following CP commands have been updated:
- TRACE
- QUERY TRACE
- SET CPTRACE
- TRSOURCE

The following CP trace codes have been added: 8B00, 8B01, 8B02, 9A01.

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Messages and Codes*
- *z/VM: CP Programming Services*
- *z/VM: Virtual Machine Operation*
- *z/VM: Diagnosis Guide*.

[V6.4] IBM z13 (Driver D27) and z13s I/O Support

z/VM compatibility support for the IBM z13 (driver D27) and z13s includes support for the user-defined identifier (UID) field for PCI functions. 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.

The following CP commands have been added for this support:

- QUERY IO_OPT
- SET IO_OPT

The following CP commands have been updated for this support:

- ATTACH
- DEFINE PCIFUNCTION
- QUERY PCIFUNCTION

Additional program requirements for UID support:

- CP support on V6.2, V6.3, and V6.4 requires APAR VM65865.
- 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.)

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Messages and Codes*
- *z/VM: I/O Configuration*

[V6.4] Expanded Storage (XSTORE) Support Removed

IBM z13 and z13s are the last IBM Z servers to support expanded storage (XSTORE). z/VM V6.4 does not support XSTORE for either host or guest usage.

The XSTORE user directory statement is supported for compatibility only. The statement has no function and has been removed from the documentation. If a directory entry includes the XSTORE statement, message HCP1401I will be issued when the guest logs on to state that XSTORE is not supported.

The following CP commands have been removed:

- DETACH XSTORE
- QUERY VIRTUAL XSTORE
- QUERY XSTORE
- RETAIN XSTORE

The following CP commands have been updated:

- ATTACH
- DETACH
- INDICATE LOAD
- INDICATE NSS
- INDICATE PAGING
- INDICATE SPACES
- INDICATE USER
- LOGOFF
- LOGON
- MONITOR SAMPLE
- QUERY MDCACHE
- QUERY RESERVED
- QUERY SRM
- QUERY VIRTUAL ALL
- SET CPTRACE
- SET MDCACHE
- SET SRM
- VMRELOCATE

The following DirMaint operand has been removed:

- XSTORE

The following DirMaint messages have been removed:

- DVH3261I
- DVH3324E

The following DIAGNOSE code has been updated:

- DIAGNOSE Code X'84' - Directory Update-in-Place

  The XSTORE operation is supported for compatibility only; the operation has no function.

The following CP monitor records are no longer available:

- Domain 0 Record 5 - MRSYTXSP - Expanded Storage Data (per processor)
- Domain 1 Record 17 - MRMTRXSG - Expanded Storage Data
- Domain 3 Record 9 - MRSTOXSG - Expanded Storage Data
- Domain 3 Record 10 - MRSTOXSU - Expanded Storage Data (per user)

The following CP monitor records have been changed:

- Domain 0 Record 14 - MRSYTXSG - Minidisk Cache
- Domain 1 Record 16 - MRMTRSCH - Scheduler Settings - Sample Record
- Domain 2 Record 4 - MRSCLADL - Add User To Dispatch List - Event Record
- Domain 2 Record 5 - MRSCLDDL - Drop User From Dispatch List - Event Record
- Domain 2 Record 6 - MRSCLAEL - Add User To Eligible List - Event Record
- Domain 2 Record 7 - MRSCLSRM - SET SRM Changes - Event Record
- Domain 3 Record 3 - MRSTOSHR - Shared Storage Management (per NSS or DCSS)
- Domain 3 Record 14 - MRSTOASI - Address Space Information Record
- Domain 3 Record 16 - MRSTOSHD - NSS/DCSS/SSP Removed From Storage
- Domain 4 Record 2 - MRUSELOF - User Logoff Data - Event Record
- Domain 4 Record 3 - MRUSEACT - User Activity Data
- Domain 4 Record 9 - MRUSEATE - User Activity Data at Transaction End - Event Record

[V6.4] 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-Express5S or later 1000BASE-T Ethernet adapters for primary and backup connectivity.

Prior versions of z/VM are not supported. There is no FICON ECKD or FICON CTC support in DPM at this time. Therefore, single system image (SSI) clusters are not supported.

[V6.4 APAR] Concurrent I/O Support for IBM XIV Storage System

With the PTF for APAR VM65929, the z/VM SCSI container enables multiple I/O requests to be issued concurrently to EDEVICEs backed by IBM XIV System Storage hardware, which may improve performance. This support particularly benefits EDEVICE paging I/O or volumes containing multiple minidisks.

[V6.4 APAR] z/VM Support for IBM z14

With the PTF for APAR VM65942, z/VM V6.4 provides support that enables guests to exploit function supported by z/VM on IBM z14™ (z14™), which includes:

- **z/Architecture support**

  The z/VM Stand-Alone Program Loader (SAPL) utility, DASD Dump Restore (DDR), and the Stand-Alone Dump utility have been enhanced to run entirely in z/Architecture mode. z/Architecture support for the SAPL and DDR utilities is in the base of z/VM V6.4 and provided for z/VM V6.3 with the available PTF for APAR VM65856. z/Architecture support for the Stand-Alone Dump utility is in the base of z/VM V6.4 and provided for z/VM V6.3 with the available PTFs for APARs VM65921 and VM65922.

- **New hardware facilities**

  z/VM enables guest use of new instructions and capabilities available on IBM z14, including the following facilities:

  – Miscellaneous-Instruction-Extensions Facility 2
  – Vector Enhancements Facility 1
  – Vector Packed Decimal Facility
  – Message-Security-Assist Extensions 6, 7, and 8

  Note: TRACE, DISPLAY I, and VMDUMPTL display support for interpreting the new instructions associated with these facilities is not provided.

- **ESA/390-compatibility mode for guests**

  IBM z14 does not support the full ESA/390 architectural mode. However, IBM z14 does provide ESA/390-compatibility mode, a hybrid architectural mode that supports a subset of ESA/390 needed to run CMS and GCS and their applications, but lacks more advanced functions needed for full function ESA/390 operating systems.

  Specifically, ESA/390-compatibility mode does not provide dynamic address translation (DAT), access-register translation (ART), address-space-number (ASN) translation, branch tracing, and related instructions and functions. Use of ART to access VM Data Spaces in XC virtual machines remains supported. Facilities, instructions, controls, and modes defined only in z/Architecture may or may not be fenced off in ESA/390-compatibility mode; attempted use of these gives unpredictable results.

  When an ESA or XA virtual machine logs on, it will be put into either full ESA/390 mode or ESA/390-compatibility mode depending on the level of

ESA/390 capability available in the machine where the virtual machine is logged on or in the virtual machine's relocation domain.

The PTF for APAR VM65976 provides infrastructure support for ESA/390-compatibility mode within z/VM V6.2, V6.3, and V6.4, and must be installed on all members of an SSI cluster before any z/VM V6.3 or V6.4 member of the cluster is run on an IBM z14 server.

- **Support for the Crypto Express6S**

  z/VM support for the new Crypto Express6S (CEX6S) adapter is included for both shared and dedicated guest use. As with the prior crypto adapter support, the CEX6S adapter can be configured as an accelerator or as an IBM Common Cryptographic Architecture (CCA) coprocessor for shared or dedicated use by z/Architecture guests. When the CEX6S adapter is configured as an IBM Enterprise Public-Key Cryptography Standards (PKCS) #11 (EP11) coprocessor, the domains on the adapter can be dedicated to z/Architecture guests, but not shared.

  With Crypto Express6S support and support for the new and enhanced CPACF functions, z/VM V6.3 and V6.4 provide the prerequisite IBM z14 encryption support to enable exploitation by guests in support of pervasive encryption of data in flight and at rest.

- **Crypto clear key ECC operations**

  Clear key Elliptic Curve Cryptographic (ECC) operations are supported for guests enrolled in the z/VM-managed shared-crypto queue (APVIRT). The shared-crypto queue must be configured with CCA coprocessor domains in order for the function to be virtualized for guest use.

- **Dynamic I/O support**

  Dynamic I/O support is provided for managing the configuration of OSA-Express6S OSD CHPIDs, FICON Express16S+ FC and FCP CHPIDs, and Regional Crypto Enablement (RCE), zHyperLink Express, and RoCE Express2 adapters.

  The PTF for APAR VM65865 provides dynamic I/O support for the Coupling Express Long Reach (CE LR) adapter in z/VM V6.3 and V6.4 and is applicable to z13, z13s, and z14.

- **RoCE Express2 support**

  Guest exploitation support for RoCE Express2 allows the adapters to be brought online and attached to supporting guests for exploitation.

- **Improved memory management efficiency**

  The IBM z14 processor design allows greater concurrency in address translation. This improvement may increase z/VM workloads' performance compared with z13, particularly when z/VM is configured to exploit multithreading.

This z/VM support also provides instruction simulation counters by virtual CPU in the z/VM monitor data.

The following configuration statement has been updated:
- CRYPTO

The following directory statements have been updated:
- CRYPTO
- MACHINE

The following CP commands have been updated:
- DEFINE
- DEFINE DEVICE / IODEVICE

- DEFINE PCIFUNCTION
- DELETE DEVICE / IODEVICE
- LOGON
- MODIFY DEVICE / IODEVICE
- QUERY CRYPTO
- QUERY PCIFUNCTION
- QUERY SET
- QUERY VIRTUAL CRYPTO
- SET MACHINE
- XAUTOLOG

The following CP monitor records have been updated:

- Domain 0 Record 15 - MRSYTCUG - Logical Partition Configuration
- Domain 0 Record 16 - MRSYTCUP - CPU Utilization in a Logical Partition
- Domain 0 Record 17 - MRSYTCUM - Physical CPU Utilization Data for LPAR Management
- Domain 1 Record 4 - MRMTRSYS - System Configuration Data
- Domain 1 Record 27 - MRMTRPCI - PCI function Configuration Data
- Domain 4 Record 2 - MRUSELOF - User Logoff Data - Event Record
- Domain 4 Record 3 - MRUSEACT - User Activity Data
- Domain 5 Record 10 - MRPRCAPM - Crypto Performance Measurement Data
- Domain 5 Record 16 - MRPRCPUP - Park/Unpark Decision (Event)
- Domain 6 Record 42 - MRIODPAD - PCI function added to the system
- Domain 6 Record 45 - MRIODPON - Real PCI function varied on

With the PTF for APAR PI73016, the TCP/IP stack and the NETSTAT OSAINFO command are updated to support OSA-Express6S adapters.

Support for the IBM z14 also requires the following APARs. For more information, see the documentation for the product, feature, or function.

- APAR PI46151 provides ICKDSF support.
- APAR VM65952 provides EREP/VM support.
- APAR VM65843 provides HCD support.
- APAR VM65266 provides HCM support.
- APAR VM65939 provides IOCP support.
- APARs PI62275 and PI65715 provide HLASM support.

For more information, see:

- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*

Installing z/VM on IBM z14

IBM Z clients currently using z/VM V6.3, who are interested in running z/VM V6.3 on IBM z14, should first consider moving to z/VM V6.4. z/VM V6.3 must be installed on a prior IBM Z server and the PTFs for APARs VM65942, VM65921, and VM65922 must be applied prior to moving the image to a z14. After applying all required service and before IPLing z/VM V6.3 on a z14 server, an updated SAPL utility must be installed by using the updated SALIPL MODULE provided by the PTF for APAR VM65856. Otherwise, z/VM V6.3 will not successfully IPL on a z14 server. In addition, after applying the PTFs for APARs VM65921 and VM65922, the Stand-Alone Dump program must be reinstalled using the updated SDINST utility.

z/VM V6.4 can be installed directly on a z14 with an image obtained from IBM after August 25, 2017. The PTF for APAR VM65942 must be applied immediately after installing z/VM V6.4.

[V6.4 APAR] Extended Address Volume Minidisk Support

With the PTFs for APARs VM65943 (CP), VM65945 (CMS), and PI85943 (ICKDSF), enhanced extended address volume (EAV) support for 3390-A DASD devices supporting 28-bit cylinder addresses is provided. This support allows non-fullpack minidisks to reside anywhere on the volume, including beyond the current restriction of the 64 KB cylinder boundary (0-65519), and up to the 1 TB limit currently supported.

If a guest virtual machine with an EAV minidisk defined above cylinder 65520 is relocated within an SSI cluster, the destination system must have EAV support installed; otherwise, the relocation will fail.

[V6.4 APAR] Processor Scalability Efficiency Improvements

With the PTF for APAR VM65988, the z/VM hypervisor is enhanced to manage its spinlocks more efficiently and thereby reduce system overhead. This enhancement will contribute to improved performance and throughput, and thereby help to improve overall system capacity by allowing additional work to be performed. While most workload configurations will benefit to some extent, the improvements are greatest for workloads using large numbers of logical CPUs. The shared-exclusive spinlock manager was replaced with a more cache-efficient design providing greater scalability for the Scheduler Lock (SRMSLOCK). SRMSLOCK reporting in Monitor D0R23 MRSYTLCK was corrected to include time spinning in HCPDSP's internal spin loop that was previously unreported. The support is available on all hardware supported by z/VM V6.4, including IBM z14. However, the design for some spinlocks take advantage of efficiencies unique to the z14.

The following CP monitor record was added:

- Domain 5 Record 22 - MRPRCSXL - Shared-Exclusive Spin Lock Utilization (per-processor)

Under-reporting of scheduler lock spin time was corrected by including the informal spin time in the dispatcher in the lock statistics reported in the following updated monitor records:

- Domain 0 Record 2 - MRSYTPRP - Processor data (Per Processor):
  SYTPRP_PFXSPINT includes the informal spin time.
  YTPRP_PFXSPINC includes the count of informal spins.
- Domain 0 Record 10 - MRSYTSCG - Scheduler Activity
  SYTSCG_CALSLKTM includes the informal spin time.
  SYTSCG_CALSLKCT includes the count of informal spins.
- Domain 0 Record 23 - MRSYTLCK - Formal Spin Lock Data
  SYTLCK_CALSTIME includes the informal spin time.
  SYTLCK_CALSSCNT includes the count of informal spins.

As a result reduced spinlock overhead from the shared-exclusive spinlock manager improvements may appear to reduce system overhead that was not previously reported as spinlock spin time.

[V6.4 APAR] Alternate Subchannel Set Dynamic I/O Support

With the PTF for APAR VM65942, z/VM V6.4 provides dynamic I/O commands within the z/VM hypervisor for supporting define, modify, and delete of a device within either the default or alternate subchannel set. HCD support is currently not available.

[V6.4 APAR] Coupling over RoCE (CL5) Channel Path Support

With the PTF for APAR VM65865, z/VM V6.3 and V6.4 provide dynamic I/O support for the new Coupling over RoCE (CL5) channel path type, which is associated with the Coupling Express Long Reach (CE LR) adapter. CE LR is a two port adapter for the IBM z13, z13s, and z14 that provides long distance coupling connectivity.

The following CP commands have been updated for this support:

- DEFINE CHPID/PATH
- QUERY CHPID

For more information, see *z/VM: CP Commands and Utilities Reference*.

Support for the CL5 CHPID also requires the following APARs:

- HCD support requires APAR VM65849.
- HCM support requires APAR VM65208.
- IOCP support requires APAR VM65880.

Connectivity and Networking

These topics describe changes that can affect how z/VM systems, applications, and guest operating systems communicate with each other.

Notes:

1. For information about z/VM support for hardware connectivity facilities, see "Support and Exploitation of Hardware and Architectures" on page 49.
2. This section does not include TCP/IP changes prior to z/VM V4.2. For that information, see *z/VM: TCP/IP Planning and Customization*. For TCP/IP migration considerations and other TCP/IP topics, also see IBM: TCP/IP for z/VM and VM/ESA (www.ibm.com/vm/related/tcpip/).

[V2.3] MQSeries Client for VM/ESA

MQSeries® Client for VM/ESA enables applications using message queuing to communicate across different platforms using client-server technology.

MQSeries Client for VM/ESA executes in any CMS virtual machine and provides a simple user interface to the server environment. It is not a full function queue manager. Therefore it needs to be connected through TCP/IP or APPC to other queue managers.

The application Language Bindings supported are:

- IBM C for VM/ESA
- IBM VS Cobol II
- IBM PL/I Compiler
- REXX/VM
- IBM Assembler

For more information, see *MQSeries: Application Programming Guide*, SC33-0807.

[V4.2] Guest LAN

z/VM supports a network of virtual adapters connecting guests within a z/VM system. The CP component of z/VM has been enhanced to provide:

- Virtual HiperSockets adapters

  CP now offers a virtual network interface card (NIC) that emulates the HiperSockets adapter. A guest should be able to operate the virtual adapter using the same software that would have been used to drive the equivalent hardware.
- Connectivity for virtual adapters

  CP now 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 an emulated LAN should be able to communicate using the same software that would have been used to communicate over an equivalent physical LAN segment.

The following CP commands have been added:
- DEFINE LAN
- DEFINE NIC
- DETACH LAN
- DETACH NIC
- QUERY LAN
- QUERY NIC
- QUERY VMLAN
- SET LAN
- UNCOUPLE

The following CP functions have been updated:
- COUPLE command
- DEFINE (Virtual Device) command
- QUERY (Virtual Device) command
- QUERY VIRTUAL OSA command
- SPECIAL directory control statement

[V4.2] TCP/IP Stack Vulnerability Reduction

Function has been added to improve the performance and reliablity of the TCP/IP stack by recording and reporting some Denial of Service (DOS) attacks and preventing their propagation. These attacks include:

- Smurf — ICMP Echo Request packets sent to IP broadcast or multicast addresses
- Fraggle — UDP Echo Request packets sent to IP broadcast or multicast addresses
- Ping-o-Death — ICMP Echo Request packets that are too large.

[V4.3] Guest LAN Enhancements

The guest LAN implementation has been extended to support multicast transmission across HiperSockets guest LANs. Also, guest LANs can now be defined to function as either HiperSockets transport media or OSA-Express QDIO transport media. The addition of QDIO simulation allows the virtualization of a QDIO LAN environment and provides support for the broadcast capability that is part of the QDIO architecture. As with the z/VM HiperSockets simulation, no real hardware is required to support the OSA-Express QDIO simulation.

Serviceability aids have been added in the form of an AIF trace option to trace the occurrence of Adapter Interruption Facility events, and external symbols are defined to mark key points for adding TRSOURCE traps.

The following new CP commands have been added for this support:
- SET VMLAN
- TRACE AIF

The following CP functions have been updated:
- DEFINE LAN system configuration statement
- SPECIAL directory control statement
- COUPLE command
- DEFINE LAN command
- DEFINE NIC command
- QUERY LAN command
- QUERY NIC command
- QUERY VMLAN command
- SET LAN command

For more information, see:
- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*

[V4.3] TCP/IP Device and Stack Performance Improvements

Key sections of TCP/IP routing, device driver, and other processes have been optimized. By primarily recoding these processes in Assembler Language, improvements in performance of high-use code paths in the TCP/IP stack have been achieved. In addition to recoding, some algorithms were changed to provide equivalent functions at lower CPU costs.

When TCP/IP is configured to use the equal-cost multipath support, the routing table will maintain up to four equal-cost paths to a particular destination, providing load balance support and better performance. See *z/VM: TCP/IP Planning and Customization* for further information.

Additional support has been added to the NETSTAT command. Functions of OBEYFILE such as starting and stopping devices and performing other simple control tasks can be accomplished by using the OBEY function of the NETSTAT command.

A new FTP PASSIVE subcommand has been added to control whether the client or the server establishes connections for data transfers. See *z/VM: TCP/IP User's Guide* for further information.

The TCP/IP stack has been enhanced to support the following:
- Multicast for HiperSockets
- Broadcast for QDIO

[V4.3] TCP/IP Dynamic Stack Configuration

Authorized users now can define, change or display the TCP/IP configuration dynamically. These changes are temporary and are discarded when the TCP/IP stack virtual machine is restarted. This support involves the following commands:

IFCONFIG
Added with z/VM V4.3; see *z/VM: TCP/IP Planning and Customization* for more information.

NETSTAT DEVLINKS
Updated to display additional information about the devices and links defined for the TCPIP virtual machine; see *z/VM: TCP/IP User's Guide* for more information.

QUERY (Real Device)
Updated with a new operand, ID, which displays the sense ID information returned by a device and its control unit; see *z/VM: CP Commands and Utilities Reference* for more information.

The IFCONFIG command can also be used to generate syntactically-correct configuration statements for inclusion in the PROFILE TCPIP file in order to make permanent changes to the network configuration.

[V4.3] TCP/IP Stack Vulnerability Reduction

Function has been added to improve the performance and reliability of the TCP/IP stack by recording and reporting additional Denial-of-Service (DoS) attacks and preventing their propagation. These attacks include:

- Kiss-of-Death (KOD) — an IGMP based attack that depletes the stack's large envelopes
- KOX — a version of the KOD attack that also has source IP address spoofing
- Stream — an attack in which TCP packets are sent to the stack with no header flags set
- R4P3D — an augmented version of the Stream attack
- Blat — a version of the Land attack that also has the URG flag turned on in the TCP header and has the ability to incrementally spoof the source IP address
- SynFlood — an attack in which the initiator floods the TCP/IP stack with SYN packets that have spoofed source IP addresses, resulting in the server never receiving the final ACKs needed to complete the three-way handshake in the connection process.

The Smurf DoS attack has also been updated to address three variants of the attack. Smurf is a DoS attack in which an ICMP Echo Request is sent to a broadcast or multicast address. The three variants are:

- Smurf-IC — where "IC" denotes that incoming packets are using the TCP/IP stack to launch an attack
- Smurf-OB — where "OB" denotes that an outbound ICMP Echo Request matched the description of a Smurf attack
- Smurf-RP — where "RP" denotes that ICMP Echo Reply packets being received by the stack do not match any Echo Requests that were sent.

[V4.4] Guest LAN Support for IPv6

Guest LAN support for OSA-Express simulation in QDIO mode has been updated for Internet Protocol version 6 (IPv6). Virtual machines in the guest LAN environment can define and use simulated OSA-Express devices that support both the IPv4 and IPv6 protocols.

[V4.4] IEEE VLAN Support

IEEE Virtual LAN (VLAN) standard 802.1q describes a mechanism that is designed to enable groups of hosts to be logically connected into a single local area network (LAN) even though they are connected to different switches in different physical locations. Network hosts can be organized into LAN segments which fit an organization's network traffic patterns, rather than being dependent on physical location. In addition, by connecting a router to multiple VLANs, savings can be realized by providing router connections to new LAN segments without having to add additional network interfaces to the routers.

To support IEEE VLANs, z/VM V4.4 provides:

- Enhancements to TCP/IP for z/VM to enable membership in a VLAN for IBM zSeries OSA-Express (QDIO) and HiperSockets adapters that support IEEE 802.1q
- Enhancements to z/VM's virtual QDIO and HiperSockets network interface simulation to support VLAN frame tagging as described by IEEE 802.1q
- Management and control of the VLAN identifiers (VIDs) that can be used by guest virtual machines

[V4.4] TCP/IP Broadcast Support for HiperSockets

TCP/IP broadcast support is available for the HiperSockets environment when utilizing Internet Protocol version 4 (IPv4). Applications that use the broadcast function can propagate the broadcast frames to all TCP/IP applications when using either HiperSockets or OSA-Express QDIO. Broadcast support is automatically enabled for the TCP/IP stack.

[V4.4] TCP/IP IMAP User Authentication Exit

Previously, TCP/IP Internet Message Access Protocol (IMAP) users required a VM user ID and password to access the IMAP mail store on z/VM. Through a new user-written IMAP authentication exit, a VM user ID and password, with the accompanying limitation of eight characters, are no longer required. The authentication exit runs in a separate virtual machine (IMAPAUTH), authenticates user IDs and passwords, and maps the IMAP user's ID to an eight-character SFS ID by which the user accesses the mail store.

You have the flexibility of using whatever authentication program you want to use. The authentication exit is enabled through a new TCP/IP configuration statement, AUTHENTICATEID. If enabled, the exit is called to handle the following types of requests:

- Validating LOGIN commands
- Mapping a long name to a short name
- Pre-authorizing connections

If the exit is not enabled, the IMAP server validates IMAP clients through CP or an external security manager.

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

[V4.4] TCP/IP SSL Server Upgrade

The TCP/IP Secure Sockets Layer (SSL) server is now compatible with the Linux 2.4 kernel. The upgraded SSL server provides appropriate Red Hat Package Manager (RPM) packages for the SuSE Linux 2.4.7 SLES 7 and 2.4.19 SLES 8

distributions. This upgrade helps avoid potential support issues with the backlevel 2.2.16 Linux kernel and allows the SSL server to exploit the many performance improvements in the Linux 2.4 kernel.

[V4.4] TCP/IP Stack Performance Improvements

The performance of the TCP/IP stack has been improved by optimizing additional high-use code paths and adding virtual multiprocessing capabilities. Individual device drivers can now be associated with particular virtual processors. A new CPU option on the DEVICE configuration statement designates the CPU where the driver for a particular device will be dispatched, allowing up to seven virtual processors to be exploited. This enables the TCP/IP load to be spread across multiple real processors and, in high-activity environments, can improve responsiveness and throughput. If your TCP/IP load ordinarily uses a substantial portion of a single processor, there may be benefits to creating a multiprocessor configuration. For more information, see *z/VM: TCP/IP Planning and Customization*.

[V4.4] TCP/IP Stack Security Improvements

The overall security and auditability of the TCP/IP stack and the integrity of the z/VM system have been improved to aid self-protection by providing better controls, monitoring, and defaults, which include:

- Logging of all TCP/IP administrative commands (including NETSTAT, OBEYFILE, and IFCONFIG) that alter, or attempt to alter, the active IP or CP configuration.
- Changing the defaults of the ASSORTEDPARMS statement within the TCP/IP configuration to:
  - **RESTRICTLOWPORTS**, to prevent impersonation ("spoofing") of well-known z/VM TCP/IP applications by CMS users or guest operating systems
  - **VARSUBNETTING**, to avoid creating unnecessary TCP/IP routing table entries

The security of the TCP/IP stack has been improved by making the RESTRICTLOWPORTS operand of the ASSORTEDPARMS statement active ***by default.*** Thus, all TCP/IP applications that listen on "well-known" ports (ports 1 through 1023) must be given permission to do so. Such permission can be granted by customizing the TCP/IP server configuration file (PROFILE TCPIP, or its equivalent) in one of three ways:

1. Use the PORT statement to reserve the specific port (or ports) required by each application (virtual machine) used on your system. **This is the preferred method.** Note that with TCP/IP Level 440, ports can reserved within a specific range, in addition to being reserved on an individual basis.
2. Modify the OBEY statement to include the affected virtual machines in the TCP/IP obey list.
3. Include the FREELOWPORTS operand as part of an ASSORTEDPARMS statement. Note that *this method removes the default protection for all well-known ports.*

Note: When the RESTRICTLOWPORTS default is in effect and appropriate port authorizations have not been provided, applications that rely upon well-known ports (for example, VM-based web servers or remote printing functions such as **lpr**) are likely to report "Unable to open port(s)" or "Permission denied" conditions.

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

[V4.4] Virtual FICON CTCAs

CP's virtual I/O subsystem supports the simulation of FICON Channel-to Channel Adapters (CTCAs). This support enhances previous virtual-CTCA support by adding the FICON protocol as an option for guest operating systems. Virtual CTCAs now support the 3088, ESCON, and FICON protocols. A new CTCA device subclass, FCTC, has been defined for the virtual FICON CTCA.

The following CP functions have been updated:
- SPECIAL directory statement
- DEFINE CTCA command
- QUERY VIRTUAL CTCA command

For more information, see:
- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*

[V4.4] Virtual Switch

z/VM further enhances virtualization technology by introducing a virtual IP switch that is capable of bridging a guest LAN to an associated real LAN connected by an OSA-Express adapter. The virtual switch is designed to help eliminate the need for virtual machines acting as routers to provide IPv4 connectivity to a physical LAN through an OSA-Express adapter. Further, it eliminates the need to define a separate routable subnet for the exclusive use of the members of a guest LAN. Using the virtual switch, the convenience of a guest LAN is maintained while allowing the guests to be assigned IP addresses in the real LAN subnet.

Virtual routers consume valuable processor cycles to process incoming and outgoing packets, requiring additional copying of the data being transported. The virtual switch helps alleviate this problem by moving the data directly between the real network adapter and the target or originating guest data buffers.

Centralized network configuration and control of the virtual switch within CP allows the guest LAN administrator to more easily grant and revoke access to the real network and to manage the configuration of guest LAN VLAN segments. While the z/VM system can be a member of multiple VLANs, the guest LAN administrator can control which guests belong to which real VLAN, without requiring additional network adapters or switch port configuration. If a guest does not support IEEE 802.1q, z/VM will transparently join the virtual network interface into the desired VLAN.

The virtual switch employs transparent bridging to enable the switch to dynamically determine and maintain node connectivity so that the LAN administrator has less network maintenance to perform.

The following new CP functions have been added for this support:
- DEFINE VSWITCH system configuration statement
- MODIFY VSWITCH system configuration statement
- MODIFY LAN system configuration statement
- NICDEF directory control statement
- DEFINE VSWITCH command
- DETACH VSWITCH command

- QUERY CONTROLLER command
- QUERY VSWITCH command
- SET VSWITCH command

The following CP functions have been updated:
- VMLAN system configuration statement
- IUCV directory statement
- SPECIAL directory statement
- COUPLE command
- DEFINE NIC command
- QUERY NIC command
- QUERY LAN command
- QUERY VMLAN command
- UNCOUPLE command

The following new Directory Maintenance Facility functions have been added:
- NICDEF command
- SPECIAL command

The following Directory Maintenance Facility functions have been updated:
- IUCV command

For more information, see:
- *z/VM: Connectivity*
- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: Directory Maintenance Facility Commands Reference*

[V5.1] Enhanced IPv6 Support

z/VM V5.1 enhances its IPv6 support by allowing the z/VM TCP/IP stack to be configured for IPv6 networks connected through OSA-Express operating in QDIO mode. The stack can be configured to provide static routing of IPv6 packets and to send IPv6 Router Advertisements. The native z/VM TCP/IP applications that have been enhanced to support IPv6 are TRACERTE and PING. In addition, support is being provided to help application developers to create socket applications for IPv6 communications. This support includes updates to the C/C++ Language sockets through Language Environment and to the OpenExtensions callable services library (CSL) socket APIs.

For more information, see:
- *z/VM: TCP/IP User's Guide*
- *z/VM: OpenExtensions Callable Services Reference*
- *XL C/C++ for z/VM: Runtime Library Reference*, SC09-7624

[V5.1] Enhanced Virtual Switch Support — Failover

The virtual switch has been improved in z/VM V5.1 to provide enhanced failover support for less disruptive recovery for some common network failures. Enhancements include the ability to:
- Recover from the failure of a virtual switch's network connection by swapping from a failing OSA-Express device to a partially initialized backup device in less time, helping to reduce data loss
- Detect a stalled OSA-Express device associated with a virtual switch, with failover to a backup device

- Detect a nonfunctioning z/VM TCP/IP controller and failover to a backup controller
- Limit the number of SETIP requests sent to an OSA-Express to help prevent overload of the device

For more information about the virtual switch, see *z/VM: Connectivity*.

[V5.1] Enhanced Virtual Switch Support — VLAN

The following changes have been made for deploying VLAN technology on z/VM:

- The VLAN ANY operand has been removed from the SET VSWITCH command and MODIFY VSWITCH statement.
- A new VLAN awareness attribute on a virtual switch has been defined. You can define a virtual switch as VLAN AWARE or VLAN UNAWARE. The awareness attribute can be specified on the DEFINE VSWITCH command and statement for IP and ETHERNET virtual switches.
- A new PORT attribute is defined on the DEFINE VSWITCH and SET VSWITCH commands. The supported ports are ACCESS and TRUNK.

See the following for more information:

- *z/VM: Connectivity* for information about the virtual switch
- *z/VM: CP Planning and Administration* for the changes to the DEFINE VSWITCH and MODIFY VSWITCH statements
- *z/VM: CP Commands and Utilities Reference* for the changes to DEFINE VSWITCH, SET VSWITCH, QUERY VSWITCH, QUERY LAN, and QUERY NIC.

[V5.1] Virtual Switch Exploitation of Layer 2 Support

z/VM V5.1 supports a new transport option for the virtual switch that defines the virtual switch as operating in Layer 2 mode. Layer 2 mode is supported for the OSA-Express and OSA-Express2.

In Layer 2 mode, each port on the virtual switch is referenced by its Media Access Control (MAC) address instead of by Internet Protocol (IP) address. Data is transported and delivered in Ethernet frames, providing the ability to handle protocol-independent traffic for both IP (IPv4 or IPv6) and non-IP, such as IPX, NetBIOS, or SNA. Coupled with the Layer 2 support in Linux for zSeries and the OSA-Express and OSA-Express2 support for the z890 and z990, Linux images deployed as guests of z/VM can operate more efficiently.

Port sharing on the OSA-Express or OSA-Express2 allows many stacks to share the same physical OSA NIC card, and provides a communication path between the LPAR and z/VM without having to send the traffic flow over the physical wire. There is a restriction with the Layer 2 support in that port sharing is supported only between virtual switches that are of the same transport type, such as Layer 2 with Layer 2 and Layer 3 with Layer 3. A Layer 2 guest cannot communicate directly with a Layer 3 guest on the same adapter.

For more information about the virtual switch, see *z/VM: Connectivity*.

[V5.1] Additional Device Connections for TCP/IP for z/VM

- TCP/IP stack utilization improvements for OSA-Express

  For the z890 and z990 servers, the number of TCP/IP stacks that can share an OSA-Express adapter has been increased from 84 to 160. This increase provides

additional connections to enable more virtual machines to be connected to an external network. z/VM V5.1 supports this hardware enhancement.

- Improved virtualization supporting more TCP/IP stacks with OSA-Express2

  With the new OSA-Express2 feature for the z890 and z990, the maximum number of connections (TCP/IP stacks) allowed can be increased up to 640. Support for 640 TCP/IP stacks is provided by the OSA-Express2 features (Gigabit Ethernet SX, Gigabit Ethernet LX, Gigabit Ethernet LR) running in QDIO mode only. z/VM V5.1 supports this new capability to allow additional connections to virtual machines, particularly Linux images.

[V5.2] Enhanced Dynamic Routing Capabilities with a New MPRoute Server

A new MPRoute server has been adapted from z/OS V1.7 to operate on z/VM V5.2 and supports the following protocols:

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

The new MPRoute server removes the limit of four equal-cost paths and allows the generation of up to 16 equal-cost routes to a destination, thus providing improved load-balancing support.

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

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

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

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

[V5.2] Improved Problem Determination for Guest LANs and Virtual Switches

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

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

  When a Linux guest is deployed, traffic can be traced, recorded, and analyzed by existing tools directly from the guest virtual machine. This Linux guest must

be authorized to use this capability through CP commands. The authorized guest can then use CP commands or the Linux device driver (when available) to put the guest NIC in "Promiscuous Mode". The collected data can be printed or written to a file using tools such as **tcpdump** (available with Linux for z Systems) to determine the causes of the network problems

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

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

The following CP command has been added:
- SET NIC

The following CP functions have been updated:
- QUERY LAN command
- QUERY VSWITCH command
- QUERY TRSOURCE command
- SET LAN command
- SET VSWITCH command
- TRACERED utility
- TRSOURCE ID command

For more information, see:
- *z/VM: Connectivity*
- *z/VM: CP Commands and Utilities Reference*

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

[V5.2] IPv6 HiperSockets Support

z/VM V5.2 allows the z/VM TCP/IP stack to be configured for IPv6 networks connected through HiperSockets operating in QDIO mode. The stack can be configured to provide static routing of IPv6 packets and to send IPv6 Router Advertisements. The NETSTAT command has been updated to support IPv6 for HIPERS devices. The LINK for HiperSockets, GATEWAY, HOME and ROUTERADV configuration statements have been updated.

For more information, see:
- *z/VM: TCP/IP User's Guide*
- *z/VM: TCP/IP Planning and Customization*

[V5.2] Predefined Virtual Switch Controllers

There is no longer a need to define and configure z/VM TCP/IP stacks to be available to control virtual switch connections to real LAN segments. Two controllers, DTCVSW1 and DTCVSW2, are predefined for you.

You can remove this capability from existing stacks by deleting the VSWITCH CONTROLLER statement from your TCP/IP configuration files. If you do this, make sure that you do not have any VSWITCH definitions that specify CONTROLLER *userid* for those stacks. Instead, allow them to default to CONTROLLER*.

[V5.2] Removal of OSA-Express Port Sharing Restriction on System z9

In OSA-Express on zSeries servers (z990, z890, z900, and z800), port sharing is supported only between LPARs or virtual switches that are of the same transport mode, Layer 2 with Layer 2 or Layer 3 with Layer 3. Attempted communications between a Layer 2 virtual switch and a Layer 3 virtual switch sharing the same OSA-Express adapter on a zSeries server will result in a network timeout condition. This restriction has been removed for System z9 servers. For more information, see *z/VM: Connectivity*.

[V5.2] SSL Server Upgrade

The SSL server included in TCP/IP for z/VM, level 520, supports the following Linux distributions:

- SUSE SLES8 Service Pack 3 (31-bit)
- SUSE SLES9 Service Pack 2 (31-bit)
- SUSE SLES9 Service Pack 2 (64-bit)
- Red Hat Enterprise Linux AS V3 (31-bit)
- Red Hat Enterprise Linux AS V3 (64-bit)
- Red Hat Enterprise Linux AS V4 (31-bit)
- Red Hat Enterprise Linux AS V4 (64-bit)

Additionally, this upgraded SSL server includes:

- A variety of industry-standard encryption algorithms, including DES, triple-DES, RC2, and RC4, with keys up to 128 bits in length. Hashes are provided by SHA-1 and MD5.
- Certificate activation and removal without server restart.

  This capability allows a new certificate to be added or an existing one to be deactivated while existing SSL-secured sessions are active, removing the need to shut down and restart the SSL server when certificates are added or deleted.
- Federal Information Processing Standard (FIPS 140-2) operational mode support.

  The SSL server may be operated in FIPS mode, restricting connections to those that employ FIPS-approved cipher suites.

[V5.3] Delete DEVICE and LINK Support

This support allows you to delete devices and links from the TCP/IP stack dynamically:

- The IFCONFIG command has changed to allow you to remove an interface from the TCP/IP configuration.
- The IOCTL call in C/C++ and REXX has changed to allow a program to remove an interface from the TCP/IP configuration.

[V5.3] Dynamic SSL/TLS

Dynamic SSL/TLS provides support for:

- Pascal applications through the use of several new APIs.
- Outbound handshakes.

[V5.3] Failover Support Enhancements for IPv4 and IPv6 Devices

Failover support for IPv4 and IPv6 devices has been improved in z/VM V5.3. When the z/VM TCP/IP stack has two (or more) Queued Direct Input/Output (QDIO) or LAN Channel Station (LCS) Ethernet devices on the same network and one device is stopped or fails, another device takes over responsibility for traffic destined for the failing device (or any devices the failing device had previously taken over). This failover support includes OSA-Express devices (in QDIO Ethernet or LCS Ethernet mode), Virtual IP Addresses (VIPAs), and addresses for which PROXYARP services are being provided through a takeover-eligible device. In addition to the basic failover support, one takeover-eligible device on that network will be responsible for informing other nodes on that network which hardware (MAC) address should be used to reach VIPA addresses on the TCP/IP stack, both when the stack initializes and when an IP takeover event occurs.

[V5.3] IP Takeover

IP takeover supports the following:

- Takeover capability for OSA express devices.
- Virtual IPv6 addresses to be defined for IPv6 devices.

[V5.3] LDAP Server Support

The z/VM Lightweight Directory Access Protocol (LDAP) server (adapted from z/OS V1.8) is based on a client/server model that provides client access to an LDAP server. An LDAP directory provides an easy way to maintain directory information in a central location for storage, update, retrieval, and exchange.

The LDAP server provides the following functions:

- Interoperability with any LDAP Version 2 or Version 3 directory client
- Access controls on directory information, using static, dynamic, and nested groups
- Secure Sockets Layer (SSL) communication (SSL V3 and TLS V1)
- Start TLS (Transport Layer Security) activation of secure communication
- Client and server authentication using SSL/TLS
- Password encryption
- Replication
- Referrals
- Aliases
- Change logging
- LDAP V2 and V3 protocol support
- Schema publication and update
- Native authentication
- CRAM-MD5 (Challenge-Response Authentication Method) and DIGEST-MD5 authentication
- Root DSE information
- LDAP access to information stored in RACF

Some of the z/VM LDAP server's features and capabilities include:

- **Multiple concurrent database instances** (referred to as backends). The LDAP server can be configured to serve multiple databases at the same time. This

means that a single z/VM LDAP server can respond to requests for many logically different portions of the LDAP tree. A z/VM LDAP server can be configured to provide access to RACF, as well as store application-specific information.

- **Robust general-purpose databases**. The LDAP server comes with an LDBM backend. There are no restrictions on the types of information that this backend can contain. The LDBM backend keeps its entries in memory for quick access and requires a minimum amount of setup. When the LDAP server is not running, LDBM stores its directory information in OpenExtensions files.
- **Access to RACF data**. The LDAP server can be configured to provide read/write access to RACF user, group, and connection profiles using the LDAP protocol. The LDAP server's access to RACF is managed by an additional configurable backend called SDBM.

  Note: To use SDBM for ONLY authentication (LDAP bind processing), any security manager implementing the SAF service required by the __passwd() function call can be used. To use SDBM for accessing and updating user, group, and connection profile information, RACF is required.
- **Access control**. The LDAP server provides a rich and powerful access control facility, allowing you to control access to the information in your database or databases. You can control access to entries based on LDAP authentication information, including users and groups. Group membership can be either static, dynamic, or nested. Access control is configurable down to individual attributes within entries. Also, access controls can be set up to explicitly deny access to information.
- **Threads**. The LDAP server is threaded for optimal performance. A single multi-threaded z/VM LDAP server process handles all incoming requests, reducing the amount of system overhead required.
- **Referrals**. The LDAP server provides the ability to refer clients to additional directory servers. By using referrals, you can distribute processing overhead, distribute administration of data along organizational boundaries, and provide potential for widespread interconnection beyond an organization's own boundaries.
- **Aliases**. An alias entry can be created in the directory to point to another entry in the directory. During search operations, an alias entry can provide a convenient public name for an entry or subtree, hiding the more complex actual name of the entry or subtree. It can also avoid the need to duplicate an entry in multiple subtrees.
- **Change Logging**. The LDAP server can be configured to create change log entries in the GDBM backend. Each change log entry contains information about a change to an entry in an LDBM backend, to the LDAP server schema, or to a RACF user, group, or connection profile.
- **Configuration**. The LDAP server is highly configurable through a single configuration file which allows you to change just about everything you would ever want to change. Configuration options have reasonable defaults, making your job much easier.
- **Secure communications**. The LDAP server can be configured to encrypt data to and from LDAP clients using SSL. The LDAP server supports the Start TLS extended operation to switch a non-secure connection to a secure connection. It has a variety of ciphers for encryption to choose from, all of which provide server and optionally client authentication through the use of X.509 certificates.
- **Native authentication**. The z/VM LDAP server allows clients to bind to entries in an LDBM backend by using the system for verifying the authentication

attempt. The client can perform a simple bind supplying an LDAP distinguished name of an entry in an LDBM backend along with a security manager-maintained password. Password authentication is then performed by the security manager.

Note: To use native authentication, any security manager implementing the SAF service required by the __passwd() function call can be used.

- **LDAP Version 3 protocol support**. The LDAP server provides support for Version 3 of the LDAP protocol in addition to the LDAP Version 2 protocol. Version 3 includes:
  - All protocol operations
  - Implicit bind
  - Certificate (or Simple Authentication and Security Layer) bind
  - Version 3 referrals
  - Aliases
  - Controls
  - Root DSE support
  - Internationalization (UTF-8) support
  - Modify name supported for all entries including subtree move
  - Schema publication
  - Additional syntax support
  - Online schema update capability
- **Dynamic schema**. The LDAP server allows the schema to be changed dynamically through the LDAP protocol.
- **Internationalization (UTF-8) support**. The LDAP server allows storage, update and retrieval, through LDAP operations, of national language data using LDAP Version 3 protocol.
- **SASL external bind and client and server authentication**. The LDAP server allows client applications to use a certificate when communicating with the server using SSL/TLS communications. In order to use a certificate on bind, the server must be configured to perform both client and server authentication. This ensures both entities are who they claim to be.
- **SASL CRAM-MD5 and DIGEST-MD5 authentication**. The LDAP server allows clients to bind to the server using DIGEST-MD5 (RFC 2831) and CRAM-MD5 (Challenge-Response Authentication Method - RFC 2195) authentication bind methods.
- **Support for root DSE**. The LDAP server supports search operations, including subtree search, against the root of the directory tree as described in IETF RFC 2251, *The Lightweight Directory Access Protocol (V3)*. The so-called Root DSE can be accessed using LDAP V3 search operations.
- **Extended group membership searching**. The LDAP server supports extended group membership searching which allows the LDAP server to find a distinguished name that may be a member of static and nested groups in a backend (LDBM) where the distinguished name does not reside. The LDAP server can find the group memberships for the distinguished names in the other backends that are configured.
- **Supported server controls**. The LDAP server supports the following:

  authenticateOnly
  IBMLdapProxyControl
  IBMModifyDNRealignDNAttributesControl

IBMModifyDNTimelimitControl
IBMSchemaReplaceByValueControl
manageDsaIT
PersistentSearch
replicateOperationalAttributes

- **Attribute encryption**. The LDAP server supports encryption of the values of several critical attributes to prevent unauthorized access to these attribute values in LDBM backends. The attributes that can be encrypted are as follows:

  replicaCredentials
  secretKey
  userPassword
- **Multiple socket ports**. The LDAP server can be configured to listen for secure and nonsecure connections from clients on one or more IPv4 or IPv6 interfaces on a system. With the listen configuration option on the LDAP server, the hostname or the IPv4 or IPv6 address, along with the port number, can target one or multiple IPv4 or IPv6 interfaces on a system.
- **Persistent search**. The LDAP server provides an event notification mechanism for applications, directories and meta directories that need to maintain a cache of directory information or to synchronize directories when changes are made to an LDAP directory. Persistent search will allow these applications to be notified when a change has occurred.
- **ibm-entryuuid attribute**. The LDAP server now generates a unique identifier for any entry that is created or modified and does not already have a unique identifier assigned. The unique identifier is stored in the ibm-entryuuid attribute. The ibm-entryuuid attribute is replicated to servers that support the ibm-entryuuid attribute.
- **ibm-allMembers and ibm-allGroups**. The LDAP server now supports the querying of the members of static, dynamic, and nested groups in an LDBM backend by using the ibm-allMembers operational attribute. The LDAP server also supports the querying of the static, dynamic, and nested groups that a user belongs to with the ibm-allGroups operational attributes.
- **Remote services**. The LDAP server includes a set of remote services that enable access to RACF auditing and authorization, giving you more flexibility in your auditing and authorization management systems.

LDAP support introduces new messages, commands, and utilities. Commands and utilities introduced are:

- LDAPSRV (initializes the LDAP server)
- SMSG interface to the LDAP server (for server administration)
- gskkyman
- gsktrace
- ldapcompare
- ldapdelete
- ldapadd
- ldapmodify
- ldapmodrdn
- ldapsearch

The support includes a new manual for server administration: *z/VM: TCP/IP LDAP Administration Guide*.

[V5.3] OSA-Express2 IEEE 802.3ad Link Aggregation

z/VM virtual switch controlled (VSWITCH-controlled) link aggregation (IEEE 802.3ad) is designed to allow you to dedicate an OSA port to the z/VM operating system when the port is participating in an aggregated group. Link aggregation (trunking) is designed to allow you to combine multiple physical OSA-Express2 ports into a single logical link for increased throughput and for nondisruptive failover in the event that a port becomes unavailable.

The OSA-Express2 adapters that comprise the link aggregation group must be connected to the same physical switch and LAN segment. If there is a connectivity problem with any of the OSA-Express2 devices within the group, the virtual switch will detect the failure and packet transmission will continue over the other devices in the group. Thus, this support provides the capability for a quicker failover for devices within the group. In the event of a failure in which the entire link aggregation group is unavailable, the ability to specify a failover OSA-Express2 adapter will be supported. Similar to the existing failover support in z/VM V5.2, failover devices can be specified with the virtual switch definition. Only a single failover device will be deployed at a time, although multiple devices can be specified for the failover capability. For optimum connectivity, the failover devices should be connected to a separate physical switch.

z/VM Link Aggregation support will be provided as follows:

- Only for Layer 2 virtual switches (that is, defined with the ETHERNET option and using guest hosts that support Layer 2)
- Only on z9 EC and z9 BC with the required OSA-Express2 link aggregation function available with the May 11, 2007 level of the Licensed Internal Code (LIC)
- Allow up to 8 OSA-Express2 adapters to be aggregated per virtual switch
- Allow failover to a single OSA-Express2 adapter in the event of failure of the link aggregation group

[V5.3] ROUTED and BOOTP Servers Withdrawn

The ROUTED and BOOTP servers have been removed from z/VM V5.3. MPROUTE is the only dynamic routing server supported by TCP/IP for z/VM FL530. If you are currently running ROUTED on your system, you must migrate to the MPROUTE server before migrating to z/VM V5.3.

[V5.3] RSCS Networking for z/VM

IBM RSCS Networking for z/VM, function level 530, is offered as an optional feature of z/VM V5.3. It is the successor to RSCS V3.2 (5684-096).

RSCS Networking for z/VM includes an optional facility called RSCS dynamic authorization (RSCSAUTH), which allows you to authorize users for RSCS commands and links without updating the RSCS configuration file. Therefore a restart of the RSCS server is not required. RSCSAUTH maintains the authorization data in a set of tables. You can authorize a user for any set of commands and links. You can add, update, or remove authorizations simply by issuing RSCSAUTH commands to update the tables. These capabilities provide you with much greater flexibility in authorizing users for RSCS tasks.

The RSCS documentation has been extensively revised and updated. Information about the RSCS Data Interchange Manager, formerly provided in a separate publication, has been merged into the other RSCS documents. The RSCS Networking documents are:

- *z/VM: RSCS Networking Diagnosis*, GC24-6151
- *z/VM: RSCS Networking Exit Customization*, SC24-6152
- *z/VM: RSCS Networking Messages and Codes*, GC24-6153
- *z/VM: RSCS Networking Operation and Use*, SC24-6154
- *z/VM: RSCS Networking Planning and Configuration*, SC24-6155

The standalone RSCS V3.2 product is planned to be withdrawn from marketing effective May 26, 2008.

[V5.3] SMTP Enhancements

In previous releases, the method used by the SMTP server to resolve host names was based on the definition of the name servers in the TCPIP DATA file. If name servers were defined with the NSINTERADDR statement, then host name resolution was performed by the name servers. If name servers were not defined, host name resolution was performed using the site tables created by the HOSTS LOCAL file (at that time, support for the use of the ETC HOSTS file was not available).

With this release, the method used by the SMTP server to resolve host names (either domain name server or local site tables, or both), and the order in which these methods are used, are specified using the DOMAINLOOKUP statement in the TCPIP DATA configuration file. SMTP host name resolution now uses the domain specified with the DOMAINORIGIN statement; SMTP does not use the domains specified on the DOMAINSEARCH statement. In addition, if local site tables are used for name resolution, SMTP uses the ETC HOSTS file, if present, or the site tables created by the HOSTS LOCAL file, if ETC HOSTS is not present.

[V5.3] SSL Server Enhancements

The SSL server included in TCP/IP for z/VM, level 530, supports the following Linux distributions:

- Novell SUSE Linux Enterprise Server (SLES) 9 Service Pack 3 (64-bit)
- Novell SUSE Linux Enterprise Server (SLES) 9 Service Pack 3 (31-bit)
- Red Hat Enterprise Linux (RHEL) AS 4 Update 4 (64-bit)
- Red Hat Enterprise Linux (RHEL) AS 4 Update 4 (31-bit)

Attention: Older levels of the SSL server are not compatible with TCP/IP for z/VM, level 530.

The z/VM SSL server has been enhanced to allow the host Linux guest system to remain active after a critical error is encountered during server operations.

SSL server and administrative support has been added that will accommodate the renewal of a previously acquired Certifying Authority (CA) certificate.

Also, the SSLADMIN command has been enhanced to:

- Allow the specification of the number of days that a self-signed certificate is valid

- Improve the management of the SSL server LOG files, by providing the ability to:
  - Maintain log information in a file named other than SSLADMIN LOG
  - Specify a maximum size to be established for the SSL server log
  - Purge log information accumulated by the SSL server

[V5.3] TRACERTE Authorization Enhancement

The level of authorization needed to execute TRACERTE has changed. The TRACERTE command no longer requires OBEY authority. The TCP/IP administrator may want to review which users need to be in the OBEY list.

[V5.3] VIPA Support for IPv6

Virtual IP Address support in the TCP/IP stack has been extended in z/VM V5.3 to support IPv6 addresses. It is now possible to enable and configure a virtual device for IPv6, as well as to associate real IPv6-capable network adapters with a specific IPv6 virtual link for determining the source address used in outgoing packets. Support for VIPA is designed to improve the capability of the TCP/IP stack to maintain connections in the event that a real network device fails.

[V5.3] Virtual Network Management Improvements

z/VM V5.3 establishes a method for providing Simple Network Management Protocol (SNMP) data for virtual networking devices. Specifically, it provides an SNMP subagent that runs in a separate virtual machine from the SNMP agent and extends the functionality of the agent by supporting a specific set of Management Information Base (MIB) variables. A preconfigured subagent and exit routine are provided in z/VM V5.3 to supply bridge Management Information Base (BRIDGE-MIB) data, as documented in RFC 1493, for the z/VM virtual switch. This subagent, through the use of a Network Management System client, can acquire BRIDGE-MIB data for the z/VM virtual switch. In addition, this support provides a programming interface to obtain information about virtual networks.

[V5.3] Virtual Switch and Guest LAN Usability Enhancements

z/VM V5.3 provides usability enhancements for the virtual switch and guest LAN environments including:

- Enhanced ease-of-use for Virtual LAN (VLAN) and promiscuous mode configuration changes

  Changes to the authorized VLAN ID (VID) set and to promiscuous mode authorization are now effective immediately instead of requiring a revoke, a grant, and an uncouple/couple in order for the changes to take effect.
- New capability to configure a native VLAN ID

  This support provides the ability to specify a native VLAN identifier for untagged traffic and a default VLAN identifier for guest ports. The DEFINE VSWITCH command now supports the specification of a native VLAN identifier.
- New virtual NIC monitor domain

  Existing counts maintained for the virtual NIC, such as inbound packets, outbound bytes, and frame counts per MAC/VLAN, are now included in records in a new Virtual Network monitor domain. These new monitor records provide data for a virtual NIC that is coupled to any guest LAN or virtual switch.

[V5.3] VSWITCH Port Isolation Support

z/VM V5.3 supports a new port isolation security mechanism that provides the ability to restrict guest-to-guest communications within a virtual switch. In support of security zones, the virtual switch sends all guest port traffic over the OSA port for filtering through an external firewall or switch mechanism.

[V5.4] LDAP Server Upgrade

The LDAP server and utilities have been upgraded to a level equivalent to z/OS 1.10. In addition, the LDAP server supports:

- Server plug-ins, software modules that extend the functionality of your directory server. For information about using server plug-ins, see *z/VM: TCP/IP Programmer's Reference*. For details on the plug-in application service routines, the following publication has been added to the z/VM library:
  - *z/OS: IBM Tivoli Directory Server Plug-in Reference for z/OS*
- RACF change logging and password/phrase enveloping, enabling bidirectional propagation of security data between z/VM and z/OS.

[V5.4] Path MTU Discovery

Path MTU discovery is designed to decrease network resource waste by allowing the z/VM TCP/IP server to dynamically discover the MTU of a given IPv4 or IPv6 communication path.

- For IPv4 communication paths, traffic will be fragmented once only, at the source of the datagram using the smallest MTU. Path MTU discovery is optional and will be disabled by default for IPv4 links.
- For IPv6 communication paths, traffic will be fragmented by the source host. Path MTU discovery cannot be disabled; and therefore, is enabled for IPv6 links.

The following changes have been made to support the use of path MTU discovery:

- TCP/IP commands:
  - (**IPv4 only**) IFCONFIG has been updated to include two new operands; PATHMTU and NOPATHMTU to enable or disable path MTU discovery.
  - NETSTAT GATE has been updated to include two new flags; one indicates if the MTU was modified by path MTU discovery for a given route and the other if the route was created as a result of path MTU discovery.
  - (**IPv4 only**) NETSTATE DEVLINKS has been updated for all non-VIPA devices to display path MTU discovery status.
- TCP/IP configuration file statements:
  - (**IPv4 only**) The ASSORTEDPARMS statement has been updated to include the PATHMTU operand to enable path MTU discovery by default for links which have not explicitly been configured.
  - (**IPv4 only**) Various LINK statements have been updated to include two new operands; PATHMTU and NOPATHMTU to enable or disable path MTU discovery on a link-by-link basis.
  - The PATHMTUAGE statement has been added to specify how long (in minutes) path MTU discovery information will be retained for a given route.

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

[V5.4] Virtual Switch Networking Management

z/VM V5.4 has improved ease of use in managing virtual networks:

- Tracking of virtual switch traffic at the VLAN level
- Subsetting of QUERY VSWITCH and QUERY LAN commands
- Command line interface to generate SNMP Trap messages

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

[V5.4] TCP/IP OSD Layer 2

z/VM V5.4 provides layer 2 support for an OSA-Express adapter. IPWIZARD has been updated to allow users to specify the transport type (layer 2 or layer 3) for QDIO interfaces. IFCONFIG has been updated to allow the transport type (IP or ETHERNET). A new interface operand is added. IFCONFIG has been updated to display the transport type for QDIO Ethernet interfaces and the MAC address for QDIO Ethernet interfaces in layer 2 mode. The SET NIC command has been updated to allow a MAC ID to be associated with a virtual device number. QUERY NIC has been updated to display information including MAC address associated with real networking devices. The QUERY OSA command has been updated to display the MAC address associated with a virtual device number.

[V5.4] TCP/IP Telnet IPv6 Support

The TCP/IP Telnet server and client have been upgraded to support IPv6 addresses. This upgrade provides new IPv6 versions for the Pascal APIs that Telnet uses to establish connections. For more information, see *z/VM: TCP/IP User's Guide*.

[V5.4] TCP/IP Support Withdrawal

The following functions are not supported by TCP/IP for z/VM, Level 540:

- Network Database (NDB)
- SNALINK
- Trivial File Transfer Protocol (TFTP) server

  Note: The TFTP client is still supported.
- X.25

Information about these functions has been removed from the z/VM V5.4 publications and HELP files have been removed from the z/VM HELP Facility. (For a list of the deleted HELP files, see "HELP Facility" on page 395.)

This satisfies the statement of direction made in the Software Announcement dated February 6, 2007.

[V5.4] CMS-Based z/VM SSL Server

The z/VM SSL server does not require a Linux distribution or Linux skills to maintain. The CMS-based SSL server might enable encryption services to be deployed more quickly and can help make installation, service, and release-to-release migration simpler. Other enhancements to the CMS-based SSL server include:

- Network-free SSL server administration

  The SSL server can be managed without requiring a network connection between the SSL server administrator and the SSL server.
- Encryption and decryption engine

The SSL server uses z/OS V1.10 System SSL technology for encryption, decryption, and certificate management.

- Certificate-management services

  The System SSL **gskkyman** utility is used to manage the SSL server certificate database. Services available for the SSL server include certificate renewal, certificate signing, and certificate exportation with or without the private key. The **gskkyman** application also manages certificates for the z/VM LDAP server.

Additional functional changes associated with the CMS-based SSL server include:

- Cipher suite changes

  Several cipher suites, including suites that provide 128-bit and 256-bit AES encryption, have been added. Two ciphers – RC4_EXP1024_56_SHA and DES_EXP1024_56_SHA – have been removed. All other previously supported cipher suites have been renamed to more closely match specifications in RFCs 2246 and 4346.

- Certificate database stored in the z/VM byte file system (BFS)

  The z/VM SSL server references a certificate database stored in a BFS subdirectory, rather than one that resides in a Linux environment.

- DTCPARMS configuration changes

  Additional or different DTCPARMS file configuration tags and parameters are used for configuration of the SSL server, including the :Admin_ID_List. tag (to grant administrative authority to privileged users) and the VMSSL command KEYFILE parameter (a required parameter used to specify the location of the certificate database).

- FTP CCC subcommand support

  The SSL server will now support moving an FTP control connection from a secure state to a clear state through the CCC subcommand.

- GSKADMIN user ID

  The GSKADMIN virtual machine has been added to the User Directory. This user ID is authorized to perform certificate management operations through **gskkyman** for any SSL server database specified, as well as administrate a running server (given the appropriate TCP/IP configuration).

- Packaging, installation and service

  The SSL server is installed and serviced through the same means as other CMS-based TCP/IP servers; to this end, the TCP/IP installation user ID 493 minidisk has been removed, and installation of an updated RPM file is no longer necessary.

- Stack-to-server communication

  Administrative command communications now are performed without use of the TCP/IP stack server. Thus, the server administrative port (previously defined at port number 9999) is no longer used and has been removed from the TCP/IP server configuration and ETC SERVICES sample files.

For general information about setting up the new z/VM SSL server, see Configuring the SSL Server in *z/VM: TCP/IP Planning and Customization*. For more detailed information regarding the migration of certificates from a prior level z/VM SSL server, as well as supplementary SSL server configuration information, see IBM: TCP/IP for z/VM SSL server configuration information and requirements (www.vm.ibm.com/related/tcpip/vmsslinf.html).

[V5.4] OSA-Express QDIO Data Connection Isolation

z/VM supports port isolation security that provides the ability to restrict guest-to-guest communications within a z/VM virtual switch (VSWITCH) and between other partitions sharing the OSA-Express adapters used by the VSWITCH. Virtual switch port isolation and QDIO data connection isolation can help you design virtual networks that adhere to strict traffic-separation policies.

For more information on QDIO data connection isolation, see the IBM System z10 Hardware Announcements dated October 21, 2008.

Traffic isolation on shared OSA-Express adapters is available for OSA-Express2 features (on a System z9 or z10) and OSA-Express3 features (on a System z10 only) and is exclusive to CHPID type OSD with the following minimum MCLs:

- OSA-Express2 on z9 requires G40946.008
- OSA-Express2 on z10 requires N10953.002
- OSA-Express3 on z10 requires N10959.004 and N10967.055

The following CP functions have been updated:

- MODIFY VSWITCH system configuration statement
- QUERY CONTROLLER command
- QUERY VSWITCH command
- SET VSWITCH command
- DIAGNOSE code X'26C'
- Monitor record Domain 6 Record 21 - MRIODVSW - Virtual Switch Activity

For more information, see:

- *z/VM: Connectivity*
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Planning and Administration*
- *z/VM: CP Programming Services*

[V6.1] Prefetch Guest Data into Processor Cache

Guest LAN and virtual switch support has been updated in z/VM V6.1 to use cache prefetch capabilities that are exclusive to the IBM System z10 servers and later in order to give the hardware hints about likely memory access patterns. This enables the hardware to prefetch data into the processor cache so that the processor does not have to wait for data to be moved from main memory. Avoidance of a “cache miss” might help improve the performance of heavy guest-to-guest streaming workloads.

[V6.1] SSL Server Enhancements

z/VM provides enhancements to the z/VM SSL server that improve the ability to provide concurrent secure connectivity by increasing the server's overall capacity and decreasing the amount of system resources the server requires:

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

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

[V6.1] SSL Server Upgrade and FIPS Support

The z/VM SSL server is upgraded to use z/OS V1.11 System SSL technology and is enhanced with support for the Federal Information Protection Standard (FIPS) 140-2. FIPS is a government required procedure for authenticating the security of a certificate database that requires signing of the database and using separate cipher suites for traffic encryption and decryption.

For more information, see:

- *z/VM: TCP/IP Planning and Customization*
- *z/VM: TCP/IP User's Guide*.

[V6.1] Support for IBM zEnterprise System Networks

z/VM enables access to the IBM zEnterprise System intraensemble data network (IEDN) and intranode management network (INMN) through z/VM's real and virtual networking capabilities. As a result, the deployment and management of z/VM network topology is integrated into the zEnterprise environment and can be managed by the IBM z Unified Resource Manager.

The QUERY NIC command is new for this support.

The following functions have been updated:

- System configuration statements:
  - DEFINE VSWITCH
  - MODIFY LAN
  - MODIFY VSWITCH
  - VMLAN
- User directory statements:
  - NICDEF
- CP commands:
  - COUPLE
  - DEFINE NIC
  - DEFINE VSWITCH
  - QUERY CONTROLLER
  - QUERY LAN
  - QUERY VIRTUAL NIC
  - QUERY VMLAN
  - QUERY VSWITCH
  - SET LAN
  - SET NIC
  - SET VMLAN
  - SET VSWITCH
- DIAGNOSE codes:
  - X'26C'
- CP monitor records:
  - Domain 1 Record 19 - MRMTRQDC - QDIO Device Configuration
  - Domain 6 Record 23 - MRIODVSR - Virtual Switch Recovery
  - Domain 6 Record 25 - MRIODQDA - QDIO Device Activation
  - Domain 6 Record 27 - MRIODQDD - QDIO Device Deactivation

  - Domain 8 Record 1 - MRVNDSES - Virtual NIC Session Activity
  - Domain 8 Record 2 - MRVNDLSU - Virtual NIC Link Up
  - Domain 8 Record 3 - MRVNDLSD - Virtual NIC Link Down
- TCP/IP configuration statements:
  - VSWITCH CONTROLLER
- TCP/IP commands:
  - NETSTAT

[V6.2] Additional IPv6 Support

Support for IPv6 has been included in:

- z/VM TCP/IP FTP client and server
- z/VM TCP/IP SMTP client and server

  Attention: If you are using SMTP assembler exits, you need to modify those exits to work with the updated SMTP. Because the address fields have been expanded, the offset of the return code field in the parameter list has changed. For more information, see *z/VM: TCP/IP Programmer's Reference*.
- CMS NOTE and SENDFILE commands

[V6.2] IPFORMAT Supports Conversion to PCAP Format

TRSOURCE traces of both TYPE GT and TYPE LAN can now be exported to the popular PCAP format. For details see the description of the new FORMAT option for IPFORMAT in *z/VM: TCP/IP Diagnosis Guide*.

[V6.2] LDAP Server Upgrade

The LDAP server and utilities have been upgraded to a level equivalent to z/OS 1.12, which includes the following new functions:

- Enhanced replication
- CDBM backend
- Access to RACF resource profiles
- Password policy
- Binding with SDBM using password policy
- Using access control
- Schema updates
- Salted SHA
- Activity log enhancements
- LDAPEXOP utility.

[V6.2] MPROUTE Enhancements

The MPROUTE server has been enhanced and upgraded from z/OS level 1.8 to z/OS level 1.12. The following functions have changed:

- MPRPOUTE now supports RFC 4191 and RFC 5175.
- The MPROUTE configuration file now supports INCLUDE statements.
- MPROUTE now reports and helps prevent futile neighbor state loops.
- The SMSG command has been updated to include DELETED, ACTIVATE, and SUSPEND keywords.
- The ROUTERADV statement now includes a PREFERENCE value that determines whether router advertisements are sent with a HIGH, MEDIUM, or LOW preference value.

- The NETSTAT CONFIG command now includes the ROUTERADV option. The NETSTAT GATE command and the NETSTAT CONFIG HELP command now include new output fields.
- Numeric interface names containing a decimal point (for example, 123.456) are no longer allowed when using MPROUTE.

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

[V6.2] Multiple Access Ports per Guest

This support allows a z/VM guest to have multiple unique access ports connected to the same virtual switch (VSWITCH) instance, where each access port has its own attributes (promiscuous and OSDSIM authority) and VLAN ID. This support also provides a way to associate guest NICs and VSWITCH ports, and allows specific ports to be designated for use. These changes better facilitate the use of SNMP to monitor VSWITCH status because up-front mapping of systems to virtual switch ports now is possible, and more closely resembles techniques that can be applied to physical switches.

In addition, support for multiple access ports per guest can help reduce system administrative complexity, because servers no longer need to be configured to provide connectivity to multiple VLANs, and z/VM guests no longer must be configured to be VLAN aware. By allowing multiple unique access ports for a guest, these configuration requirements are eased, because the guest is not aware of the VLANs (tagging is performed by the virtual switch). This support can also reduce the potential for configuration errors that result in servers being unable to communicate on a LAN segment.

[V6.2] Query OSA Address Table

This support provides TCP/IP with the capability to display the content of the OSA Address Table (OAT) via the QDIO interface for both OSA cards and VSWITCH controllers. The NETSTAT command now includes the OSAINFO option. See *z/VM: TCP/IP User's Guide*.

[V6.2] SSL Server Upgrade and Implementation Changes

The z/VM SSL server is upgraded to use z/OS V1.12 System SSL technology. In addition, only a multiple SSL server "pool" is defined as part of the z/VM version 6 release 2 System Deliverable. A single SSL server (SSLSERV) no longer is provided. However, continued use of a single-instance server still is possible, and remains supported.

For more information, see:

- *z/VM: TCP/IP Planning and Customization*
- *z/VM: TCP/IP User's Guide*.

[V6.2] TCP/IP Support Withdrawal

The following functions have been removed from TCP/IP for z/VM, Level 620:

- z/VM Kerberos servers (VMKERB and ADMSERV)

  The Kerberos protocol is no longer supported.
- z/VM Domain Name System server (NAMESRV)

  The Domain Name System (DNS) protocol is still supported using non-z/VM servers.

Information about these z/VM servers has been removed from the z/VM V6.2 publications, and the corresponding HELP files have been removed from the z/VM HELP Facility. (For a list of the deleted HELP files, see "HELP Facility" on page 395.)

This satisfies the statement of direction made in the Software Announcements dated August 5, 2008, July 7, 2009, October 20, 2009, and July 22, 2010.

[V6.2] Virtual Networking Support for an SSI Cluster

This support extends the virtual switch's existing layer 2 MAC service to coordinate MAC address assignments with all active members of an SSI cluster. The following existing and new functional areas have been modified or added by SSI support:

- Existing assigned MAC table logic.
- New SSI cluster *NETWORK Communication Manager.
- Add MAC address protection to SSI cluster.
- Network Device Equivalency Determination for Live Guest Relocation.

[V6.2] *VMEVENT Enhancements

This support provides the ability for the *VMEVENT system service to indicate asynchronous completion of networking commands (such as DEFINE VSWITCH, SET VSWITCH, and SET PORT GROUP), thereby allowing subsequent commands to be issued successfully. A new class (Class 4, for networking events) has been added to *VMEVENT. See *z/VM: CP Programming Services*.

[V6.2] Virtual Switch HiperSockets Bridge

With the PTFs for APARs VM65042 and PM46988, the z/VM virtual switch is enhanced to transparently bridge a guest virtual machine network connection on a HiperSockets LAN segment. This bridge allows a single HiperSockets guest virtual machine network connection to also directly communicate with the following:

- Other guest virtual machines on the z/VM virtual switch
- External network hosts through the virtual switch OSA UPLINK port.

The following system configuration statements have been updated for this support:

- DEFINE VSWITCH
- MODIFY VSWITCH

The following CP commands have been updated for this support:

- DEFINE CHPID
- DEFINE VSWITCH
- DETACH
- QUERY CONTROLLER
- QUERY OSA
- QUERY VSWITCH
- SET CPTRACE
- SET NIC
- SET PORT GROUP
- SET VMLAN
- SET VSWITCH

The following TCP/IP functions have been updated for this support:

- TCPIP server VSWITCH CONTROLLER configuration statement

- NETSTAT DEVLINKS command

The following DIAGNOSE code has been updated for this support:
- DIAGNOSE Code X'26C'

The following monitor records have been added for this support:
- Domain 6 Record 33 - MRIODBPA - Virtual Switch Bridge Port Activation
- Domain 6 Record 34 - MRIODBPD - Virtual Switch Bridge Port Deactivation
- Domain 6 Record 35 - MRIODBPS - Virtual Switch Bridge Port Activity

The following monitor records have been updated for this support:
- Domain 8 Record 1 - MRVNDSES - Virtual NIC Session Activity
- Domain 8 Record 2 - MRVNDLSU - Virtual NIC Guest Link State - Link Up
- Domain 8 Record 3 - MRVNDLSD - Virtual NIC Guest Link State - Link Down

The following trace codes have been added for this support:
- 7402 - Virtual Switch Bridge Port Outbound Transfer
- 7403 - Virtual Switch Bridge Port Completion Event
- 7404 - Virtual Switch Bridge Port Inbound Transfer

[V6.2] Client Certificate Validation for z/VM SSL Server

With the PTF for APAR PM52716, the z/VM SSL server is enhanced to accommodate client certificate validation as part of handshaking for dynamic SSL connections. This provides a means for Transport Layer Security (TLS) handshaking to incorporate mutual authentication of certificates before the establishment of a secure connection.

z/VM V6.2 with APAR PM52716 is designed to meet the Common Criteria certification of z/VM V6.1. For more information, see IBM z/VM Security and Integrity resources (www.ibm.com/vm/security/).

The following z/VM TCP/IP functions have been updated for this support:
- New CLIENTCERTCHECK option on the INTERNALCLIENTPARMS configuration statement for the TCPIP server
- Client Cert Check added to the Internal Client Settings in the response for the NETSTAT CONFIG PARMS command

For more information, see:
- *z/VM: TCP/IP Planning and Customization*
- *z/VM: TCP/IP User's Guide*

[V6.3] Live Guest Relocation Support for Port Based Virtual Switches

This support builds on the existing live guest relocation framework to:
- Allow relocation of a port based interface
- Prevent relocation of an interface that will clearly be unable to establish network connectivity
- Adjust the destination virtual switch configuration (when possible) by inheriting virtual switch authorization from the origin.

The support adds the checks that are needed to safely relocate a guest with a port based virtual switch interface. Eligibility tests are performed before relocation to determine if the destination network is equivalent.

The support improves the preexisting base network eligibility test by testing other attributes that would prevent the guest from establishing connectivity. Without the improved checks, if you accidentally mislabel two networks as equivalent, a guest might be relocated to the destination node but find it impossible to establish network connectivity, and the reason for the loss of connectivity might not be clear.

[V6.3] MPROUTE Enhancements

The MPROUTE server has been enhanced and upgraded from z/OS level 1.12 to z/OS level 1.13.

[V6.3] SSL Server Upgrade and TLS 1.2 Support

The z/VM SSL server has been upgraded to z/OS V1.13 equivalency. This upgrade includes support for Transport Layer Security (TLS) protocol, version 1.2, which provides support for SHA-256 certificates. A new PROTOCOL operand on the VMSSL command allows the system administrator to enable and disable SSL and TLS protocols for cryptographic use in the operation of the SSL server.

The z/VM SSL server also has been enhanced to support secure IPv6 SSL connections. The following servers and clients have added support for secure IPv6 connections:

- TELNET server and client
- FTP server and client
- SMTP server

Furthermore, NETSTAT IDENT SSL and SSLADMIN have been enhanced to display secure IPv6 connections.

For more information, see:

- *z/VM: TCP/IP User's Guide*
- *z/VM: TCP/IP Planning and Customization*

[V6.3] TCP/IP Support Withdrawal

The following functions have been removed from TCP/IP for z/VM, Level 630:

- A220 and CLAW device drivers.
- z/VM Dynamic Host Configuration Protocol Daemon (DHCPD)

  The DHCP Daemon is no longer supported.
- z/VM Line Printer Daemon (LPD)

  The Line Printer Daemon (LPD) protocol is still supported using non-z/VM LPD servers, and LPD related tools are still shipped for this purpose. An LPD link driver is still provided in the z/VM RSCS feature.

Information about these unsupported z/VM servers has been removed from the z/VM V6.3 publications, and the corresponding HELP files have been removed from the z/VM HELP Facility. (For a list of the deleted HELP files, see "HELP Facility" on page 395.)

Any customer who chooses to continue using an unsupported server will need to copy the appropriate server/class definitions from IBM DTCPARMS to their system/node ID DTCPARMS file and maintain them there.

[V6.3] Virtual Switch Recovery and Stall Prevention

A new SET VSWITCH UPLINK SWITCHOVER command provides a fast mechanism to switch an uplink port's network connection from its current device to one of its configured backup devices. The use of this new command is less disruptive to virtual switch external network connectivity when there is a need to apply service to an OSA-Express feature.

Additionally, support is added to automatically enable or disable the z/VM Missing Interruption Handler (MIH) function for an OSA-Express or HiperSockets read, write, and data uplink device. MIH is used to detect and recover (failover) from network connectivity problems due to unresponsive or broken hardware.

The following CP command has been updated for this support:
- SET VSWITCH

The following DIAGNOSE code has been updated for this support:
- DIAGNOSE Code X'26C'

The following monitor record has been updated for this support:
- Domain 6 Record 22 - MRIODVSF - Virtual Switch Failover

[V6.3] Virtual Switch Support for Virtual Edge Port Aggregator (VEPA) Mode

Virtual Edge Port Aggregator (VEPA) is part of the IEEE 802.1Qbg standardization effort and is designed to reduce the complexities associated with highly virtualized deployments such as hypervisor virtual switches bridging many virtual machines. VEPA provides the capability to take all virtual machine traffic sent by the server and send it to an adjacent network switch.

The following commands have been updated for this support:
- SET VSWITCH
- QUERY VSWITCH
- QUERY CONTROLLER
- NETSTAT OSAINFO

The following configuration statement is changed:
- MODIFY VSWITCH

The following monitor records have been updated:
- D6R21 - MRIODVSW - Virtual Switch Activity
- D6R22 - MRIODVSF - Virtual Switch Failure
- D6R35 - MRIODBPS - Virtual Switch Bridge Port Activity

The following diagnose codes are changed:
- Diagnose x'26C' - Access Certain System Information, Subcode x'10'
- Diagnose x'26C' - Access Certain System Information, Subcode x'20'
- Diagnose x'26C' - Access Certain System Information, Subcode x'24'

[V6.3] Duplicate IP Address Handling on a Virtual Switch

The IP (Layer 3) virtual switch has been updated with a change in the way a duplicate IP address is handled. When a guest coupled to an IP virtual switch registers its IP address, the virtual switch will forward this setting to the Uplink

for registration with the OSA-Express. The OSA-Express will verify that the same IP address is not already registered on the LAN segment. If it is, the OSA-Express rejects the request.

Prior to z/VM V6.3, the virtual switch marked this IP address as "Local", and the guest IP address was reachable only on the local simulated LAN segment. In z/VM V6.3, the IP virtual switch now disables the IP address and marks it as "Conflict". The guest IP address will no longer be reachable by guests on the simulated LAN segment.

This change allows the configuration problem to be highlighted more quickly. In the event that an existing configuration has been configured with a duplicate IP address for local connectivity within the virtual switch LAN, a unique IP address should be assigned, or the virtual switch should be configured without an Uplink port.

[V6.3] Unified Resource Manager Support Withdrawn

The IBM z Unified Resource Manager, first supported in z/VM V6.1, is no longer supported in z/VM V6.3. Therefore z/VM V6.3 cannot participate as a member of an ensemble.

[V6.3 APAR] Additional SSL Server Enhancements

With the PTF for APAR PM93363, the z/VM SSL server has been upgraded to enforce configurations that meet specific cryptographic standards. This includes continued support for FIPS 140-2 compliance as well as support for NIST SP 800-131A configurations, which mandates a minimum size of 2048 bits for asymmetric keys.

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

[V6.3 APAR] z/VM V6.3 System SSL Cryptographic Module Receives FIPS 140-2 Certification

With the PTF for APAR PI04999, the z/VM V6.3 System SSL module has been validated as conforming to the Federal Information Processing Standard (FIPS) 140-2. This industry-recognized cryptographic standard mandates modern digital key sizes and integrity checking for SSL and TLS operations. z/VM 6.3 System SSL is used by both the z/VM LDAP Server and z/VM SSL-TLS Server. This satisfied the statement of direction made in the IBM Software Announcement dated July 23, 2013.

[V6.3 APAR] Multi-VSwitch Link Aggregation Support

With the PTFs for APARs VM65583 and PI21053, z/VM provides support for Multi-VSwitch Link Aggregation on the IBM z13. This support allows a port group of OSA-Express features (OSA-Express4S or later) to span multiple virtual switches within a single z/VM system or between multiple z/VM systems. Sharing a Link Aggregation Port Group (LAG) with multiple virtual switches increases optimization and utilization of the OSA-Express when handling larger traffic loads. Higher adapter utilization protects customer investments, which is increasingly important as 10 Gigabit deployments become more prevalent. With this support, a port group is no longer required to be dedicated to a single virtual switch.

Systems management API support for Multi-VSwitch Link Aggregation requires APAR VM65670. See "[V6.3 APAR] Systems Management API Support for Multi-VSwitch Link Aggregation" on page 261.

Performance Toolkit support for Multi-VSwitch Link Aggregation requires APAR VM65528. See "[V6.3 APAR] Additional Performance Toolkit Enhancements" on page 219.

The following configuration statements have been updated for this support:
- DEFINE VSWITCH
- MODIFY PORT

The following CP commands have been updated for this support:
- ATTACH
- COUPLE
- DEFINE VSWITCH
- DETACH VSWITCH
- QUERY CONTROLLER
- QUERY OSA
- QUERY PORT
- QUERY (Real Device)
- QUERY VMLAN
- QUERY VSWITCH
- SET CPTRACE
- SET PORT GROUP
- SET VMLAN
- SET VSWITCH

The following monitor record has been added:
- Domain 8 Record 4 - MRVNDGLB - Global Virtual Switch Activity

The following monitor records have been updated:
- Domain 6 Record 21 - MRIODVSW - Virtual Switch Activity
- Domain 6 Record 23 - MRIODVSR - Virtual Switch Recovery
- Domain 6 Record 33 - MRIODBPA - Virtual Switch Bridge Port Activation
- Domain 8 Record 1 - MRVNDSES - Virtual NIC Session Activity
- Domain 8 Record 2 - MRVNDLSU - Virtual NIC Guest Link State - Link Up
- Domain 8 Record 3 - MRVNDLSD - Virtual NIC Guest Link State - Link Down

The following diagnose code has been updated:
- DIAGNOSE Code X'26C' - Access Certain System Information

The following CP System Service has been updated:
- *VMEVENT - new Class 4 records

For more information, see:
- *z/VM: Connectivity*
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Planning and Administration*
- *z/VM: CP Programming Services*

[V6.3 APAR] System SSL Upgrade and AES GCM for the TLS/SSL Server

With the PTFs for APARs PI40702 (TCP/IP and System SSL), VM65717 (CMS), and VM65718 (Language Environment), the z/VM System SSL cryptographic module is upgraded to z/OS V2.1 equivalency. This upgrade includes support for AES Galois/Counter Mode (GCM). These ciphers are available only when using TLS 1.2

in the TLS/SSL Server. Updates have also been made to the TLS/SSL Server's MODE NIST-800-131a to allow for DSA certificates of key length 2048.

Additionally, changes have been made to the default protocol levels and cipher suites presented by the TLS/SSL Server on initialization. SSLv3 is no longer enabled by default. NULL, RC4, and EXPORT ciphers are disabled by default. These mechanisms are available for legacy purposes only, and their use is not encouraged for production environments. They can be re-enabled using the PROTOCOL or ENABLE operands.

Finally, the GSKKYMAN application has been enhanced to add support for Elliptic Curve Cryptography (ECC) certificates and keys. ECC support for the TLS/SSL Server and LDAP Server is planned for a future release.

For more information, see:

- *z/VM: TCP/IP Planning and Customization*
- *z/VM: TCP/IP User's Guide*

[V6.3 APAR] Encryption of TCPNJE Connections

With the PTFs for APARs VM65788 and PI56474, 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.

For more information, see:

- *z/VM: RSCS Networking Operation and Use*
- *z/VM: RSCS Networking Planning and Configuration*
- *z/VM: RSCS Networking Messages and Codes*
- *z/VM: TCP/IP Programmer's Reference*

[V6.4] Client/Server TLS/SSL Support

With the PTF for APAR VM65892, new ioctl() commands are added to allow clients and servers that communicate through the TCP/IP stack to have their data flow through the TLS/SSL server. The new ioctls are supported for C programs as well as assembler programs which use IUCV.

For more information, see *z/VM: TCP/IP Programmer's Reference*

[V6.4] z/VM LDAP Server and Client Utilities z/OS V2.2 Equivalency

The z/VM LDAP server and utilities are upgraded to a level equivalent to z/OS V2.2. The updated z/VM support includes:

- Paged or 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

For more information, see:

- *z/VM: TCP/IP Planning and Customization*
- *z/VM: TCP/IP LDAP Administration Guide*
- *z/VM: TCP/IP User's Guide*
- *z/VM: TCP/IP Programmer's Reference*
- *z/VM: TCP/IP Messages and Codes*

[V6.4] LOGON Security Improvement

The CP LOGON command has been updated to prevent user ID enumeration without authenticating to the system. This means that 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.

For more information, see *z/VM: CP Commands and Utilities Reference*.

[V6.4] z/VM MPROUTE Server z/OS V2.2 Equivalency

The z/VM MPROUTE server is upgraded to a level equivalent to z/OS V2.2. 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

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

[V6.4] Domain Name System (DNS) IPv6 Support

z/VM TCP/IP supports DNS servers that have IPv6 addresses.

For more information, see:

- *z/VM: TCP/IP User's Guide*
- *z/VM: TCP/IP Planning and Customization*
- *z/VM: TCP/IP Programmer's Reference*

[V6.4] 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. See "CP System Services" on page 368 for more details.

For more information, see:

- *z/VM: CP Programming Services*
- *z/VM: RACF Security Server Security Administrator's Guide*

[V6.4] 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.

The following CP commands have been updated:
- QUERY PORT
- QUERY VIRTUAL NIC
- QUERY VSWITCH
- SET VSWITCH

For more information, see *z/VM: CP Commands and Utilities Reference*.

[V6.4] SET DIALDROP Command

This support provides the ability to control whether devices DIALed to a virtual machine are dropped or stay connected when an implicit or explicit virtual machine reset occurs via the SET DIALDROP command.

The following CP command has been added for this support:
- SET DIALDROP

The following CP command has been updated for this support:
- QUERY SET

For more information, see *z/VM: CP Commands and Utilities Reference*.

[V6.4 APAR] Crypto Express APVIRT Support for z/VM TLS/SSL Server and LDAP/VM

With the PTF for APAR PI72106, the z/VM System SSL cryptographic library is updated to offload cryptographic operations to Crypto Express hardware associated with your IBM Z or LinuxONE hardware, which may improve performance. This support is intended for clear-key RSA operations. To enable this support, add the CRYPTO APVIRTUAL statement to the pertinent service virtual machine entries in the z/VM user directory.

For more information, see "Configuring the LDAP Server" and "Configuring the SSL Server" in *z/VM: TCP/IP Planning and Customization*.

[V6.4 APAR] Distributed IUCV Enhancements

With the PTF for APAR VM65872, the rules for Distributed IUCV CONNECT in a single system image (SSI) environment are revised.

This support changes the rules for IUCV CONNECT with DISTRIBUTE IUCV NO:
- IUCV CONNECT to USERID=*userid* with no target node previously failed if the invoker was a multiconfiguration virtual machine (IDENTITY) user. This is now allowed when *userid* is a single-configuration user in the SSI cluster. This fails if *userid* is an IDENTITY user because *userid* without a node does not uniquely identify the instance of the IDENTITY user in the cluster.
- IUCV CONNECT to USERID=*userid* TARGET=*nodeid* previously failed if *userid* or the invoker was an IDENTITY user. This is now allowed because *nodeid* provides the information needed to identify the target user.

This support also changes the rules for IUCV CONNECT with DISTRIBUTE IUCV TOLERATE or YES:

- IUCV CONNECT to USERID=*userid* with no target node previously failed if the invoker was a multiconfiguration (IDENTITY) user. This was an unintended side effect of the code that restricted IDENTITY users from connecting to a single-configuration user who might not be able to automatically connect back. This is now corrected, so an IDENTITY user can initiate the IUCV CONNECT to a single-configuration user.

This support also makes it easier for an administrator to change the Distributed IUCV policy for an SSI cluster. Previously, the Distributed IUCV policy within an active SSI cluster could be changed only by shutting down all members at the same time. As this new support is applied to each system, it will be possible for that member to join the cluster regardless of its Distributed IUCV configuration.

The following CP interfaces have been updated:
- DISTRIBUTE IUCV system configuration statement
- IUCV macro CONNECT function

[V6.4 APAR] NICDEF Security Controls

With the PTF for APAR VM65925, the NICDEF user directory statement is enhanced to provide a set of new operands referred to as Directory Network Authorization (DNA). With DNA, a system administrator can configure and consolidate a virtual NIC device and its network properties in a centralized location -- the z/VM user directory.

Operational differences between PORTBASED and USERBASED VSwitches have been eliminated with this support. A system administrator has the option to manage a VSwitch by user, by port number, or using a combination of the two methods. While the management of USERBASED and PORTBASED VSWITCHes is simplified, Live Guest Relocation of a guest connected to a VSwitch still requires the destination system to have a VSwitch with a PORTBASED or USERBASED designation matching that of the source system.

The following statements have has been updated:
- DEFINE VSWITCH configuration statement
- MODIFY VSWITCH configuration statement
- VMLAN configuration statement
- NICDEF directory statement

The following CP commands have been updated:
- COUPLE
- DEFINE VSWITCH
- LOGON
- QUERY LAN
- QUERY VMLAN
- QUERY VSWITCH
- SET VMLAN
- SET VSWITCH
- VMRELOCATE

The following diagnose code has been updated:
- DIAGNOSE Code X'26C' - Access Certain System Information

With the PTF for APAR VM65926, the Directory Maintenance Facility (DIRMAINT) NICDEF command is updated to provide new options:

- PORTNUMber *portnum*
- PORTType ACCESS | TRUNK
- VLAN *vidset*
- PROmiscuous | NOPROmiscuous

With the PTF for APAR VM65931, RACF and RPIDIRCT have been updated to support NICDEF enhancements.

For more information, see:

- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Messages and Codes*
- *z/VM: Connectivity*
- *z/VM: Directory Maintenance Facility Commands Reference*
- *z/VM: Directory Maintenance Facility Messages*
- *z/VM: Directory Maintenance Facility Tailoring and Administration Guide*
- *z/VM: RACF Security Server Security Administrator's Guide*

[V6.4 APAR] Firewall Friendly FTP

With the PTF for APAR PI80912, the z/VM TCP/IP FTP client is configured to determine the party initially responsible for initiating data connections, improve compliance with RFC 2428, and reflect the active/passive FTP mode and the EPSV4 setting.

The following FTP subcommands have been updated:

- LOCSITE
- LOCSTAT
- PASSIVE
- SENDPORT

The following FTP DATA file statement has been added:

- FWFRIENDLY

The following FTP DATA file statement has been updated:

- EPSV4

The following unnumbered message has been added:

- FWFRIENDLY

For more information, see:

- *z/VM: TCP/IP User's Guide*
- *z/VM: TCP/IP Messages and Codes*

System Administration and Operation

These topics describe changes that can affect how you manage the administration and operation of z/VM and guest operating systems.

Note: Enhancements in TCP/IP administration are described under "Connectivity and Networking" on page 125.

[V2.2] Year 2000 Support

VM/ESA V2.2 provides support for the year 2000 and beyond. CP and CMS commands that accept dates as input or display dates as output have been enhanced to support dates with 4-digit years. In addition to commands, CMS Pipelines and application program interfaces such as REXX and CSL routines have been enhanced to support dates with 4-digit years. Existing dates consisting of 2-digit years are resolved with correct 4-digit year information. A new CSL routine, DateTimeSubtract, has been added to assist users in modifying their user-written programs to work with 4-digit date support. The TOD clock support is also enhanced to aid in IPLing a VM system in the year 2000 or later.

The capability to set a default date format for CP and CMS commands is provided on a system-wide basis and also on a user (virtual machine) basis. A new system configuration statement, SYSTEM_DATEFORMAT, sets the system-wide default date format. A new directory control statement, DATEFORMAT, sets the default date format for a user (virtual machine). Two new CP commands, SET DATEFORMAT and QUERY DATEFORMAT, set and query the default date formats for the system and individual users. Default date formats for the FILELIST and RDRLIST commands can also be set with the CMS DEFAULTS command.

Range of Dates Supported for Files

Minidisk files and SFS files support dates only in the range of 01/01/1900 to 12/31/2099.

Files on a Back Level File Pool Server

4-digit years are not supported on a file pool server from a VM release prior to VM/ESA V2.2. The new system will interpret all dates associated with a file on a back level server as 19*yy*, where *yy* is the 2-digit year.

Remote Systems That Are Not Year 2000 Ready

If a remote system that is not Year 2000 ready accesses minidisks on the new system, all the 4-digit years for existing files will appear to be 19*yy*, even if they have been created after 1999.

If you copy a file from a remote system to a minidisk or SFS directory on the new system using the OLDDATE option, the 4-digit year for the file may not be correct if the remote system is not Year 2000 ready. Therefore, an incorrect date might be propagated onto your system.

If you use the new system to create a file on a minidisk on a remote system that is not Year 2000 ready, the new system will see the correct 4-digit year.

Profile for CMS Productivity Aids

If you have a pre-VM/ESA V2.2 XEDIT profile for FILELIST, RDRLIST, SENDFILE, or PEEK on your A-disk or in your search order accessed ahead of the system disk, some PF key functions may not work correctly with the new FULLDATE and ISODATE options. To ensure that the new date functions operate correctly, you should erase your old profiles and create new ones. The recommended method for customizing the operation of the PF keys for these commands (as well as other productivity aids that use profiles) is to build a user profile that first calls the profile from the system disk, followed by your customized changes. For more

information, see *z/VM: CMS Commands and Utilities Reference*.

[V2.2] CMS Productivity Aids Performance Enhancement

The CMS productivity aids FILELIST, RDRLIST, and PEEK have been rewritten from EXEC 2 to REXX and compiled. They have also been placed into the CMSINST logical saved segment, which has been moved above the 16 MB line. (See "[V2.2] HELPINST Saved Segment Replaced by HELPSEG and INSTSEG.") This will improve the performance of these commands and also aid in the servicing of these parts.

[V2.2] HELPINST Saved Segment Replaced by HELPSEG and INSTSEG

The default CMS installation saved segment (CMSINST) and the HELP logical saved segment are no longer defined in a single physical saved segment called HELPINST. CMSINST is now defined in the INSTSEG physical saved segment, which has been moved above the 16 MB line. HELP is now defined in the HELPSEG physical saved segment, which is still located below 16 MB. This change provides more room for help files in the HELP saved segment.

More of the CMS productivity aids (execs and XEDIT macros that reside on the S-disk) have been added to CMSINST. For a list of the contents of CMSINST, see the CMSINST LSEG file.

To allow CMSINST to be moved above 16 MB, execs included in CMSINST that were formerly written in EXEC or EXEC 2 have been rewritten in REXX. If you have user modifications to any of these rewritten execs, you must rewrite your modifications in REXX. You cannot add EXEC or EXEC 2 execs to CMSINST. If you have EXEC or EXEC 2 execs you want to provide in a logical saved segment, you must add them to a segment that is loaded below 16 MB.

[V2.2] Removable Media Services Free Drive Support

DIAGNOSE code X'254', Access Real Subsystem, allows the Removable Media Services (RMS) virtual machine to issue nondrive dependent I/O to a 3494 or 3495 tape library subsystem, without requiring the DFSMS/VM RMS machine to have a CP-attached tape drive. For more information, see *z/VM: CP Programming Services*.

[V2.2] SFS AUDIT Enhancement

You can now specify AUDIT OFF CLOSE *fn ft* or AUDIT CLOSE *fn ft* to create a file to capture a "snapshot" of the security audit trace file. This prevents overwriting the security audit trace file and allows you to get audit information while auditing remains on. For more information, see *z/VM: CMS File Pool Planning, Administration, and Operation*.

[V2.3] Additional Year 2000 Support

VM/ESA V2.3 provides extended support for year 2000 and beyond. A new CMS pipelines stage called DATECONVERT allows date conversion, validation, and windowing functions. It supports all of the REXX date formats, plus additional formats. It is a powerful front-end to the DateTimeSubtract CSL routine. There are two new data migration and testing aid commands: FIXCENT and FILESERV FIXCENT. These commands can be used to set the internal century information stored with minidisk and SFS files. Note that this is not needed as part of the general processing. There is a new conversion command that converts dates in the

user's NETLOG file called NETLCNVT as well as message display enhancements to the NETDATA command. When IPLing VM, the 'SET DATE' prompt has been enhanced to accept 4-digit year dates.

For more information, see:

- *z/VM: CMS Pipelines Reference*
- *z/VM: CMS Commands and Utilities Reference*
- *z/VM: CMS File Pool Planning, Administration, and Operation*

[V2.3] ISO Date Used in Default CMS IPL Heading

The date field in the default CMS IPL heading has been changed to ISO format (*yyyy-mm-dd*). Each time you IPL CMS, a heading is displayed to identify the VM release level and time stamp of the CMS system being IPLed. A default heading is constructed when the CMS system is generated if the VERSION= parameter in the DEFNUC macro is specified without a value (which is the default).

[V2.3] Logical Device Limit Relief

Support has been added to CP to let you to change the maximum number of logical devices allowed on the system. The old limit of 4096 devices is now the default, but you can use the new CP SET MAXLDEV command to reset the limit up to 32768 devices. The new QUERY MAXLDEV command allows you to determine the current setting.

[V2.3] Allocation of Real Storage for Segment Tables

To represent the virtual storage for each primary address space (virtual machine) larger than 32 MB and each non-primary address space (data space), CP creates segment tables in real storage:

- For each virtual machine larger than 32 MB but less than or equal to 1024 MB, one real storage frame is allocated for the segment table.
- For each data space less than or equal to 1024 MB, one real storage frame is allocated for the segment table.
- For each virtual machine or data space larger than 1024 MB, two consecutive real storage frames are allocated for the segment table.

CP creates the segment table at the start of the storage frame. But in previous releases, the entire storage frame was allocated for the segment table regardless of the actual size of the table. This could have resulted in constraints on real storage availability if large numbers of virtual machines had arbitrarily large storage sizes, or were allowed to create many data spaces and/or very large data spaces, or loaded saved segments or saved systems defined at very high addresses outside the virtual machines.

On the new system, although real storage frames are still nominally allocated for the segment tables as indicated above, only the actual amount of real storage necessary to contain the segment tables is used. Any storage remaining beyond the end of a segment table may be used for CP free storage.

[V2.3] TCP/IP Awareness

Support has been added to the following CP functions to detect TCP/IP addresses and include those addresses in queries and responses:

- QUERY LDEVS command
- QUERY VIRTUAL CONSOLE command

- DIAGNOSE code X'7C' (Logical Device Support Facility)
- CP accounting records
- Access control interface (ACI)

For more information, see:

- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Programming Services*

Changes to Accounting Records for TCP/IP

The formats of the CP accounting records for journaling (Record types 04, 05, 06, and 08) have been changed to supply TCP/IP data, as indicated in Table 4.

Table 4. Changes to Accounting Records for TCP/IP Awareness

Record	Release	Columns	Contents
04	Old	57-70	Reserved
		71-78	LUNAME for SNA terminal
	New	57-62	Reserved
		63-70	Network qualifier for SNA terminal or host virtual machine name for TCP/IP terminal
		71-78	LUNAME for SNA terminal or IP address for TCP/IP terminal
05	Old	58-70	Reserved
		71-78	LUNAME for SNA terminal
	New	58-62	Reserved
		63-70	Network qualifier for SNA terminal or host virtual machine name for TCP/IP terminal
		71-78	LUNAME for SNA terminal or IP address for TCP/IP terminal
06	Old	61-70	Reserved
		71-78	LUNAME for SNA terminal
	New	61-62	Reserved
		63-70	Network qualifier for SNA terminal or host virtual machine name for TCP/IP terminal
		71-78	LUNAME for SNA terminal or IP address for TCP/IP terminal
08	Old	29-64	Reserved
	New	29-48	Reserved
		49-56	LUNAME for SNA terminal or IP address for TCP/IP terminal
		57-64	Network qualifier for SNA terminal or host virtual machine name for TCP/IP terminal

Changes to the ACI for TCP/IP

In the access control interface (ACI) to an external security manager (ESM), the ACIPARMS control block has been changed. In the ACIPARMS parameter list for authorization checking on the LOGON command, a new ACILOGIP option can be set under ACILGOPT (if the ACILOGCL option is also set) to indicate a logical terminal with an IP address. The IP address is specified in ACITRMID.

[V2.3] IPLing with the NODIRECT Option

When you IPL your new system with the NODIRECT option, the OPERATOR user ID is logged on in ESA mode with 2047 MB of virtual storage. In releases prior to VM/ESA V2.3, OPERATOR was logged on in 370 mode with 32 MB of virtual storage.

[V2.3] Product Enablement Support

You can use the CP SET PRODUCT command and the PRODUCT configuration statement to define products and features to your VM system and to determine whether the product or feature can run on that system. You can then use the CP QUERY PRODUCT command to display information about the products that are defined to the system. You can also use DIAGNOSE code X'27C' to request information about the enablement status of a single product or feature.

For more information, see:

- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Programming Services*

[V2.4] Dynamic CP Exits

You can define CP exit points dynamically using the DEFINE EXIT command or system configuration statement. A dynamic CP exit point behaves just like a formally-defined exit point, except that its ability to influence subsequent processing in the module containing the exit point is limited, and it does not carry over an IPL. Dynamic exits provide a convenient way to collect diagnostic or other information or to handle many situations in which the flow of control of a CP module does not need to be changed extensively.

The MODIFY EXIT command or system configuration statement allows you to change the definition of an existing dynamic CP exit point or remove it from the system. The QUERY EXITS command has been enhanced to display additional information about a dynamic exit: exit location, characteristics, and parameter definitions.

For more information, see:

- *z/VM: CP Exit Customization*
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Planning and Administration*

[V2.4] Euro Support

Support has been added to the following:

- Almost 200 new translation tables have been added to VM TCP/IP.
- A new CSL routine, DTCXLATE, provides an application programming interface (API) for programs to use the translation information contained in the TCP/IP translation tables.
- OPENVM GETBFS and PUTBFS commands were updated, increasing the number of code pages that could be specified.
- The FTP server and VM/ESA V2.3 TCP/IP FL310 NFS feature have been updated to allow specification of particular translation tables.
- Language Environment 1.6 and 1.8
- LANRES

- LFS
- The TCP/IP LPR client and LPD server applications have been enhanced to provide users with the ability to specify alternate translation tables.
- RSCS LPR/LPD
- SMTP

For more information, see IBM: z/VM and VM/ESA are EuroReady (www.ibm.com/vm/euro/).

[V3.1] CP Nucleus

An IPLable CP nucleus is no longer supported. Only a nucleus load module may be used to hold CP nucleus code.

z/VM now provides two CP images:

- 32-bit

  You can IPL this CP image on an ESA/390 or z/Architecture server.
- 64-bit

  You can IPL this CP image only on an z/Architecture server.

A dual-image CP is also provided with z/VM. If the dual-image CP is installed, it automatically determines at IPL if the server is z/Architecture-capable. If it is, the 64-bit CP image is loaded; otherwise, the 32-bit CP image is loaded. This can be overridden on the IPL to load the 32-bit CP on a z/Architecture server.

[V3.1] CP Exit Modifications

CP exits may need to be modified for correct operation. For more information, see *z/VM: CP Exit Customization*.

[V3.1] High Level Assembler V1.4

The IBM High Level Assembler V1.4 is required for:

- Installation:
  - Adding devices that cannot be sensed (updating HCPRIO ASSEMBLE)
  - Local modifications
- Servicing the CP Loader (HCPLDR)
- Assembling any CP modules
- Using CP exit routines

[V3.1] RIO370

RIO370 is no longer supported.

[V3.1] 64-Bit Support

z/VM provides support for 64-bit addressing in:

- Storage configuration
- Display/dump/store of storage, registers, and PSWs
- Tracing

Storage Configuration

The CP DEFINE STORAGE command allows larger extents and discontiguous storage definitions. The display of the configurations has been changed.

Two new Directory Control Statements, STORAGE and MAXSTORAGE, allow you to define a default and a maximum virtual storage size for users.

The USER statement now allows a maximum storage of 16 exabytes.

Display/Dump/Store

All commands for displaying, dumping, or storing data or registers accommodate larger registers and storage areas.

Tracing

Trace commands show new instructions, larger registers, PSWs, and storage areas. The new TRACE GG command allows you to trace changes in general-purpose registers for either a 64-bit or a 32-bit image of CP.

[V3.1] Graphical User Interface Facility Changes

The following changes have been made to the Graphical User Interface (GUI) Facility:

- The GUI workstation agents, along with their Help files, are not shipped with z/VM. They are available with limited support from IBM: VM Downloads (www.ibm.com/vm/download/).
- The CMSDESK application modules and message repository are also available as part of the GUI package from the VM Download Library.
- The CMSDESK command and nucleus routines remain in CMS. Documentation is available from the GUI download package. To use these functions, you must download the workstation agents and GUI modules and put them on a CMS minidisk or SFS directory.

[V3.1] VMLINK Improvements

The VMLINK command has been overhauled with both code and documentation changes. Many enhancements and some additions were made to increase the usability of this command:

- Changes were made to better validate some of the values in the VMLINK CONTROL file.
- The PUSH and POP options were changed to ensure the behaviors were consistent and could be documented.
- File mode support was made more complete and the documentation lists all the supported file mode definitions.
- Previously, when multiple INVOKE commands were processed and one or more failed, it was difficult to determine which ones failed because only one message and one return code were issued. VMLINK has been updated to return an error message and return code for each failing INVOKE command.
- Previously, the parameters passed to exits were automatically uppercased before the exit was called. These parameters will now remain in the case entered by the user, allowing mixed-case strings to be passed.
- The .EX and .PX VMLINK variables have been added to complete support for all possible NAMES file tags to be retrieved.
- Messages displayed when disks and directories are detached and released were changed for consistency.
- Message DMS2062I has been improved to include the name of the nickname which is being used when the QUERY option is being processed. This can be especially useful when more than one nickname is specified.

- The documentation has been completely reorganized. Guide information has been removed from the command description in *z/VM: CMS Commands and Utilities Reference* and moved to *z/VM: CMS User's Guide*.

[V4.1] Vector Facility Support

CMS support for the Vector Facility has been removed.

[V4.1] 370-Mode Virtual Machines Not Supported

z/VM no longer supports 370-mode virtual machines. However, the 370 Accommodation Facility allows many CMS applications written for 370 virtual machines to run in XA, ESA, or XC virtual machines. For more information about the 370 Accommodation Facility, see *z/VM: CP Programming Services*.

If a user attempts to define a 370-mode virtual machine, the following will occur:

- If a user issues the SET MACHINE 370 command, the command is rejected and message HCP1016E is issued.
- If a user with a MACHINE 370 statement in their directory logs on, the logon continues, but the virtual machine is defined as XA-mode. In addition, CP's 370 accommodation support (370ACCOM) is set ON and message HCP1512E is issued. Note that doing certain operations, such as resetting the system or defining storage, will set 370ACCOM OFF.

[V4.1] Fast CCW Translation

A new fast channel control word (CCW) translator has been provided that targets Network Adapter channel programs. This enhancement improves performance of I/O to network adapters, allowing any guest, such as Linux for S/390, to take advantage of this new fast-CCW translation processes, which improves I/O performance by reducing CP overhead.

[V4.1] Enhanced Page Fault Handling for Guests

Page fault handling support within CP has been enhanced to allow 31-bit or 64-bit guests to take full advantage of page fault notifications, allowing the guest to continue processing while the page fault is handled by CP.

[V4.1] System Default Language Set Dynamically

All of the IBM-translated languages are now included in CP. You can set the system default language by specifying the CPLANGUAGE operand on the USER_DEFAULTS statement in the system configuration file. You no longer have to rebuild CP to change the system default language. You can dynamically change the default language on the running system by using the SET CPLANGUAGE command.

[V4.2] SET OBSERVER Support

This support allows one virtual machine to observe the console activity of another without affecting the behavior of the observed machine.The following CP functions have been added or updated:

- SET OBSERVER command
- QUERY OBSERVER command
- SET SECUSER command
- CONSOLE directory control statement

[V4.2] DDR Compression

The LZCOMPACT option has been added to the I/O Definition Control Statement for the DDR command to allow users to select an alternate compression algorithm.

[V4.3] Automated Shutdown Support

By virtualizing a hardware interface, automated shutdown enables guests to shut themselves down when signaled to do so by a CP command. This support is exploited by z/VM when running as a guest and can also be exploited by other guest operating systems. The amount of elapsed time during which a guest is allowed to process a termination signal is limited by a timeout interval. automated shutdown can also be used to cause z/VM itself to shut down when a hardware event (for example, deactivation of the logical partition) is initiated while z/VM is active. In this situation, guest virtual machines may, at customer discretion, be given an opportunity to shut themselves down before z/VM terminates.

The following new CP functions have been added for this support:
- SET SIGNAL system configuration statement
- QUERY SHUTDOWNTIME command
- QUERY SIGNAL command
- QUERY SIGNALS command
- SET SHUTDOWNTIME command
- SET SIGNAL command
- SIGNAL command

The following CP commands have been updated:
- FORCE
- SHUTDOWN

For more information, see:
- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*

[V4.3] Enhanced Timer Management

The performance of CP timer management has been improved for environments where a large number of requests are scheduled, particularly for short intervals, and where timer requests are frequently canceled before they become due. A system with large numbers of Linux or z/OS guests with certain workloads would be an example of such an environment. Master-processor overhead has been reduced by allowing timer events to be handled on any processor. Also, clock comparator settings are now tracked and managed across all processors to eliminate duplicate or unnecessary timer interruptions.

[V4.3] Improved Utilization of Large Real Storage

System performance may be improved with better utilization of large real storage when Expanded Storage is unavailable, full, or nearly full while CP is replenishing its available-page list. This is accomplished by moving pages to storage above the 2 GB line when:
- Pages would otherwise go to DASD
- No Expanded Storage has been made available
- Expanded Storage is not available when requested
- Excess storage is available above the 2 GB line

[V4.3] Shared Tape Support

Shared tape support allows multiple guest operating systems to serially share the same tape device (IBM 3424, 3480, 3490, or 3590). It removes the need to manage the movement of a dedicated tape device from one guest to another, because the tape device can now remain attached to several guests concurrently. This support is intended for z/OS, z/OS.e, and OS/390 guests that use ATS Star or IEFAUTOS for automatic tape switching. In principle, however, it will work for any guest that manages its own assignment of tape devices. Traditionally, a tape device would be attached to such a guest with the CP ATTACH command using the NOASSIGN option. This support is not intended for CMS unless some external means of managing assignments or serializing access to the tape device among the sharing users is explicitly implemented.

The MULTIUSER option has been added to the CP ATTACH command and the DEDICATE directory control statement to specify that a tape device is to be shared. The QUERY ALL, QUERY (Real Device), and QUERY TAPES commands have been enhanced to return information about which tape devices are attached MULTIUSER. Assign and Unassign CCWs are simulated for tape devices attached MULTIUSER so that guests running under the same VM image can serialize access to the tape devices they share. The ALL option has been added to the CP DETACH command to allow removal of a shared tape device from all sharing guests with one command invocation. Third party assignment and multiple system assignment (Control Access CCW) are not supported.

For more information, see:

- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*

[V4.3] Virtual Machine Accounting Improvements

CP has been enhanced with additional capability to account for the use of system resources by virtual machines, including those running Linux. A new accounting record is produced that tracks a virtual machine's use of virtual network resources, including virtual channel-to-channel adapters (CTCAs), inter-user communication vehicle (IUCV) or advanced program-to-program (APPC) connections, and virtual (guest LAN) network interface cards (NICs).

The following CP functions have been updated:

- DEFINE LAN system configuration statement
- VMLAN system configuration statement
- OPTION directory control statement
- ACNT command
- DEFINE LAN command
- QUERY LAN command
- QUERY VMLAN command
- SET LAN command
- SET VMLAN command
- DIAGNOSE code X'4C' (Generate Accounting Records)
- Accounting records — Accounting Records for Network Data Transmissions (Record Type 0C) has been added.

For more information, see:

- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Programming Services*

[V4.3] Virtual Machine Resource Manager

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

- Manage guest performance

  Capabilities have been added that allow z/VM to manage guest performance to meet customer-defined goals. A new 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

  I/O management facilities have been added that enable z/VM to exploit the hardware I/O Priority Queueing facility to prioritize guest and host I/O operations. A virtual equivalent of the hardware facility is provided, allowing 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 defined by a new CP command. z/VM will automatically set a priority for I/O operations initiated by virtual machines that do not exploit this function.

The following CP functions have been added for this support:
- IOPRIORITY directory control statement
- QUERY IOPRIORITY command
- SET IOPRIORITY command

For more information, see:
- *z/VM: Performance*
- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*

[V4.4] Automated SFS Shutdown

Utilizing the automated shutdown support added in z/VM V4.3, SFS file pool servers now shut down automatically when CP is shut down. This function helps ease z/VM system administration and helps maintain the integrity of the SFS and its data.

The new SHUTDOWNSIGNAL start-up parameter (the default) enables an SFS file pool server to receive a shutdown signal from a CP SHUTDOWN, FORCE, or SIGNAL command. When an enabled SFS file pool server receives this shutdown signal from CP, the file pool operator STOP command (with no operands) is automatically issued to shut down the server before CP shuts down. To override this automation, the NOSHUTDOWNSIGNAL start-up parameter must be explicitly specified in the SFS file pool server's DMSPARMS file. For more information, see *z/VM: CMS File Pool Planning, Administration, and Operation*.

Message DMS3108I is issued to indicate that the SFS server is stopping as the result of a signal from CP. Message DMS3109I is issued at SFS initialization if the SFS server is not enabled to receive a shutdown signal from CP. For more information, see *z/VM: CMS and REXX/VM Messages and Codes*.

[V4.4] CP Command Response Suppression

The CP SILENTLY command executes another CP command and suppresses its responses, if that command is so enabled. A command is enabled for response suppression by specifying the SILENT option on the DEFINE COMMAND / CMD

or MODIFY COMMAND / CMD command or system configuration statement. Response suppression is supported only for the ATTACH, DETACH, and GIVE commands.

The following CP function has been added for this support:
- SILENTLY command

The following CP functions have been updated:
- DEFINE COMMAND / CMD system configuration statement
- MODIFY COMMAND / CMD system configuration statement
- DEFINE COMMAND / CMD command
- MODIFY COMMAND / CMD command

For more information, see:
- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*

[V4.4] Guest IPL from SCSI Disks

z/VM allows IPLing from FCP-attached SCSI disks for Linux and other guest operating systems that contain support for IPLing from FCP-attached SCSI disks, when z/VM is running on a z990, z900, or z800 server equipped with the function for IPLing from FCP-attached SCSI disks. When this hardware function is available, Linux guests can be started and run completely from FCP-attached SCSI disks in your hardware configuration. z/VM continues to require ESCON-attached or FICON-attached disk or tape for its own IPL and other functions.

The following new CP functions have been added for this support:
- LOADDEV directory statement
- QUERY LOADDEV command
- SET LOADDEV command

The following CP functions have been updated:
- IPL command

The following new Directory Maintenance Facility functions have been added:
- LOADDEV command

For more information, see:
- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: Directory Maintenance Facility Commands Reference*

[V4.4] Hardware Configuration Definition and Hardware Configuration Manager for z/VM

Hardware Configuration Definition and Hardware Configuration Manager for z/VM (HCD and HCM for z/VM) provides a comprehensive I/O configuration management environment similar to that available with the z/OS operating system. The HCD and HCM programs work together to help you create and manage the hardware and software aspects of your z/VM I/O configuration.

HCM runs on a Windows based personal computer connected to the z/VM system through TCP/IP. HCM provides a graphical user interface (GUI) as well as

commands to help you configure your system. You provide the 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. HCM provides the primary user interface to HCD, but HCD provides a backup user interface on your z/VM host for certain I/O configuration tasks and for when HCM is not available.

HCD and HCM for z/VM can be used to create and manage the hardware aspects of the I/O configuration definition for all of the LPARs in your zSeries or S/390 server. For the software aspects, HCD and HCM for z/VM creates and manages the I/O configuration definition for only the z/VM system on which it is running. The software aspects of the I/O configuration definition for operating systems in other LPARs are handled by those operating systems.

HCD and HCM for z/VM provides an easy-to-use alternative to z/VM's existing method of I/O configuration using the I/O Configuration Program (IOCP) and CP's dynamic-I/O configuration commands. The two methods are not intended to be used together, so you should choose one method or the other for your installation.

The following new CP functions have been added for this support:
- QUERY HCD command
- DISABLE HCD command
- DIAGNOSE code X'2AC' – HCD Dynamic I/O

The following CP functions have been updated:
- IODF system configuration statement
- SET IPLPARMS command

For more information, see *z/VM: I/O Configuration*.

[V4.4] HELP Facility Enhancements

The HELP Facility has been enhanced as follows:

- Help files are now provided for all z/VM TCP/IP commands and subcommands. The following new HELP components have been defined for TCP/IP:

  BOOTPD
  BOOTP server administrative (BOOTPD command) subcommands

  DHCPD
  DHCP server administrative (DHCPD command) subcommands

  DNS DNS server administrative commands

  GDDMXD
  GDDMXD/VM interface subcommands

  IMAPADM
  IMAP server administrative commands

  LPD LPD administrative commands

  MPROUTE
  MPROUTE server administrative commands

  NFS NFS server administrative commands

RTE RouteD server administrative commands

SMTP SMTP server administrative commands

SNMP
SNMP administrative commands

SSLADMIN
SSL server administrative commands

TCPIP TCP/IP commands and related functions

TELNET
Telnet protocol client subcommands

TFTD TFTP server administrative (TFTPD command) subcommands

TFTP TFTP client (TFTP command) subcommands

UFTD UFT server administrative (UFTD command) subcommands

X25IPI
X.25 interface administrative commands

You can access the top level of help for TCP/IP by entering one of the following commands:

```
help tcpip menu
help tcpip task
```

Menus are also provided for each of the new HELP components listed above.

Note: For two of the new TCP/IP HELP components, the menu name is different from the component name:

- The menu for the RouteD server administrative commands is ROUTED HELPMENU (not RTE HELPMENU).
- The menu for the TFTP server administrative (TFTPD command) subcommands is TFTPD HELPMENU (not TFTD HELPMENU).

- A new HELP component, LE, has been defined for the Language Environment component. Help files are provided for the following Language Environment commands:

  C370LIB
  CMOD
  CPLINK
  GENXLT
  ICONV
  LINKLOAD

  An LE HELPMENU file is also provided.
- Help for privileged CP commands (formerly HELP component CPOTHER) has been merged with the nonprivileged CP commands (HELP component CP). This can make it easier to display the help file for a CP command, because you do not need to know whether the command is privileged when you enter the HELP command. For example, to display the help file for the ATTACH command (privilege class B), previously you had to enter:

  ```
  help cpother attach
  ```

  Now, to display the help file for the ATTACH command, enter:

  ```
  help cp attach
  ```

To display a menu of help for all CP commands, enter:

```
help cp menu
```

[V4.4] Linux Guest Capacity Improvements

Improvements to the CP scheduler increase the number of Linux and other guest virtual machines that can be managed concurrently. While this increase in capacity can be experienced on all zSeries hardware configurations, it may be more dramatic on larger n-way servers. The scheduler lock is used to serialize scheduler activities, timer request block management, and processor local dispatch vector (PLDV) management. Because contention on the lock is lowered by creating a new lock for timer request block management, CP overhead can be decreased.

[V4.4] Performance Toolkit for VM

The Performance Toolkit for VM is offered as an optional feature of z/VM. Derived from the FCON/ESA program (5788-LGA), the Performance Toolkit is designed to assist system operators, systems programmers, and performance analysts in the following areas:

- Operation of the system operator console in full-screen mode
- Management of multiple z/VM systems (local or remote)
- Post-processing of Performance Toolkit history files and VM monitor data captured by the MONWRITE utility
- Viewing of performance monitor data using either web browsers or PC-based 3270 emulator graphics
- TCP/IP performance reporting

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

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

Although the RealTime Monitor (RTM) and Performance Reporting Facility (PRF) optional features continue to be supported in z/VM V4.4, IBM plans to make future performance management enhancements primarily to the Performance Toolkit. IBM plans that z/VM V4.4 will be the last release in which the RTM and PRF features will be available with z/VM, and plans to withdraw both the RTM and PRF features from marketing in a future z/VM release.

[V4.4] VMRM Enhancements

VMRM has been enhanced to provide the infrastructure necessary to support more extensive workload and systems resource management features that may be delivered in future releases of z/VM. To accomplish this objective, the following have been added:

- Extended syntax for user ID specification in the VMRM configuration file.

  In the WORKLOAD statement, a new wildcard character (*) allows matching up to seven characters of a user ID. For example, `WORKLOAD WORK1 USER` *`user1 vmrm`* `zvm*` would search for any user ID names starting with 'zvm' as the first three characters.

- Continuation of configuration file statements. Configuration file statements can be continued on multiple lines, using a comma as a continuation character at the end of the line being continued.
- Monitor Application Data Sample Records that now contain VMRM workload information.
- The ability to dynamically change the VMRM configuration file, in order to change workloads, workload characteristics, and performance goals, without manually stopping and restarting the VMRM service virtual machine.
- Improvements to debug messages, log file entries, and the format of the log file for better readability and serviceability.
- Performance improvements for searching through lists when analyzing monitor data.
- A SYNCHECK option on the IRMSERV command allows syntax checking of the configuration file.

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

[V5.1] 370 Functions Removed

370 virtual machines are not supported on z/VM Version 4 or later. Obsolete operands, responses, and messages related to 370 virtual machines have been removed. The following CP functions have been modified:

- MACHINE directory control statement
- STORE PSW/CAW/CSW command
- CPFMTXA utility
- HCPLDR utility

[V5.1] Authorization Enhanced for Guest LANs and Virtual Switches

z/VM V5.1 enhances the authorization capabilities for z/VM guest LANs and virtual switches by using the Resource Access Control Facility (RACF) or any equivalent external security manager (ESM) that supports this new authorization function. It is designed to provide ESM-centralized control of authorizations and Virtual LAN (VLAN) assignment.

For more information, see:

- *z/VM: CP Planning and Administration*
- *z/VM: RACF Security Server Security Administrator's Guide*

[V5.1] Contiguous Frame Management

Various algorithms in CP frame management have been adjusted to improve system performance and reduce hang conditions.

[V5.1] CP Functions Removed

z/VM V5.1 operates only on IBM servers that support the IBM 64-bit z/Architecture. Certain CP functions are not provided by z/VM V5.1:

- IPL from a 31-bit image of the CP nucleus
- Preferred (V=R and V=F) virtual machines
- Paging of the CP nucleus

The following CP commands have been removed:

- QUERY IOASSIST

- QUERY VRFREE
- QUERY V=R
- SET CCWTRAN
- SET IOASSIST
- SET NOTRANS

The following CP functions have been modified:
- CPXLOAD configuration statement
- STORAGE configuration statement
- DEDICATE directory control statement
- OPTION directory control statement
- ATTACH command
- AUTOLOG command
- COMMIT command
- CPXLOAD command
- DEDICATE command
- DEFINE command
- DESTAGE command
- DETACH command
- DISCARD PINNED command
- DUPLEX command
- FLASHCOPY command
- INDICATE PAGING command
- INDICATE USER command
- IPL command
- LOCATE (Storage) command
- LOCATE SYMBOL command
- LOCK command
- LOGON command
- MONITOR command
- QUERY CACHE command
- QUERY CACHEFW command
- QUERY CPLEVEL command
- QUERY DASDFW command
- QUERY DUMP command
- QUERY DUPLEX command
- QUERY FENCES command
- QUERY FRAMES command
- QUERY NVS command
- QUERY PINNED command
- QUERY (Real Device) command
- QUERY RSAW command
- QUERY SET command
- QUERY VIRTUAL CTCA command
- QUERY VIRTUAL DASD command
- QUERY (Virtual Device) command
- QUERY VIRTUAL DUPLEX command
- QUERY VIRTUAL GRAF command
- QUERY VIRTUAL LINES command
- QUERY VIRTUAL OSA command
- QUERY VIRTUAL PRINTER command
- QUERY VIRTUAL PUNCH command
- QUERY VIRTUAL READER command
- QUERY VIRTUAL TAPES command
- REDEFINE command
- RESET command

- SAVESEG command
- SET CACHE command
- SET CACHEFW command
- SET DASDFW command
- SET DUMP command
- SET IPLPARMS command
- SET NVS command
- SET SCMEASURE command
- SHUTDOWN command
- SYNCMDRS command
- TRACE command
- UNDEDICATE command
- UNLOCK command
- XAUTOLOG command
- DDR utility
- HCPLDR utility
- SALIPL utility
- DIAGNOSE code X'84'
- DIAGNOSE code X'90'
- DIAGNOSE code X'98'
- MDLATENT macro
- SYSRES macro
- SYSSTORE macro
- *RPI system service
- OPTIONS directive for CPXLOAD

[V5.1] Deploying Linux on zSeries with z/VM

z/VM: Getting Started with Linux on zSeries, SC24-6194, describes z/VM basics and how to configure and use z/VM functions and facilities to create and manage Linux servers running on zSeries processors. This document is designed to help systems personnel (system programmers, administrators, and operators) with limited knowledge of z/VM to deploy Linux servers on z/VM more quickly and more easily. It provides some requirements and guidelines for z/VM installation, but for the most part assumes that you have installed z/VM and are ready to deploy Linux servers in z/VM virtual machines. Specific topics covered include:

- Configuring, administering, and servicing a z/VM system
- Configuring TCP/IP for z/VM
- Creating and cloning Linux virtual machines
- Setting up basic system automation
- Monitoring performance and capacity

[V5.1] Performance Toolkit Enhancements

The Performance Toolkit feature available with z/VM V4.4 included functional equivalence to the RealTime Monitor (RTM) optional feature. The Performance Toolkit has been enhanced in z/VM V5.1 by providing functional equivalence to the Performance Reporting Facility (PRF) optional feature, thereby eliminating the need for separate products (PRF and RTM) to help manage your performance more efficiently. Other new function includes:

- High-level Linux reports based on Application Monitor records from Linux guests
- Report for SCSI disks

The Performance Toolkit also supports:

- Monitor records created by the SUSE LINUX 2.6 kernel
- The collection of data on Queued Direct I/O (QDIO) activity, such as OSA-Express and FCP

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

z/VM HELP files are now provided for all Performance Toolkit commands and messages.

[V5.1] SCSI FCP Disk Support

z/VM V5.1 supports SCSI FCP disk logical units (SCSI disks) for both system and guest use. SCSI disks supported are those within an IBM ESS when it is connected to a fibre-channel (FC) fabric by zSeries FCP channels.

SCSI disks can be used directly by a guest operating system when an FCP subchannel is dedicated to the guest. Such a guest must provide its own SCSI device support. Linux for zSeries is designed to be one such guest.

SCSI disks can also be used as emulated 9336 Model 20 fixed-block-architecture (FBA) disks. z/VM CP and CMS rely almost exclusively on this emulated-FBA support for their SCSI use. Specifically, this use includes system paging, spooling, directory services, minidisks, and all other system functions and programming services that can use FBA disks. Guests that support FBA disks (such as CMS, GCS, RSCS, and VSE) also can use SCSI disks through the emulated-FBA support, without requiring their own SCSI support.

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

z/VM supports an individual emulated FBA disk up to 381 GB in size. However, directory, paging, and spooling allocations must reside within the first 64 GB of a CP-formatted volume. Other kinds of CP allocations (TDSK, PERM, and PARM) may be defined beyond the first 64 GB.

The following new CP functions have been added for this support:
- EDEVICE system configuration statement
- DELETE EDEVICE command
- QUERY EDEVICE command
- SET EDEVICE command

The following CP functions have been updated:
- MDISK directory statement
- QUERY ALLOC command
- CPFMTXA utility
- DIRECTXA utility
- DISKMAP utility

Screen layouts for the following CMS functions have been adjusted to accommodate larger file and disk sizes:
- BROWSE command
- FLIST command

- SADT command
- DIRMAP utility
- QSYSOWN utility

For more information, see:

- *z/VM: CP Planning and Administration*
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CMS Commands and Utilities Reference*

The Directory Maintenance Facility (DirMaint) feature has also been updated to support the larger disk sizes. Larger values may appear in the disk space reports generated by commands such as DIRMAP, FREEXT, and USEDEXT.

[V5.1] SPTAPE Command Removed

The SPTAPE command has been removed. The SPXTAPE command offers superior function and speed for backup and recovery of spool files. Backup tapes made with SPTAPE cannot be restored to a z/VM V5.1 or later system.

In addition, the following CP commands have been modified:

- QUERY ALL
- QUERY READER/PRINTER/PUNCH
- QUERY (Real Device)
- QUERY TAPES

[V5.1] System Administration Facility To Be Withdrawn

IBM intends to withdraw the System Administration Facility from a future release of z/VM. Customers using the System Administration Facility to create and manage Linux server images as guests of z/VM should start using other systems management facilities of z/VM, write a client application using the systems management APIs introduced in z/VM V4.4, or acquire a system management application from an IBM solution provider. (Also see "[V5.1] Deploying Linux on zSeries with z/VM" on page 179.) Documentation is available in *z/VM: System Administration Facility*, SC24-6034, but this publication is not included in the z/VM V5.1 library.

[V5.1] FCP LUN Access Control by z/VM with Linux for zSeries Guests

z/VM V5.1 supports logical unit number (LUN) access control for SCSI FCP disks connected to a FICON Express on the z990 or z890. The LUN access control function of zSeries FCP channels is expected to provide added security for SCSI devices connected to a SAN. For a logical partition running z/VM V5.1, you can specify separate permissions for the z/VM system itself (when z/VM is using SCSI disks for system operations) and for each Linux for zSeries guest. In addition, if you run a second-level z/VM V5.1 system, you can specify separate permissions for that system and for each of its Linux for zSeries guests.

The following CP functions have been updated for this support:

- DEDICATE directory control statement
- ATTACH command
- QUERY FCP command

For more information, see:

- *z/VM: CP Planning and Administration*

- *z/VM: CP Commands and Utilities Reference*

[V5.1] Large Disk Support on DS8000 and DS6000

On the DS8000 and DS6000, z/VM V5.1 supports the definition of emulated FBA disks (SCSI disks) up to 2,147,483,640 512-byte blocks (1 TB minus 1 page) in size. z/VM V5.1 also supports the definition of a 3390 Model 9 disk up to 65520 cylinders in size, with the exception of CMS and GCS, which allow access to a 3390 Model 9 up to 32767 cylinders.

CMS Minidisk Size Limitation

Because the CMS file system requires file status and control information to reside below 16 MB in virtual storage, there is a practical limitation on the size of CMS minidisks. As a minidisk increases in size, or more files reside on the disk, the amount of virtual storage associated with the disk for CMS file system status and control increases in storage below 16 MB. The current ECKD DASD limitation is 32767 cylinders for a 3390 disk on an IBM DASD subsystem, or about 22 GB of data. IBM suggests that customers defining FBA SCSI disks for use by CMS should set a practical limit of about 22 GB. If larger disks are defined, they should be limited to contain very few files, or the CMS file system may not be able to obtain enough virtual storage below 16 MB to format or access those disks. For more information, see the ACCESS command in *z/VM: CMS Commands and Utilities Reference*. The maximum size for FBA SCSI disks supported for use by CMS or GCS is 381 GB.

[V5.2] 64-bit Exploitation

z/VM V5.2 further exploits the 64-bit addressing capability of IBM z/Architecture, providing major scalability improvements when operating on IBM Z servers with large storage configurations. CP now uses storage above 2 GB for a much broader set of operations. Previously, guest pages had to be moved below 2 GB for many reasons. For example, guest I/O buffers for both standard I/O and QDIO, used for networking and FCP devices, were moved below 2 GB when an I/O operation was initiated. For QDIO, CP queue-control structures had to reside below 2 GB, as did guest queue-control structures, which were moved below 2 GB when needed. Now, I/O can be done using buffers anywhere in real storage, and QDIO structures can reside above 2 GB, as can most CP control blocks. These improvements offer constraint relief for large-real-storage virtual-server environments that are storage-intensive.

Prior to z/VM V5.2, storage addresses in the system execution space (the address space in which CP executes) were identity mapped to real storage. That is, a logical storage address was identical to the real storage address. Now, only the CP nucleus and prefix pages are identity mapped. Most of the system execution space (also known as host logical storage) is dynamic. Frames in real storage are used to back host logical storage pages, but backing frames are not necessarily contiguous. Frames below 2 GB are used only when required, freeing CP to exploit backing frames above 2 GB for most operations. For more information about host logical storage, see *z/VM: CP Planning and Administration*.

The following CP functions have been added:
- LOCATE SXSTE command
- QUERY SXSPAGES command
- QUERY SXSSTORAGE command
- DIAGNOSE code X'290'

The following CP functions have been updated:

- DISPLAY (Host Storage) command
- DUMP (Host Storage) command
- INDICATE LOAD command
- INDICATE USER command
- LOCATE CMDBK command
- LOCATE DGNBK command
- LOCATE FILID command
- LOCATE FRAMETBL command
- LOCATE ICLBK command
- LOCATE LDEV command
- LOCATE RDEV command
- LOCATE SHPBK command
- LOCATE SNABK command
- LOCATE SPFBK command
- LOCATE (Storage) command
- LOCATE SYMBOL command
- LOCATE VDEV command
- LOCATE VMDBK command
- LOCATE VSMBK command
- LOCATE XITBK command
- LOCK command
- QUERY CPCMDS command
- QUERY FRAMES command
- SET CPTRACE command
- SET RESERVED command
- STORE (Host Storage) command
- UNLOCK command
- DIAGNOSE code X'04'
- DIAGNOSE code X'90'
- DIAGNOSE code X'98'
- DIAGNOSE code X'250'

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Programming Services*

The CP functions that produce hard and soft abend dumps and SNAPDUMPs have been enhanced to dump storage above 2 GB if such storage exists. Likewise, a standalone dump of a z/VM system or a VMDUMP of a virtual machine in z/Architecture mode can now include storage above 2 GB. See "[V5.2] 64-Bit Dump Support" on page 265.

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

- **TCP/IP for z/VM:** Enhanced QDIO device driver to use I/O buffers above 2 GB when possible. This will help to increase system reliability by reducing contention for storage below 2 GB.
- **Block I/O (Diagnose X'250'):** Extended to allow a virtual machine to specify parameter addresses and I/O buffers with addresses above 2 GB. This will benefit operating systems that use DIAGNOSE code X'250' in a 64-bit environment.
- **VM Dump Tool:** Enhanced to support the new dump formats.

[V5.2] CP Storage Addressing and Module Linkage Changes

Exploitation of 64-bit addressing has resulted in fundamental changes to CP storage addressing and module linkage.

ATTENTION

Existing customer-written CP exits will most likely need to be updated and recompiled.

The following CP changes might affect customer-written CP exits:

Storage Addresses
A central storage address reference by CP is typically termed a host logical address, or HLA. HLAs map the system execution space (SXS), which is the address space where CP runs. CP control blocks, the CP nucleus, and files loaded with CPXLOAD are all mapped in the SXS and are addressed by HLAs.

The SXS is not called CP Virtual Storage, even though the HLA undergoes hardware dynamic address translation (DAT) to generate a host real address (HRA) and address prefixing in order to generate a host absolute address (HAA). You can ignore HRAs and HAAs unless the need for such an address is explicitly stated. For example, the HRA and HAA would be used for hardware addresses like CCW addresses. You can ignore CP Virtual Storage unless the need for such an address is explicitly stated. Generally, write your programs and refer to CP control blocks just like always.

Module and entry point attributes
Runtime attributes are extended to include PSW Translation Mode (PSW bits 16-17) and PSW Addressing Mode (PSW bits 31-32), called Tmode and Amode respectively. Specification of runtime attributes in HCPMDLAT or in the component ID macro xxxMDLAT dictates how the CP linkage macros (HCPCALL, HCPGOTO, HCPENTER, HCPEXIT, HCPLCALL, HCPLENTR, and HCPLEXIT) expand.

Formal entry requirements
Every formal execution entry into a module must be through an HCPENTER or HCPLENTR macro. The term “formal” means using HCPCALL, HCPLCALL, or HCPGOTO. This includes HCPCALL TYPE=INDIRECT, deferred HCPCALLs, and deferred HCPGOTOs. Previously, you could stack (by calling HCPSTK, for example) a deferred transfer of control (that is, a deferred HCPCALL or HCPGOTO) to an arbitrary location in a module. Now, that location must be defined by HCPENTER or HCPLENTR.

4 KB constraint relief
Module CSECTs are arranged (by use of High Level Assembler LOCTR statements) so that literal data (LTORG data) and certain other macro-generated data are placed early in the CSECT. Relative-and-Immediate instructions are used extensively throughout CP, particularly in instructions generated by macros. These together mean that a single module base register frequently can be adequate for module addressability even when the module grows well beyond 4 KB in size. The HCPPROLG macro should specify a single base register, BASE=(R12).

Changes to CP linkage macros
In the HCPPROLG macro, the ATTR parameter no longer allows the

DIRECTCALL and INDIRECTCALL keywords. The new HCPCALL parameter is to be used (HCPCALL=DIRECT or HCPCALL=INDIRECT). Also, the TRSOURCE parameter is no longer allowed. The new RESTRICT parameter is to be used (RESTRICT=TRSOURCE).

In the HCPCALL macro, the ATTR parameter may be used to specify only entry point attributes, not module running attributes.

MDLATENT macro restructure
The syntax of the MDLATENT macro has been extensively changed. Parameters previously retained only for compatibility (FULLREG and NUCLEUS) have been removed. Parameters have been added for specifying the addressing mode (Amode) and translation mode (Tmode). In addition, the relationship between module attributes and entry point attributes and how they are specified in MDLATENT have changed.

For more information about the changes to the CP linkage macros and MDLATENT, see *z/VM: CP Exit Customization*.

[V5.2] DirMaint Enhancements

The DirMaint, function level 510, optional feature has been enhanced as follows:

- A new SUBSCRIBE operand is added to the DIRMAINT command to manage subscriptions to directory updates. Specifically, you can now add new subscriptions, delete old subscriptions, and query existing subscriptions.
- New exits and configuration data files are added to create a more seamless interface between DirMaint and RACF/VM, for such functions as user creation and deletion, password management, POSIX segment management, ACI group management, permission requirement for facilities that require additional coordinated CP and RACF privileges, and discrete resource profile creation and deletion.
- Support for new systems management APIs (see "[V5.2] Enhanced Systems Management APIs" on page 254).
- A new PAVALIAS operand is added to both the DASDOPT and MINIOPT commands, to create one or more parallel access volume devices for the minidisk base device. Also, a new PAVALIAS command is added to request that the PAVALIAS operand of the DASDOPT or MINIOPT directory statement associated with a specific device be added, changed, or deleted.
- A new NOMEMASSIST operand is added to both the OPTION and SETOPTN commands, to disable MEMASSIST support for a user at logon time.

For more information, see *z/VM: Directory Maintenance Facility Commands Reference*.

[V5.2] Large Disk Support on IBM DASD Subsystems

z/VM supports very large emulated 3390 Model 9 DASD (known as Model 27 and Model 54) on IBM DS8000, DS6000, and ESS DASD subsystems and supports disks up to 65520 cylinders on those devices. However, as described in "CMS Minidisk Size Limitation" on page 182, the limitation for CMS and GCS disks is 32767 cylinders. In addition, a minidisk defined with greater than 32767 cylinders, including a temporary disk, is not eligible for minidisk cache.

[V5.2] Performance Toolkit Enhancements

The Performance Toolkit feature is enhanced in z/VM V5.2 with the following new screens (along with their corresponding subcommands):

- Available List Log Screen (FCX254) – AVAILLOG

- Demand Scan Log Screen (FCX259) – DEMNDLOG
- QEBSM Activity Screen (FCX255) – QEBSM
- QEBSM Activity Log (per Device, FCX256) – QEBSMLOG
- Storage Management Logs Menu (FCX260) – STORMENU
- Storage Utilization Log Screen (FCX253) – STORLOG
- SXS Available Page Queues Management Log Screen (FCX261) – SXSAVAIL
- SXS Deferred Tasks Management Log Screen (FCX263) – SXSDEFER
- SXS Page Management Log Screen (FCX262) – SXSPAGE
- SXS Storage Utilization Log Screen (FCX264) – SXSUTIL
- User QDIO Activity Screen (FCX257) – UQDIO
- User QDIO Activity Log (per User), FCX258 – UQDIOLOG

Additional enhancements:

- The FC RELOAD subcommand now has an additional SYSTEMS parameter to allow the FCONRMT SYSTEMS file to be reloaded without interrupting the running server.
- The Performance Toolkit supports the Open Systems Adapter for NCP, CHPID type OSN.

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

[V5.2] System Administration Facility Removed

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

[V5.2] Vector Facility Support Removed

The Vector Facility is an S/390 hardware feature that is not available on z/Architecture servers. CP support for the Vector Facility, including the CP virtual vector facility, has been removed.

The following CP commands have been removed:

- DEFINE VECTOR
- DETACH VECTOR
- DISPLAY VECTOR
- DUMP VECTOR
- INDICATE VECTOR
- QUERY VECTOR
- QUERY VIRTUAL VECTOR
- STORE VECTOR
- VARY VECTOR

The following CP functions have been updated:

- DEDICATE command
- DEFINE command
- DEFINE CPU command
- DETACH command
- DETACH CPU command
- DETACH CRYPTO command
- DISPLAY command
- DUMP command

- DUMP (Guest Storage - ESA/XC) command
- DUMP (Guest Storage - ESA/390) command
- DUMP (Guest Storage - z/Architecture) command
- INDICATE LOAD command
- INDICATE USER command
- LOGON command
- QUERY TIME command
- QUERY VIRTUAL ALL command
- SET CPTRACE command
- STORE command
- STORE (Guest Storage - ESA/390) command
- STORE (Guest Storage - z/Architecture) command
- TRACE: Command Responses
- VARY PROCESSOR command
- VMDUMP command

[V5.3] Host Real Storage Improvements

Changes to page table allocation in z/VM V5.3 allow z/VM images to support significantly more real storage (memory) than the prior limit of 128 GB, as well as more virtual storage for guests. To reduce the constraints on host real storage below 2 GB, the page management blocks (PGMBKs) have been moved above 2 GB. In addition, CP has been enhanced to improve the management of contiguous frames of host real storage.

The following CP external interfaces have been updated:

- The SET CPTRACE command trace codes 4003 and 4004 have been replaced with four new trace codes, 4005 - 4008. For more information about these trace codes, see *z/VM: Diagnosis Guide*.
- The following monitor records have been updated:
  - Domain 0 Record 3 - MRSYTRSG - Real Storage Data (global)
    - The following fields have been added:
      - SYTRSG_RSASNG2G - Count of all single frames that are currently on the >2G available list.
      - SYTRSG_RSAANGAV - Count of all single frames that are currently on the <2G available list.
    - The descriptions of the following fields have changed:
      - SYTRSG_RSAAVAIL - Count of all frames that are currently on the <2G available lists. This is a total value (single frames + contiguous frames) as in previous releases.
      - SYTRSG_RSA2GAVL - Count of all frames that are currently on the <2G available lists. This is a total value (single frames + contiguous frames) as in previous releases.
  - Domain 0 Record 23 - MRSYTLCK - Formal Spin Lock Data (global)

    Note: See "[V5.3] Support for Up to 32 CPUs" on page 189.
    - This support uses the following fields of this new record:
      - SYTLCK_CALRAVLLK - Entry for <2G available list lock (RSAAVLLK).
      - SYTLCK_CALRAVCLK - Entry for <2G available contiguous list lock (RSAAVCLK).
      - SYTLCK_CALR2GLCK - Entry for >2G available list lock (RSA2GLCK).
      - SYTLCK_CALR2GCLK - Entry for >2G available contiguous list lock (RSA2GCLK).

- Domain 3 Record 1 - MRSTORSG - Real Storage Management (global)
  - The following fields have been added:
    - STORSG_RSAAVCLT - Count of the <2G contiguous available list low threshold value. This indicates when the <2G contiguous frame demand scan should be invoked. This provides a buffer of frames for the system to use while demand scan is replenishing the <2G contiguous frame available list.
    - STORSG_RSAAVCHT - Count of the <2G contiguous available list high threshold value. This is the number of frames that should be taken by the <2G contiguous available list replenishment functions to satisfy the current user and system needs.
    - STORSG_RSAAVCLG - Count of the >2G contiguous available list low threshold value. This indicates when the >2G contiguous frame demand scan should be invoked. This provides a buffer of frames for the system to use while demand scan is replenishing the >2G contiguous frame available list.
    - STORSG_RSAAVCHG - Count of the >2G contiguous available list high threshold value. This is the number of frames that should be taken by the >2G contiguous available list replenishment function to satisfy the current user and system needs.
  - The descriptions of the following fields have changed:
    - STORSG_RSAAVLLT - Count of the <2G single frame available list low threshold value. This indicates when the <2G single frame demand scan should be invoked. This provides a buffer of frames for the system to use while the demand scan is replenishing the <2G single frame available list.
    - STORSG_RSAAVLHT - Count of the <2G single frame available list high threshold value. This is the number of frames that should be taken by the <2G single frame available list replenishment function to satisfy the current user and system needs.
    - STORSG_RSAAVLHG - Count of the >2G single frame available list high threshold value. This is the number of frames that should be taken by the >2G single frame available list replenishment function to satisfy the current user and system needs.
    - STORSG_RSAAVLLG - Count of the >2G single frame available list low threshold value. This indicates when the >2G single frame demand scan should be invoked. This provides a buffer of frames for the system to use while the demand scan is replenishing the >2G single frame available list.
- Domain 3 Record 2 - MRSTORSP - Real Storage Activity (per processor)
  - The following fields have been added:
    - STORSP_PLSASFCL - Count of the times a single frame request <2G took a frame from the <2G contiguous frame available list.
    - STORSP_PLSASFCG - Count of the times a single frame request >2G took a frame from the >2G contiguous frame available list.
- Domain 4 Record 3 - MRUSEACT - User Activity Data
  - The following fields have been added:
    - USEACT_VMDTMORD - Last time (sequence) the PGMBK List was reordered.
    - USEACT_VMDPBKCT - Count of the PGMBKs that are currently on the PGMBK List.

- USEACT_VMDCRPGM - Count of the referenced PGMBKs on the PGMBK List at the last reorder.
- USEACT_VMDCUPGM - Count of the unreferenced PGMBKs on the PGMBK List at the last reorder.

For specific information on the effect of these changes and on the new upper bound of real storage supported, see IBM: z/VM Performance Report (www.ibm.com/vm/perf/reports/).

[V5.3] Support for Up to 32 CPUs

z/VM V5.3 can support up to 32 real processors in a single z/VM image on an IBM Z server. To help reduce scheduler lock contention, the scheduler lock has been changed from an exclusive spin lock to a shared/exclusive spin lock.

To allow easier debugging of spin lock contention problems, a new CP monitor record has been created and others have been updated . The new monitor record is MRSYTLCK (Domain 0, Record 23), Formal Spin Lock Data. This record contains an entry for each formal spin lock. The first three fields in the record specify how many spin lock entries are included in the record, how big each entry is, and the displacement from the beginning of the record to the first entry. Each spin lock entry contains the following fields:

- SYTLCK_CALLCKID - Lock identifier

  Note: Once an identifier value is assigned to a particular lock, it will always be used for that lock and will never be reused.
- SYTLCK_CALSSCNT - Total count of spin routine completions for shared requests
- SYTLCK_CALSTIME - Total elapsed spin time for shared requests
- SYTLCK_CALXSCNT - Total count of spin routine completions for exclusive requests
- SYTLCK_CALXTIME - Total elapsed spin time for exclusive requests

Previously, data was collected only for the scheduler and TRQBK locks and was recorded in monitor record MRSYTSCG (Domand 0, Record 10), Scheduler Activity. The existing data for these two locks (total spinning completions and total spin time) will continue to be collected in MRSYTSCG, and the new data for these locks will be included in MRSYTLCK.

Data for the following formal spin locks will be collected in the new MRSYTLCK record:

BUTDLKEY - Spin lock for HCPBUTDL

DCTLLOK - Directory control lock

FSDVMLK - Free storage lock

HCPPGDAL - CPVOL list lock

HCPPGDML - Exposure MLOADs and system total MLOAD lock

HCPPGDPL - Page allocation lock

HCPPGDTL - T-disk allocation lock

HCPPGDSL - SPOOL allocation lock

HCPRCCMA - Switch master lock

HCPRCCSL - Symbolic-lock list lock

HCPTMFLK - HCPTIMER lock

HCPTRQLK - TRQBK queue lock
NSUIMGLK - IMG SNTBK lock
NSUNLSLK - NLS SNTBK lock
RCCSFQL - Save frame queue lock
RSAAVLLK - <2G available list lock
RSAAVCLK - <2G available contiguous list lock
RSA2GLCK - >2G available list lock
RSA2GCLK - >2G available contiguous list lock
RSACALLK - Global cleared available list lock
RSANOQLK - No-owned resource queue lock
RSASXQLK - SXSPM queue lock
SRMALOCK - Adjunct list lock
SRMATDLK - Spin lock for updating SRMATOD and SRMATOD2
SRMSLOCK - Scheduler lock
SYSDATLK - SYSDATE/SYSTODMD lock

The following monitor records have been updated:

- MRSYTSCG (Domain 0, Record 10)
  - SYTSCG_SRMCPUWT - This field has been increased to a doubleword to allow a bit for each processor. This field is the instantaneous value of the bit mask of processors that are currently in wait state.
- MRSYTSCP (Domain 0, Record 13)
  - SYTSCP_PLSDSPCN - This new field holds the count of the number of times the dispatcher has spun on the scheduler lock, not been able to attain it, and then opened an interruption window before trying again. This is an indication of the informal scheduler lock spin time occurring in the dispatcher when it attempts to get the lock without calling the formal spin lock manager.
- MRUSEACT (Domain 4, Record 3)
  - USEACT_VMDUFOLKCT - This new field indicates the total number of times the formal spin lock manager was called to spin on this user's UFO list lock when attempting to get it in exclusive mode.
  - USEACT_VMDUFOLKTM - This new field indicates the total elapsed spin time in the formal spin lock manager for this user's UFO list lock when attempting to get it in exclusive mode, starting at zero and counting up.
- MRPRCPRP (Domain 5, Record 3)
  - PRCPRP_PLSSTLCT - This array (formerly called PRCPRP_PLSSTLNU) has been increased from 31 halfwords to 64 fullwords to allow the number of VMDBK steals from the PLDV of each processor to be counted.

The SET CPTRACE command has been updated to support CP trace codes 3400 - 3407 for tracing spin lock operations. These trace entries are generated only for spin locks whose obtain and release macros indicate they should be traced. Trace codes 3650 and 6000 - 6006 have been updated. For more information about CP trace codes, see *z/VM: Diagnosis Guide*.

For further considerations on performance in a large-scale multiprocessing environment, see IBM: z/VM Performance Report (www.ibm.com/vm/perf/reports/).

[V5.3] Asynchronous CP Command Response Support

The Asynchronous CP Command Response (*ASYNCMD) system service handles collecting the CP command responses from a target virtual machine after issuing the FOR command. The FOR command allows commands to be issued to other virtual machines and receive the command responses and return code either to the terminal or over an IUCV connection to the Asynchronous CP Command Response (*ASYNCMD) system service. The FOR command is an alternative to the SEND command.

Due to the new FOR command, the FORWARD command has been changed to increase the minimum abbreviation to 4 characters (FORW) from 2 characters.

In addition, in order to use this system service, a new operand, *ASYNCMD, has been added to the IUCV directory control statement.

For more information, see *z/VM: CP Programming Services*.

[V5.3] DirMaint Enhancements

The Directory Maintenance Facility, function level 530, includes the following enhancements:

- A new CLEANUP command is added, to unhang a DATAMOVE machine by cleaning up internal files and retrying or canceling all workunits assigned to the machine. Also, you can subsequently use this operand to delete those debug files that DirMaint creates during cleanup processing.
- A new COMMAND command is added, to add, replace, delete, or query the COMMAND statements in a directory entry.
- A new *ASYNCMD operand is added to the IUCV command, to specify the asynchronous CP command response system service.
- The ESM_PASSWORD_AUTHENTICATION_EXIT was changed from DVHDAO MODULE to DVHXPA EXEC to authenticate mixed-case passwords. For more information on mixed-case passwords, see your ESM documentation.

  Note: Passwords can be changed to mixed case only by using the ESM.

For more information, see *z/VM: Directory Maintenance Facility Commands Reference*.

[V5.3] Monitor Enhancements

The MONWRITE command has been updated to include the CLOSE operand. The CLOSE operand enhances automation solutions for gathering monitor records for post processing. Such automation includes but is not limited to enabling user exits and setting monitor record intervals.

For more information, see *z/VM: CP Commands and Utilities Reference*.

[V5.3] Password Phrase Support

Password phrase support has been added to meet demanding password length security requirements. Password phrases, in addition to traditional 8-character passwords, may now be used for system logon and other password related services. A password phrase can be up to 200 characters, including all upper and lower case letters, blanks, quotation marks, and any other special character between X'00' and X'FF'. The use of password phrases and the maximum length

are subject to support by your external security manager (ESM). Currently, RACF Security Server for z/VM supports password phrases up to 100 characters.

If your ESM does not support password phrases you may still use traditional passwords, which are limited to 8 characters and cannot contain blanks or any other special characters, in addition to being case insensitive.

Note: The word password means any password, whether it be a phrase or limited to 8 characters. The use of passwords greater than 8 characters depends on your system configuration.

Password phrase support has resulted in the following updates and additions.

- CMS routines:
  - DMSPASS is a new routine that provides functionality similar to DMSPWCHK, except it utilizes diagnose code X'88' subcode 8, to verify a particular user ID and password are valid. DMSPASS is the preferred routine for validating a password.
  - DMSPWCHK has been updated to support a token value of -1 (X'FF') for calling RACROUTE without an ACEE.
  - For more information, see *z/VM: CMS Callable Services Reference*.
- CP system services:
  - *CCS system service password data coming from VSCS is no longer upper-cased or truncated.
  - *RPI system service has been updated to support password verification using diagnose code X'88' as well as allowing your ESM to exploit the password change indicator.
    - ACIDATA has been updated to include a new code (ACIRSCHK) and a new field (ACIPASSA).
    - Two new flags (ACICHGPW and ACIUSEEP) were added to ACILOPTS.
  - For more information, see *z/VM: CP Programming Services*.
- Diagnose code X'88':
  - Subcode 8 and -1 (X'FF') have been added to support calls for password verification.
  - For more information, see *z/VM: CP Programming Services*.
- LOGON command:
  - The CHANGE operand has been added allowing a user to request a change password dialog directly from the command line.
  - Enhanced support for password use both on the system logon screen and command line.
  - For more information, see *z/VM: CP Commands and Utilities Reference*.
- TCP/IP servers:
  - Passwords may be a maximum length of 200 characters.
  - Password verification is done using DMSPASS instead of DMSPWCHK.
  - The RACF option of the FTP server will affect access control checks and no longer passwords.
  - For more information, see *z/VM: TCP/IP Planning and Customization* and *z/VM: TCP/IP User's Guide*.

[V5.3] Performance Toolkit Enhancements

The Performance Toolkit, function level 530, includes the following enhancements:

- Passphrases are supported when accessing the Performance Toolkit using the Web interface.
- The service process for Performance Toolkit has been changed from a full-part replacement MODULE to service by individual object parts, reducing the size of the service deliverable.
- New or updated displays and reports are provided to support the following new z/VM V5.3 functions:
  - Linux monitor data for virtual CPUs and steal time counters
  - Monitor data for virtual network devices and virtual switches
  - Monitor data for guest simulation of zAAPs, zIIPs, and IFLs
  - Monitor data for up to 32 processors in a z/VM image
- Correctly displays the z10 EC and z9 processor models

[V5.3] RACF Security Server for z/VM

IBM RACF Security Server for z/VM, function level 530, is offered as an optional feature of z/VM V5.3. It is the successor to the IBM Resource Access Control Facility for z/VM feature, V1R10.

The RACF Security Server for z/VM includes support for mixed-case passwords and password phrases. A password phrase is a string of up to 100 characters, including blanks, and can be used in addition to, or in place of, the traditional 8-character password. An installation exit is provided to help enable customers to define rules governing the length and content of password phrases.

Additional password management enhancements have been added, including:

- Validation of a password and password phrase using DIAGNOSE code X'88' and the new DMSPASS CSL routine
- Operation with the new industry-standard LDAP server included in z/VM V5.3 to enable remote management of passwords and selected user attributes, and to enable remote applications to perform authorization and auditing using RACF for z/VM
- Access by password phrase, allowing the removal of the 8-character password from a user ID
- Enhanced security of password reset operations, which now removes the password completely, rather than resetting the user password to the default group name, as in prior releases
- Adding the user's password to the password history list when the password is reset by an administrator
- Providing the capability for passwords to be set by administrators or authorized password management applications without the need for the user to immediately change the password on its first use, improving the auditing of password changes

To simplify analysis of the security audit trail, the RACF SMF Unload utility has been updated to store the unloaded data in industry-standard XML format, making it suitable to be examined by a variety of applications, including XML browsers and spreadsheets.

This new RACF Security Server feature will be the base for all future RACF enhancements on z/VM and works with the existing functions and features of z/VM to provide improved discretionary and mandatory access controls, separation of duties, and auditability capabilities of z/VM.

The RACF documentation has been extensively revised and updated:

- *z/VM: RACF Security Server Auditor's Guide*, SC24-6143
- *z/VM: RACF Security Server Command Language Reference*, SC24-6144
- *z/VM: RACF Security Server Diagnosis Guide*, GC24-6145
- *z/VM: RACF Security Server General User's Guide*, SC24-6146
- *z/VM: RACF Security Server Macros and Interfaces*, SC24-6147
- *z/VM: RACF Security Server Messages and Codes*, GC24-6148
- *z/VM: RACF Security Server Security Administrator's Guide*, SC24-6142
- *z/VM: RACF Security Server System Programmer's Guide*, SC24-6149
- *z/VM: Security Server RACROUTE Macro Reference*, SC24-6150

[V5.3] U.S. Daylight Saving Time Effect on z/VM

A provision of the U.S. Government's Energy Policy Act of 2005 and similar legislation enacted by the governments of Canada and Bermuda extends Daylight Saving Time (DST) by four weeks, beginning in 2007. Starting in March 2007, Daylight Saving Time in the United States, Canada, and Bermuda will begin on the second Sunday in March and end on the first Sunday in November.

New sample system configuration file statements will be shipped with z/VM V5.3. System programmers should change the dates that are specified on TIMEZONE_BOUNDARY statements in the existing system configuration files that their systems use.

For a complete description on the changes to Daylight Saving Time in 2007, see IBM Daylight Savings Time alert (www.ibm.com/support/alerts/ca/en/daylightsavingstimealert.html).

The Language Environment PTF for APAR VM64117 must be applied to z/VM V5.1 and V5.2 and provides required changes for C/C++ library functions. This PTF is integrated into z/VM V5.3.

For more detailed information on how DST changes might affect your z/VM installation, see IBM: VM Service: VM62316(www.ibm.com/vm/service/dst2007.html).

[V5.3] User Directory Commands

The new COMMAND directory control statement specifies a CP command (or multiple CP commands) to be executed after a virtual machine is logged on, and before the virtual machine is IPLed. Note that the commands to be executed may be of any class, and that privileged commands can therefore be executed this way without having to give the virtual machine the associated privilege class. For more information, see the COMMAND directory control statement in *z/VM: CP Planning and Administration*.

This control statement is specifically designed so that directories from V5.3 or later systems can still be used on backlevel systems. In this case, any COMMAND statements will be ignored without error.

[V5.3] VMRM Cooperative Memory Management

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 by issuing a CP DIAGNOSE X'10' instruction to release pages of storage.

In addition to the workload management functions for CPU and DASD I/O provided by VMRM, the following is provided for Linux guests:

- A NOTIFY statement with a MEMORY keyword in the Virtual Machine Resource Manager configuration file. Following the keyword is a user ID or a list of user IDs to be notified when virtual memory becomes constrained.
- System and storage domains are monitored for data to be used for calculating memory constraint, as well as how much memory to request the guest machine to release.
- When memory is constrained, VMRM issues a CP SMSG to notify the specified guests with the amount required to release in order to relieve the constraint.
- A message is logged in the VMRM log file, indicating which users were sent an SMSG, and the text of the SMSG buffer. Also, if MSGUSER is specified on the VMRM ADMIN statement, the same message written to the log is written to the MSGUSER user ID's console as well.

Note: Do not confuse VMRM Cooperative Memory Management with Collaborative Memory Management Assist. The latter is a machine feature introduced on IBM System z9 models that allows z/Architecture guests with the appropriate support to exchange memory usage and status information with z/VM.

For additional information on VMRM memory management, see *z/VM: Performance*.

[V5.3] z/VM Integrated Systems Management

z/VM has added a set of functions that enable you to do system management through the Hardware Management Console (HMC), which does not require you to establish additional network connections or do complex configuration of the system. Details of this local management function and how to use it are in the documentation for your HMC.

z/VM's enablement includes the following:

- A new SCLP system service (*SCLP) that allows the system to receive and transmit HMC events.
- A new proxy server that exploits this service and directs requests to the systems management API server to perform the desired function, and sends the results back to the HMC using the *SCLP interface.
- A VM event system service (*VMEVENT) that sends notifications about events that occur in the z/VM system, such as virtual machine status changes. The proxy server receives a notice when one of these events occurs and reports the event to the HMC through the *SCLP interface.
- A new CP command, QUERY IUCV, that allows you to display information about IUCV connections, which this function uses.
- An update to the IUCV directory statement to enable the *SCLP and *VMEVENT system services.

This enablement function is part of the z/VM system, requires no configuration, and does not require additional products to be installed.

For more information about the SCLP and VM event system services, see *z/VM: CP Programming Services*. For more information about the QUERY IUCV command, see *z/VM: CP Commands and Utilities Reference*. For more information about the IUCV directory statement, see *z/VM: CP Planning and Administration*.

To manage guests with the HMC, the HMC and the Support Element (SE) must be at the following level with the minimum required MCL level:

- On the System z10, the HMC and SE must be at level 2.10.0.
- On the System z9 or zSeries 990 or 890, the HMC and SE must be at level 2.9.2.

For information about the minimum required MCL level, see the z/VM subsets of the Preventive Service Planning (PSP) buckets for your particular server.

[V5.3] Improved Memory Management Algorithms

z/VM V5.3 provides improved memory management algorithms to help benefit paging workloads with large memory environments. This enhancement may be more beneficial with the faster processor speeds of the IBM System z10.

[V5.4] DAT Table Performance Enhancement

This enhancement removes the requirement that segment tables and higher level (region) tables must be allocated below 2 GB in host real storage, which could result in difficulty locating the necessary contiguous frames. Segment tables and region tables can now reside anywhere in host real storage, and CP creates them above 2 GB if possible.

[V5.4] DCSS above 2047 MB

This enhancement removes the constraint that a discontiguous saved segment (DCSS) must be defined below 2047 MB in guest storage. A DCSS (that is, a saved segment which is not a member saved segment or a segment space) can now include pages up to 512 GB. However, the maximum size of an individual DCSS is still 2047 MB. Also, CMS and VMSES/E do not support DCSSs above 2047 MB. A DCSS containing pages above 2047 MB must be loaded with DIAGNOSE X'64', using one of the new 64-bit subcodes.

The following CP functions have been updated for this support:

- DEFSEG command
- INDICATE NSS command
- QUERY NSS command
- DIAGNOSE code X'64'

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Programming Services*

[V5.4] DirMaint Enhancements

The Directory Maintenance Facility, function level 540, includes the following enhancements:

- The PRESET *password* prefix on the DIRMAINT command, as well as the *password* operand on both the PW and SETPW subcommands, are updated to support the use of External Security Manager (ESM) password phrases.

- New CONFIG* DATADVH entries allow you to make existing automatic communications with RACF server configurable, and to identify the exit to be called for password-related command processing.
- New information is added to help you process retry or stalled work units.

For more information, see:

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

[V5.4] Dynamic Storage Reconfiguration Support

z/VM provides the capability to dynamically increase the size of its memory (online real storage) by bringing designated amounts of standby storage online. No system re-IPL is required. Also, virtual memory (virtual machine storage) can be configured to define standby and reserved values. Configuring standby and reserved storage before IPLing a guest operating system in the virtual machine allows that operating system to exploit the dynamic storage reconfiguration (DSR) capabilities of the Service Call (SERVC) instruction.

The following new CP command has been added for this support:
- SET STORAGE

The following CP commands have been updated:
- DEFINE STORAGE
- QUERY STORAGE
- QUERY VIRTUAL STORAGE

For more information, see *z/VM: CP Commands and Utilities Reference*.

The following CP monitor records have been added:
- Domain 1 Record 21 - MRMTRMCC - Memory Configuration Change
- Domain 3 Record 21 - MRSTOADD - Add Central Storage

The following CP monitor records have been updated:

- Domain 1 Record 7 - MRMTRMEM - Memory Configuration Data

  New fields:
  - MTRMEM_SYSGSTBY - The standby real storage size in bytes
  - MTRMEM_SYSGSTRS - The reserved real storage size in bytes
- Domain 4 Record 3 - MRUSEACT - User Activity Data

  The USEACT_ASCDEFSZ field is updated with a new value after a guest uses the DSR functions.

[V5.4] Enhanced Dynamic Storage Reconfiguration Support

Further enhancements to z/VM storage management allow better cooperation with Linux on z Systems guests exploiting DSR. z/VM support:

- Allows operation when running second level on z/VM to be more compatible with operation when running in an LPAR
- Displays configured, standby, and reserved values for each virtual storage element with the QUERY VIRTUAL STORAGE command
- Improves z/VM handling of unexpected DSR conditions that can occur

The following CP functions have been updated:

- QUERY VIRTUAL STORAGE command
- SET STORAGE command
- SYSTEM command
- Monitor record Domain 1 Record 21 - MRMTRMCC - Memory Configuration Change

  Two fields have been added to display the count of offline frames above 2 GB and below 2 GB.

For more information, see *z/VM: CP Commands and Utilities Reference*.

[V5.4] Guest FCP Dump

z/VM V5.4 provides support for guests such as Linux to use third party dump utilities.

The following new CP functions have been added for this support:
- QUERY DUMPDEV command
- SET DUMPDEV command

The following CP function has been updated:
- IPL command

For more information, see *z/VM: CP Commands and Utilities Reference*.

[V5.4] Performance Toolkit Enhancements

The Performance Toolkit, function level 540, includes the following enhancements:

- A new option to create your own initial "Web Banner" page is added.
- Several existing displays are updated to support dynamic storage reconfiguration (as described in "[V5.4] Dynamic Storage Reconfiguration Support" on page 197) and other enhancements in z/VM V5.4.

For more information, see:

- *z/VM: Performance Toolkit Guide*
- *z/VM: Performance Toolkit Reference*.

[V5.4] RACF Security Server Enhancements

Attention

If you are migrating from z/VM V5.3 RACF FL530, or if you plan to share your z/VM V5.4 RACF FL540 database with z/VM V5.3 RACF FL530, you must apply PTF UV61030 for APAR VM64383 to your z/VM V5.3 system (and restart your RACF FL530 server) before attempting any migration or sharing. Failing to apply the PTF before migrating or sharing might result in uncorrectable and undetectable RACF database corruption.

Also, if you were an Early Support Program (ESP) customer for z/VM V5.3 and ran RACFCONV on the ESP code level of RACF FL530, you must run RACFCONV on the GA code level of RACF FL530 before proceeding with RACF FL540 initialization or applying the PTF for APAR VM64383.

The RACF Security Server for z/VM, function level 540, is updated to include support for enveloped passwords and password phrases.

The output of the LISTUSER command has been updated to indicate whether password or password phrase enveloping is active.

For more information, see:

- *z/VM: RACF Security Server Command Language Reference*
- *z/VM: RACF Security Server Security Administrator's Guide*

[V5.4] SHUTDOWN Enhancement

A new SYSTEM *systemid* operand has been added to the SHUTDOWN command to help prevent inadvertently shutting down the wrong system. It is optional at installation by a change to the system configuration file. Changes to the FEATURES statement will allow installations to ENABLE or DISABLE this new SHUTDOWN validation.

[V6.1] VMPRF Mode Not Supported by Performance Toolkit

VMPRF mode, which provided compatibility with the report file specifications of the VM Performance Reporting Facility (5684-073), is not supported by the Performance Toolkit for z/VM V6.1. All reports must now follow the current Performance Toolkit specifications.

[V6.1] HMF/VM Not Supported

Host Management Facilities/VM (HMF/VM), V1.1 (5684-157), is not supported on z/VM V6.1. Use the IBM Operations Manager for z/VM, V1.3.0 (5697-J10).

[V6.1] Alternate Method for Hard Limiting of Scheduled System Resources

z/VM supports the CP scheduler enforcing hard limiting of scheduled system resources using a new consumption method. The LIMITHARD setting only affects users with absolute maximum shares defined with the LIMITHARD option on their SHARE setting. See the SET SRM LIMITHARD command for details on setting the method that the CP scheduler will use to enforce the limit on a guest's CPU usage.

The following CP commands have been updated for this support:

- QUERY SRM
- SET SHARE
- SET SRM

The following user directory statement has been updated for this support:

- SHARE statement

[V6.1] Disabling the Guest Page Reordering Process

The SET REORDER and QUERY REORDER commands are added. SET REORDER allows you to disable the guest page reordering process for the entire z/VM system or for specific virtual machines, helping to improve performance for guests with large memory footprints.

[V6.1] Enhanced Contiguous Frame Coalescing

The contiguous frame coalescing algorithms have been enhanced to increase the probability of z/VM finding multiple adjacent frames of real memory, helping to increase system throughput.

[V6.1] Improved Page Release Serialization

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

[V6.1] Performance Toolkit Enhancements

The Performance Toolkit includes updated reports containing the following new information:

- Capacity-Change Reason and Capacity-Adjustment Indication values
- Ensemble membership
- Server Time Protocol (STP) Facility Configuration values
- Virtual NIC types, including IEDN and INMN
- Reorder settings for virtual machines

For more information, see:

- *z/VM: Performance Toolkit Guide*
- *z/VM: Performance Toolkit Reference*.

[V6.1] Shutdown Signals Enhancement

The SET SHUTSIGNAL command is added, which allows a guest that is enabled for shutdown signals to temporarily disable from receiving shutdown signals and later to re-enable for shutdown signals. The QUERY SIGNALS command can be used to determine whether a user is currently enabled to receive shutdown signals.

[V6.2] SSI Cluster Configuration and Management

Functions have been added to define a z/VM single system image (SSI) cluster and its configuration, and to manage an active z/VM SSI cluster. Managing an active SSI cluster includes coordinating the process of a member joining and leaving the cluster, maintaining each member's view of the operational mode of the cluster (“SSI mode”), and maintaining each member's state with regard to the cluster.

The following functions have been added:

- SSI system configuration statement
- BEGIN and END system configuration statements
- REPAIR IPL parameter
- QUERY SSI command
- SET SSI command
- FORMSSI utility
- DIAGNOSE code X'2CC'

The following functions have been updated:

- SYSTEM_IDENTIFIER system configuration statement
- SYSTEM_IDENTIFIER_DEFAULT system configuration statement
- CPSYNTAX utility
- SET CPTRACE command

- *VMEVENT system service (new Class 2 and Class 3)

A new CP monitor domain, Domain 11 - SSI, and new monitor records have been added:

- Domain 11 Record 1 - MRSSISCS - State Change Synchronization Activity
- Domain 11 Record 2 - MRSSISMI - State/Mode Information
- Domain 11 Record 3 - MRSSISCH - State Change Event
- Domain 11 Record 4 - MRSSISLT - Slot Definition
- Domain 11 Record 6 - MRSSIXLK - XDISK Serialization Activity
- Domain 11 Record 7 - MRSSIXDI - XDISK Activity

Note: To use the functions that define and maintain an SSI cluster, the IBM z/VM Single System Image Feature (VMSSI) must be licensed and enabled. See "[V6.2] IBM z/VM Single System Image Feature" on page 10.

[V6.2] SSI Cluster User Identity and Configuration

The z/VM V6.2 user directory is arranged to facilitate the deployment of SSI clusters and differs significantly from the directory for a pre-V6.2 release. The changes include a new type of virtual machine definition (which includes new directory statements), changes to the layout of the system minidisks, and changes to the MAINT user ID.

New Type of Virtual Machine Definition

Two types of virtual machine definitions are now supported in the z/VM directory:

Single-configuration virtual machine definition

This is the traditional type of definition, which begins with a USER statement and consists of the user entry and any included profile entry. In a z/VM SSI cluster, a user ID defined this way can log on to any member of the cluster, but only one member at a time.

As before, this type of definition is intended for general users and virtual machines that support guests.

Multiconfiguration virtual machine definition

This type of definition begins with an IDENTITY statement and consists of the identity entry, the associated subconfiguration entries, and any included profile entry. When included in the common source directory for a z/VM SSI cluster, a multiconfiguration virtual machine definition can be configured to define multiple virtual machine instances, which enables the user ID to be logged on concurrently and independently to multiple members of the cluster. Each of these virtual machine instances can have a different configuration from the others.

A subconfiguration entry begins with a SUBCONFIG statement that corresponds to a BUILD statement in the identity entry that specifies a member of the SSI cluster. Directory statements that apply to the virtual machine instance on any member of the cluster are included in the identity entry. Directory statements that apply only to the virtual machine instance on a specific member of the cluster are included in the subconfiguration entry for that member.

This type of virtual machine definition is intended for system support user IDs, service virtual machines, and servers. The definitions for many IBM-supplied user IDs (such as MAINT, OPERATOR, and TCPIP) have

been changed to multiconfiguration virtual machine definitions. For a list of these user IDs, see Changed virtual machine definitions.

In a non-SSI environment, a multiconfiguration virtual machine definition is configured for a single system (that is, the definition includes at most one BUILD statement and the corresponding subconfiguration entry) and is equivalent to a single-configuration virtual machine definition.

The z/VM source directory can be one of the following types:

SSI-enabled
A source directory that has the SSI option specified on the DIRECTORY statement and includes IDENTITY, BUILD, and SUBCONFIG statements. This type of source directory is created when a z/VM SSI cluster is installed using one of the SSI installation procedures documented in *z/VM: Installation Guide* and is shared by the members of the SSI cluster. (Each cluster member creates its own object directory.) An SSI-enabled source directory can contain single-configuration virtual machine definitions and multiconfiguration virtual machine definitions.

SSI-ready
A source directory that does not have the SSI option specified on the DIRECTORY statement but includes IDENTITY, BUILD, and SUBCONFIG statements. This type of source directory is created when a single z/VM system is installed using one of the non-SSI installation procedures documented in *z/VM: Installation Guide*. An SSI-ready source directory can contain single-configuration virtual machine definitions and multiconfiguration virtual machine definitions, but each multiconfiguration virtual machine definition can define only one virtual machine instance.

Non-SSI
A source directory that does not have the SSI option specified on the DIRECTORY statement and does not include any IDENTITY, BUILD, or SUBCONFIG statements. Although this type of source directory is supported, it is not supplied with z/VM or created by any of the z/VM installation procedures.

SYSAFFIN statements are not supported in an SSI-enabled or SSI-ready directory. If you have a CSE complex, you should convert it to an SSI cluster. For information about a possible conversion approach, see *z/VM: CP Planning and Administration*.

Table 5 shows which directory types are compatible with specific system environments.

Table 5. Directory types and system environments

	Object directory derived from this source directory type		
z/VM release and environment	**Non-SSI**	**SSI-ready**	**SSI-enabled**
Prior to V6.2	Y	Y	Not supported
V6.2 (or later), no SSI cluster defined	Y	Y	Error
V6.2 (or later) in a single-member SSI cluster	Y	Y	Y
V6.2 (or later) in a multimember SSI cluster	Error	Error	Y

Table 5. Directory types and system environments (continued)

	Object directory derived from this source directory type
Notes: 1. "Error" means CP will not bring object directories derived from this type of source directory online. 2. "Not supported" means this environment is not supported because in older z/VM releases CP does not know how to use this directory format. In older z/VM releases, CP can bring the directory online, but cannot support multiconfiguration virtual machine definitions.	

New System Minidisk Layout

The z/VM V6.2 installation DASD volumes (for an SSI installation or a non-SSI installation) are divided into three groups: common, release-specific, and system-specific. The system configuration file and source directory are located on the common volume. In an SSI cluster, the common volume is truly common, and is shared by all the members of the cluster. Likewise, in an SSI cluster there is one set of release-specific volumes for each release level in the cluster.

Corresponding to the new installation DASD structure, the layout of the system minidisks has been changed. Some new minidisks have been defined, some minidisks have been dropped, some minidisks are now owned by different user IDs, and some minidisks have been relocated to different volumes. See "[V6.2] Changes to Predefined User Directory Entries" on page 25.

Attention: It is important for future migrations and for the service tools that you preserve the volume layouts as defined by the z/VM installation process.

Changes to the MAINT User ID

System resources owned by the MAINT user ID in pre-V6.2 releases are now divided between MAINT and two new user IDs, PMAINT and MAINT*vrm*. These user IDs have the following roles:

PMAINT
: Owns the resources that will be common in an SSI cluster, such as the parm disk and the minidisks that contain the source directory files, VMSES/E system-level files, and the new VMPSFS file pool. The system configuration file has been moved to the new PMAINT CF0 minidisk. The CP utilities CPFMTXA, DIRECTXA, and DISKMAP, and the CMS utility DIRMAP have been moved from the MAINT 190 minidisk to the new cross release utilities minidisk (PMAINT 551). PMAINT is defined by a single-configuration virtual machine definition.

MAINT*vrm*
: Owns the resources that are specific to a z/VM release. For example, MAINT620 own the resources specific to z/VM V6.2. MAINT*vrm*, rather than MAINT, must now be used to service all z/VM and preinstalled products, features, and components for the *vrm* release. MAINT*vrm* is defined by a single-configuration virtual machine definition. In an SSI cluster that includes z/VM systems at multiple release levels, the directory will include a MAINT*vrm* user ID for each release.

MAINT
: Owns the resources that are specific to each system. MAINT is defined by a multiconfiguration virtual machine definition and therefore can have concurrent logon instances on the members of the SSI cluster. Most

minidisks formerly owned by MAINT have been moved to either MAINT*vrm* or PMAINT, and MAINT now has links to those minidisks. See "[V6.2] Changes to Predefined User Directory Entries" on page 25.

New and Changed Externals

The following directory statements have been added:

- BUILD
- IDENTITY
- SUBCONFIG
- VMRELOCATE

The following directory statements have been changed:

- DIRECTORY
- LOAD

The following functions have been changed:

- DISKMAP command
- DIRECTXA utility
- DIRMAP utility (CMS)
- VMUDQ macro
- DIAGNOSE code X'3C'
- DIAGNOSE code X'84'
- DIAGNOSE code X'25C'

The following monitor records have changed:

- Monitor record Domain 1 Record 15 - MRMTRUSR - Logged on User

  A new byte (MTRUSR_VMDLOGFG) is added to the end of the record that includes MTRUSR_VMDIDENT, where VMDIDENT is the name of flag in the VMDBK field VMDLOGFG that indicates the user is a multiconfiguration virtual machine.
- Monitor record Domain 4 Record 1 - MRUSELON - User Logon Data

  A new byte as defined for MRMTRUSR is added.

Note: To use the functions that define and maintain an SSI cluster, the IBM z/VM Single System Image Feature (VMSSI) must be licensed and enabled. See "[V6.2] IBM z/VM Single System Image Feature" on page 10.

[V6.2] Cross-System Spool and CP Commands in an SSI Cluster

Cross-system spool enables the CP spooling functions on each member of an SSI cluster to create, manage, and share spool files cooperatively with the other members and enables spool files to be shared among the members. Each member creates spool files only on its own spool volumes, but has access to the spool volumes owned by the other members.

Users (defined by USER definitions) have a single logical view of all their spool files. Users can log on to any member of the cluster and access all of their spool files, regardless of the member on which the files were created. (For a spool file to be accessible, the member on which it was created must be joined to the cluster.) Users can manipulate their spool files in the same ways as on a non-SSI system.

Cross-system support for the single console image facility (SCIF) allows the secondary user and primary user to be logged on to different members of the SSI cluster. Likewise, cross-system support for the observer function allows the observer and the observed user to be logged on to different members of the SSI cluster. The following commands have been updated for this support:

- QUERY OBSERVER
- QUERY SECUSER
- SEND
- SET OBSERVER
- SET SECUSER

The following CP commands have also been updated for SSI cluster support:

- FREE LOGON
- HOLD LOGON

Some CP commands are enabled for cluster use through existing AT *sysname* operands, which allow you to specify the target member. Examples are MESSAGE (MSG), MSGNOH, and SMSG.

The new AT command can be used to remotely issue most privileged CP commands (that is, not class G or class ANY) on other active members of the SSI cluster. Some CP commands have operational requirements when issued via the AT command:

- The following commands have required operands when issued via the AT command. The target user ID (or some other value) must be explicitly specified. The target cannot default to the command issuer or *.
  - CPXH
  - DEFINE LAN
  - DETACH LAN
  - DETACH SYSASCII
  - DETACH XSTORE
  - INDICATE SPACES
  - INDICATE USER
  - QUERY BYUSER
  - QUERY CFLINKS
  - QUERY CPUAFFINITY
  - QUERY IUCV
  - QUERY MAXSPOOL
  - QUERY MEMASSIST
  - QUERY OBSERVER
  - QUERY PRIVCLASS
  - QUERY QIOASSIST
  - QUERY SECUSER
  - QUERY STGEXEMPT
  - RESET RESERVE MDISK
  - SET CPUAFFINITY
  - SET LAN
  - SET MEMASSIST
  - SET QIOASSIST
  - TRSAVE
- The following commands have options that cannot be used when issued via the AT command:
  - RESET (*vdev* option)
  - INDICATE SPACES (*spaceid* option if it does not include an explicit user ID)

- The following commands have an existing AT option that is not allowed when issued via the AT command:
  - INDICATE LOAD
  - MESSAGE
  - MSGNOH
  - QUERY NAMES
  - QUERY USERID
  - SMSG
  - WARNING
- The CP command when preceded by the logical-line-end symbol (for example, #CP) cannot be issued via the AT command. If the CP command without the logical-line-end symbol is issued via the AT command, it will not cause any virtual machine on the target system to enter the CP environment. However, the CP command can still be used to preface other CP commands when issued via the AT command.

[V6.2] Live Guest Relocation in an SSI Cluster

Changes have been implemented to provide the ability for a running Linux guest to be moved transparently from one z/VM system to another within an SSI cluster, a process known as live guest relocation.

The following functions have been added:

- RELOCATION_DOMAIN system configuration statement
- VMRELOCATE directory statement
- DEFINE RELODOMAIN command
- QUERY CHPIDV command
- QUERY RELODOMAIN command
- QUERY VIRTUAL CHPID command
- QUERY VMRELOCATE command
- SET VMRELOCATE command
- VMRELOCATE command

The following functions have been changed:

- GLOBALOPTS system configuration statement
- OPTION directory statement
- ATTACH command
- CPHX command
- DEFINE STORAGE command
- DEFINE (Virtual Disk in Storage) command
- DETACH (Virtual Device) command
- DISPLAY command
- DUMP command
- FORCE command
- HOLD LOGON command
- IPL command
- LOCK command
- LOGOFF command
- QUERY SPACES command
- SET CPTRACE command

- SET LKFACR command
- SET RESERVED command
- SET VDISK command
- STORE command
- SYSTEM command
- UNLOCK command
- IMAGELIB utility

The following CP monitor records are added or changed:

Domain 0 Record 3: Replace field SYTRSG_SYSRSVPG (cardinal count of reserved pages) with SYTRSG_RSARSVPG (the reserved page counter is being moved from the SYSCM to the RSMBK).

Domain 0 Record 7: Add field SYTSHS_QDGRELOC to contain the sum of VDISK blocks reserved for relocations in progress.

Domain 0 Record 8: Add two fullword counters to this record. These are counts of the number of relocations in-flight: one count of incoming and another of outgoing relocations.

Domain 1 Record 15: Add a field that contains the name of the last system from which this virtual machine was relocated and also a field with the name of the system on which this virtual machine originally logged on.

Domain 2 Record 11 I/O priority change event record: Add a flag indicating that this event occurred because of a K relocation and cut one of these records on the destination system during the relocation if the effective I/O Priority values are different on the destination than the source.

Domain 4 Record 1: Add a flag indicating that this event occurred because of a relocation.

Domain 4 Record 2: Add a flag indicating that this event occurred because of a relocation.

Domain 4 Record 3: The following changes are made to this record

- Add an indication that this is the first record since a relocation occurred (both a flag and a timestamp of the end of the relocation).
- Add a field that contains the name of the user's relocation domain.
- Add a field that contains the domain member bit mask.
- Add a flag that indicates if CP is currently showing to the guest a virtual architecture that is not identical to the architecture of the member on which it is running.

Domain 4 Record 4: Add an indication that this is the first record since a relocation occurred (both a flag and a timestamp of the end of the relocation).

Domain 4 Record 11 New Relocation Start event record: Contains the following fields in the source system record:

- Userid of VMRELOCATE command issuer
- Userid of relocation target
- Source system name

- Destination system name
- Flag with relocation information (source or destination system indication)
- Flag with command options
  - FORCE ARCHITECTURE specified (yes or no)
  - FORCE DOMAIN specified (yes or no)
  - FORCE STORAGE specified (yes or no)
  - SYNCH or ASYNCH specified
  - MAXTOTAL time specified
  - MAXQUIESCE time specified
  - IMMEDIATE specified
- Flag with SET VMRELOCATE settings
  - VMRELOCATE set on or off -
- MAXTOTAL time value used
- MAXQUIESCE time value used
- Timestamp of start of relocation

Domain 4 Record 12 New Relocation End event record. Contains the following fields in both the source system and destination system records except as otherwise indicated:

- Userid of VMRELOCATE command issuer
- Userid of relocation target
- Source system name
- Destination system name
- Flag with relocation information
  - Source or destination system indication
  - CPU affinity suppression setting after the relocation – Indication of whether the CPU affinity suppression was changed as a result of the relocation
- Flag with command options
  - FORCE ARCHITECTURE specified (yes or no)
  - FORCE DOMAIN specified (yes or no)
  - FORCE STORAGE specified (yes or no)
  - SYNCH or ASYNCH specified
  - MAXTOTAL time specified
  - MAXQUIESCE time specified
  - IMMEDIATE specified
- Flag with SET VMRELOCATE settings
  - Does guest have VMRELOCATE set on or off
- MAXTOTAL time value used
- MAXQUIESCE time value used
- Reason code (decimal) for end of relocation
  - 0: completed successfully
  - 1: canceled by VMRELOCATE CANCEL command
  - 2: canceled by CPHX command
  - 3: canceled due to lost ISFC connection
  - 4: canceled due to MAXTOTAL time limit exceeded
  - 5: canceled due to MAXQUIESCE time limit exceeded

  - 6: canceled due to eligibility violation
  - 7: canceled due to virtual machine action
  - 8: canceled due to an internal processing error
  - 9: canceled because the CP exit rejected the command
  - 11: canceled because the CP exit gave a return code that is not valid
  - 12: canceled because the destination system detected an error
- Number of storage passes made
- Number of pages transferred on pass 1
- Average number of pages transferred during passes 2 through n-2; this field is zero when the relocation is IMMEDIATE.
- Number of pages transferred on pass n-1 (next to last pass)
- Number of pages transferred on pass n (the last pass)
- Number of virtual devices transferred
- Number of active I/Os encountered
- Number of non-QDIO type I/Os cleared
- Number of QDIO type I/Os cleared
- Timestamp of start of the relocation
- Timestamp of completion of the initial connection to the destination system
- Timestamp of completion of eligibility checks
- Timestamp of completion of initial virtual machine creation on the destination system
- Timestamp of completion of initial address space creation
- Timestamp of completion of virtual machine memory transfer prior to virtual machine quiesce (source only)
- Timestamp of completion of FCP I/O delay (source only)
- Timestamp of completion of stopping the virtual machine
- Timestamp of completion of I/O relocation
- Timestamp of completion of virtual machine state relocation
- Timestamp of completion of final VSIM eligibility checks
- Timestamp of completion of final STMGT eligibility checks
- Timestamp of completion of next-to-last memory pass (source only)
- Timestamp of completion of last memory pass
- Timestamp of completion of final I/O eligibility checks
- Timestamp of completion of virtual machine restart on destination
- Timestamp of completion of virtual machine 'logoff' (source only)
- The number of frames reserved by the virtual machine on the source system
- The number of frames reserved by the virtual machine on the destination system

Domain 8 Record 1: Add an indication that this is the first record since a relocation occurred (both a flag and a timestamp of the end of the relocation).

Domain 10 Record 1 (APPLMON event record): When a virtual machine that used diagnose x'DC' to start up APPLMON data collection is relocated, it may be necessary to give a configuration event record. This will be done if there was an outstanding configuration record on the source system and if the APPLMON monitoring has been started for this virtual machine and a user is connected to *MONITOR.

Domain 10 Record 2 (APPLMON sample record) A new flag called APLSDT_RELO1 will be set the first time this record is returned after a relocation.

Note: To use the functions that define and maintain an SSI cluster, the IBM z/VM Single System Image Feature (VMSSI) must be licensed and enabled. See "[V6.2] IBM z/VM Single System Image Feature" on page 10.

[V6.2] ACIPARMS Enhancements

In the access control interface (ACI) to an external security manager (ESM), the ACIPARMS control block has been changed. The following ACIPARMS parameter lists have been updated:

- LOGOFF
- LOGON

The following ACIPARMS parameter list has been added:

- VMRELOCATE

[V6.2] DirMaint Enhancements

The Directory Maintenance Facility (DirMaint), function level 620, includes the following enhancements to support SSI clusters:

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

Note: To use the functions that define and maintain an SSI cluster, the IBM z/VM Single System Image Feature (VMSSI) must be licensed and enabled. See "[V6.2] IBM z/VM Single System Image Feature" on page 10.

For more information, see:

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

[V6.2] ESM Access Control for Real Devices

This provides the changes to CP and RACF to enable discretionary and mandatory access controls of real devices.

The following commands have been updated for this support:

- ATTACH
- GIVE

[V6.2] LIMITHARD Default Change

The default for the SET SRM LIMITHARD option has been changed to CONSUMPTION. In most workloads, the CONSUMPTION method will give more accurate results than DEADLINE.

[V6.2] Memory Constraint Relief

Frame allocation management has improved system memory scalability and decreased memory and processor constraints for larger memory sizes. Monitor record domain 3 record 1 is changed in the following ways:

- Flags are added to indicate:
  - Pageable pages may be allocated above 2 GB.
  - Pageable pages may be allocated below 2 GB.
  - Demand Scans may be done for pageable pages above 2 GB.
  - Demand Scans may be done for pageable pages below 2 GB.
  - The decision on how to allocate pageable pages will no longer change.
- A fullword that equals the turnover rate multiplier before any minimum/maximum is applied for memory less than 2 GB (RSATURNRB).
- A fullword that equals the current turnover rate multiplier for memory less than 2 GB (RSATURN0).
- A fullword that equals the current turnover rate multiplier for memory greater than 2 GB (RSATURNG).
- A fullword that equals the current count of deferred frame request that are waiting for an available frame that can be located only above 2 GB (RSADefCTA2G).
- A fullword that equals the cumulative count of times HCPPTGRD attempted to redrive tasks waiting for an available frame that can be located anywhere (RSARDaANY).
- A fullword that equals the cumulative count of individual task redrives performed by HCPPTGRD for an available frame that can be located anywhere (RSARDtANY).
- A fullword that equals the cumulative count of times HCPPTGRD attempted to redrive tasks waiting for an available frame that can be located only above 2 GB (RSARDaA2G).
- A fullword that equals the cumulative count of individual task redrives performed by HCPPTGRD for an available frame that can be located only above 2 GB (RSARDtA2G).

[V6.2] Contiguous Page Replenishment Counters in Monitor

The results of an unsuccessful search for contiguous pages to steal are now recorded in monitor record domain 3 record 1:

- A fullword that equals the cumulative count of times demand scan could not satisfy the need for contiguous frames below 2 GB (RSADSCC0).
- A fullword that equals the cumulative count of times demand scan could not satisfy the need for contiguous frames above 2 GB (RSADSCCG).
- A fullword that equals the cumulative count of times demand scan quit pass 1 of PGMBK steal early due to a paging overrun (RSAPGMQ1).

- A fullword that equals the cumulative count of times demand scan quit pass 2 of PGMBK steal early due to a paging overrun (RSAPGMQ2).
- A fullword that equals the cumulative count of times demand scan quit emergency pass of PGMBK steal early due to a paging overrun (RSAPGMQE).

In addition, the following fields are now contained in monitor record domain 3 record 2:

- A fullword that equals the cumulative count of times demand scan completed after pass 1 of PGMBK steal (PLSPPGM1).
- A fullword that equals the cumulative count of times demand scan completed after pass 2 of PGMBK steal (PLSPPGM2).
- A fullword that equals the cumulative count of times demand scan completed after emergency pass of PGMBK steal (PLSPPGME).

[V6.2] MONITOR Changes

The size of the default MONITOR MONDCSS segment shipped with z/VM has been increased from 16 MB to 64 MB. In addition, the default size of the MONITOR SAMPLE CONFIG area has been increased from 241 pages to 4096 pages.

If your MONITOR MONDCSS segment is too small, you will receive the following error message when you try to connect using the *MONITOR system service and the MONWRITE utility:

```
HCPMOW6270E MONWRITE severed the IUCV connection, reason code 2C
HCPMOW6267I MONITOR writer connection to *MONITOR ended
```

If you receive this message, you will need to increase the size of your MONDCSS segment. For more information, see *z/VM: Performance* and the MONITOR command in *z/VM: CP Commands and Utilities Reference*.

[V6.2] Performance Toolkit Enhancements

The Performance Toolkit includes new and updated reports containing the following information:

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

The source COPY files defining extended trend record layouts are now shipped with the product in FCXGPI MACLIB. The details of these records and are no longer documented in *z/VM: Performance Toolkit Reference*, although the records are still listed there.

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

For more information, see:

- *z/VM: Performance Toolkit Guide*
- *z/VM: Performance Toolkit Reference*
- *z/VM: Performance*.

[V6.2] RACF Security Server Enhancements

RACF Security Server enhancements include support for z/VM SSI clusters and support for protected user IDs.

RACF Support for z/VM SSI Clusters

RACF is enhanced to support z/VM SSI clusters by presenting z/VM guests with a common and consistent security image across up to four nodes in an SSI cluster. All participating systems will share the RACF database.

The following commands are updated to automatically propagate to all RACF servers running in the same SSI cluster as the issuing system:

- RVARY
- SETEVENT REFRESH
- SETROPTS (all except SETROPTS LIST)

Note: To use the functions that define and maintain an SSI cluster, the IBM z/VM Single System Image Feature (VMSSI) must be licensed and enabled. See "[V6.2] IBM z/VM Single System Image Feature" on page 10.

For more information, see:

- *z/VM: RACF Security Server Command Language Reference*
- *z/VM: RACF Security Server Messages and Codes*
- *z/VM: RACF Security Server System Programmer's Guide*

RACF Support for Protected User IDs

Protected user IDs are protected from being used to log on to the system and from being revoked through inactivity or unsuccessful attempts to access the system using incorrect passwords and password phrases. The LISTUSER RACF command output is modified to include a new PROTECTED attribute.

The following RACF commands are modified to include new NOPASSWORD and NOPHRASE attributes:

- ADDUSER
- ALTUSER

See *z/VM: RACF Security Server Security Administrator's Guide* for more information.

[V6.2] SECLABEL Support for SECUSER/OBSERVER Facilities

This provides the changes to CP and RACF to enable use of SECUSER/OBSERVER facilities when mandatory access controls (security labels) are active.

The following commands have been updated for this support:

- QUERY OBSERVER
- QUERY SECUSER
- SEND
- SET OBSERVER
- SET SECUSER

[V6.2] Default Virtual Machine Type is now ESA

If you omit the MACHINE statement when you code a virtual machine definition, the user's virtual machine mode is defined by the mode specified by the GLOBALOPTS directory control statement. If no mode is specified in the GLOBALOPTS directory statement, the default mode will be ESA.

[V6.2] STORBUF Default Change

The initial (default) percentage values for SET SRM STORBUF, which partition pageable storage for users with different transaction classes, have been changed to 300%, 250%, and 200%. (The previous defaults were 125%, 105%, and 95%.)

[V6.2] Changed XEDIT Default Case for Other File Types

XEDIT supplies default settings, such as case, for files with special file types. See the list of file type defaults in *z/VM: XEDIT Commands and Macros Reference*. For files with other (not listed) file types, the XEDIT default case has been changed from uppercase to mixed case.

[V6.2] CMS File Pool Server Enhancements

A new file pool server startup parameter has been added that allows the scope of a file pool to be limited to the members of an SSI cluster.

[V6.2] Utilities Moved from the MAINT 190 Minidisk

Files CPFMTXA EXEC, DIRECTXA MODULE, DISKMAP EXEC, and DIRMAP MODULE have been moved from the CMS system disk (MAINT 190) to the cross release utilities disk (PMAINT 551). The new FORMSSI MODULE is also located on this disk. The PMAINT 551 disk must be accessed to run the CP utilities CPFMTXA, DIRECTXA, DISKMAP, and FORMSSI, and the CMS utility DIRMAP.

[V6.2] Additional Performance Toolkit Enhancements

With the PTF for APAR VM65044, the Performance Toolkit recognizes the following new monitor records:

- Domain 6 Record 32 - MRIODHPF - zHPF Feature Change
- Domain 6 Record 33 - MRIODBPA - Virtual Switch Bridge Port Activation
- Domain 6 Record 34 - MRIODBPD - Virtual Switch Bridge Port Deactivation
- Domain 6 Record 35 - MRIODBPS - Virtual Switch Bridge Port Activity.

The performance data screen FCX155 (MONDATA) has been updated to include the new monitor records.

[V6.3] Cross System Extensions (CSE) Support Removed

Support for the cross system extensions (CSE) environment has been removed. z/VM single system image (SSI) clusters provide the new technology for clustering z/VM systems.

The cross-system link (XLINK) function that was included in CSE is still supported for non-SSI systems, and "CSE" is still used in some function and object names, command responses, and messages related to XLINK. For information about cross-system link, see *z/VM: CP Planning and Administration*.

The following CP functions have been removed:
- XSPOOL_SYSTEM configuration statement
- XSPOOL_TRACE configuration statement
- XSPOOL_XLIST_INPUT configuration statement
- XSPOOL_XLIST_OUTPUT configuration statement
- XSPOOL QUERY command
- XSPOOL QUERY SHARE command
- XSPOOL REMOVE command
- XSPOOL SYNC command
- XSPOOL UNLOCK command
- XSPOOL XLIST command
- XSPOOL XLISTADD command
- XSPOOL XLISTDEL command

The following CP macros are no longer available for customer use:
- CSESYS macro
- CSETRACE macro
- CSEUSER macro

The following CP functions have been updated:
- CP_OWNED configuration statement
- DEFINE CPOWNED command

The following reserved DIAGNOSE code has been disabled:
- DIAGNOSE code X'F0'

[V6.3] HELPSEG and NLSAMENG Saved Segments No Longer Supplied with Base System

The HELPSEG and NLS American English (NLSAMENG) saved segments are no longer supplied with the z/VM base system. American English is the default language for all the z/VM components, so removing the NLSAMENG saved segment does not cause any loss of function. For preinstalled z/VM products, components, and features that placed data in these segments, the VMSES/E build lists that pertain to these segments are still provided. Similarly, the respective $PPF file part handler references necessary for building these segments have also been maintained, though in bypassed (inactive) form.

See *z/VM: Saved Segments Planning and Administration* for more information should you decide to define and build these segments for your system.

[V6.3] IPL Changes for NSS in a Linux Dump

A new option on the IPL command, NSSDATA, preserves the contents of a named saved system (NSS), if one is in use, in guest memory. This enhancement allows standalone dump tools, such as those used with the Linux operating system, to include the contents of the NSS in a standalone dump.

[V6.3] Memory Management

The z/VM memory management algorithms are redesigned to enable support for real memory up to 1 TB. These enhancements are intended to improve efficiency

for the overcommitment of virtual to real memory for guests and to improve performance. Some benefits of these memory management changes are:

- Improved efficiency in the management of memory over-commitment.
- Less need to use multiple LPARs by supporting more virtual machines in a single image and larger virtual machines in a single image, reducing system management complexity.

Preparing for the increased memory scalability of z/VM V6.3 will depend on your current hardware configuration for memory and CPUs. Depending on the workload you plan to run, you might need to add capacity to take full advantage of the 1 TB support provided with z/VM V6.3. This might mean adding memory and CPUs to your existing configuration, adding books or drawers to your system, or reconfiguring XSTORE and DASD paging space.

- To assist with planning for z/VM V6.3, additional information can be found at IBM: z/VM V6.3 Resources (www.ibm.com/vm/zvm630/).
- For XSTORE and 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*.

Two new commands are added:
- SET AGELIST
- QUERY AGELIST

The following commands are updated:
- DEFSEG
- DEFSYS
- INDICATE LOAD
- INDICATE NSS
- INDICATE SPACES
- INDICATE USER
- QUERY REORDER
- QUERY RESERVED
- SET CPTRACE
- SET REORDER
- SET RESERVED

The following configuration statements are added:
- STORAGE AGELIST
- STORAGE RESERVED SYSMAX

The following monitor records are changed:
- D0R3 - MRSYTRSG - Real Storage Data (global)
- D0R4 - MRSYTRSP - Real Storage Data (per processor)
- D0R6 - MRSYTASG - Auxiliary Storage (global)
- D0R7 - MRSYTSHS - Shared Storage Data
- D0R23 - MRSYTLCK - Formal Spin Lock Data (global)
- D1R15 - MRMTRUSR - Logged on User
- D2R4 - MRSCLADL - Add User To Dispatch List - Event Record
- D2R5 - MRSCLDDL - Drop User From Dispatch List - Event Record
- D2R6 - MRSCLAEL - Add User To Eligible List - Event Record
- D2R8 - MRSCLSTP - System Timer Pop - Event Record
- D3R1 - MRSTORSG - Real Storage Management (global)
- D3R2 - MRSTORSP - Real Storage Activity (per processor)
- D3R3 - MRSTOSHR - Shared Storage Management (per NSS or DCSS)

- D3R14 - MRSTOASI - Address Space Information Record
- D3R15 - MRSTOSHL - NSS/DCSS/SSP Loaded into Storage
- D3R16 - MRSTOSHD - NSS/DCSS/SSP Removed From Storage
- D4R2 - MRUSELOF - User Logoff Data - Event Record
- D4R3 - MRUSEACT - User Activity Data
- D4R9 - MRUSEATE - User Activity data at Transaction End - Event Record

The following CP trace codes have been removed: 4005, 4006, 4007, and 4008.

[V6.3] Performance Toolkit Enhancements

The Performance Toolkit includes new and updated reports containing the following information:
- CP Memory Scalability
- HiperDispatch
- Support for High Performance Ficon
- Support for HiperSockets Bridge
- Support for Live Guest Relocation

Support for the following monitor records have been added:
- Domain 3 Record 22 - MRSTOFRP - Frame Replenishment
- Domain 5 Record 15 - MRPRCDSV - Dispatch Vector Assignments (Event)
- Domain 5 Record 16 - MRPRCPUP - Park/Unpark Decesion (Event)
- Domain 5 Record 17 - MRPRCRCD - Real CPU Data (per CPU) (Sample)
- Domian 5 Record 18 - MRPRCDHF - Dispatch Vector High Frequency Data (Sample)

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

Support for rmfpms Stabilized

The Performance Toolkit can process Linux performance data obtained from the Resource Management Facility (RMF) Linux performance gatherer, rmfpms. Performance Toolkit support for rmfpms has been stabilized and might cease to function as the underlying Linux systems evolve. Support for the Linux rmfpms agent has been withdrawn and no new copies of it are available for installation. If you have rmfpms installed on an existing Linux image, it should continue to run on that image unsupported. There is no guarantee that a current rmfpms installation will run on future Linux image installations.

[V6.3] User Class Restructure (UCR) Support Removed

Support for the user class restructure (UCR) function and the OVERRIDE utility have been removed. Privilege classes for CP commands, DIAGNOSE codes, and other CP functions can be redefined by using MODIFY system configuration statements and MODIFY commands.

If any UCR files exist on the system, the contents of those files will not be processed by CP. However, CP will check for the existence of such files during system initialization and issue a message if any are found.

The following CP utility has been removed:
- OVERRIDE

The following CP utility has been updated:
- CVTOVRID

The following CP commands have been updated:
- LOCATE CMDBK
- PURGE UCR
- QUERY CPCMDS
- QUERY UCR

The following system configuration statements have been updated:
- MODIFY COMMAND / CMD statement
- MODIFY DIAGNOSE statement
- MODIFY PRIV_CLASS statement

The following directory statements have been updated:
- CLASS statement
- IDENTITY statement
- USER statement

The following DIAGNOSE code has been removed:
- X'C4'

[V6.3] DASD Volume Real Device Address

The real device address can be specified for a DASD volume to be included in the CP-owned list or the user volume list. An RDEV option has been added to the CP_OWNED system configuration statement, and there is a new USER_VOLUME_RDEV system configuration statement.

[V6.3 APAR] CPU Pooling Support

With the PTF for APAR VM65418, z/VM V6.3 allows you to define new entities called CPU pools. CPU pooling provides the ability to define a limit on the amount of CPU resources (in terms of real CP or IFL processors) that a group of z/VM guests are allowed to consume in aggregate. A CPU pool has a name and an associated capacity. For example, a CPU pool named GEORGE could be defined with a capacity of 1.5 IFL processors. One or more virtual machines can be assigned to a CPU pool (a virtual machine can be assigned to one CPU pool at a time) and have their aggregate CPU consumption limited to the pool's capacity. This support allows you to cap CPU utilization for sets of guests to better balance overall resource usage, for example to limit CPU charges based on what a set of guests is allowed to use.

Two types of CPU resource limits can be set for the group of users in a CPU pool:
- The group can be capped by the LIMITHARD method currently supported for individual users. (Individual user limits are set using the SET SHARE command, while group limits are set using the new DEFINE CPUPOOL and SET CPUPOOL commands.) This method caps the CPU pool to a specific percentage of the shared logical IFL or CP processors currently varied online. The group's allocation changes whenever the number of shared processors of the appropriate type changes due to actions like varying processors online or offline, or dedicating or undedicating processors.
- The group can be capped by the new CAPACITY method, which sets an amount of processor power equivalent to a specific number of real IFL or CP processors as the limit for the group. This method caps the CPU pool to an absolute

amount of processor resource, and the group's allocation is unaffected by changes in the number of online shared processors (unless the number falls below the limit for the group).

In a single system image (SSI) cluster, a guest assigned to a CPU pool can be relocated to another cluster member if a compatible CPU pool (with the same name and the same type of CPU being limited) exists on the destination member. Capacity limits are set and enforced independently on each cluster member.

For more information about using CPU pools, see *z/VM: Performance*.

The following CP commands have been added:
- DEFINE CPUPOOL
- DELETE CPUPOOL
- SET CPUPOOL
- QUERY CPUPOOL
- SCHEDULE

The following CP monitor records have been added:
- D1R28 - MRMTRCPC - CPU Pool Configuration
- D1R29 - MRMTRCPD - CPU Pool Definition - Event Record
- D4R13 - MRUSECPC - CPU Pool Change - Event Record
- D5R19 - MRPRCCPU - CPU Pool Utilization (Sample)

The following CP monitor records have been updated:
- D1R15 - MRMTRUSR - Logged On User
- D2R13 - MRSCLALL - Add VMDBK to the limit list - Event Record
- D2R14 - MRSCLDLL - Drop VMDBK from the limit list - Event Record
- D4R3 - MRUSEACT - User Activity Data

[V6.3 APAR] Environment Information Interface

With the PTF for APAR VM65419, z/VM V6.3 supports the new Store Hypervisor Information (STHYI) instruction, which 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 new interface, available to the guest in problem state, can be exploited by application tools that need to understand the configuration of the application used by z/VM guest operating systems. The applications need to understand what is deployed and the amount of capacity available to be consumed by those applications.

For more information about the STHYI instruction, see *z/VM: CP Programming Services*.

The following CP monitor record has been updated:
- D5R11 - MRPRCINS - Instruction Counts (per processor)

[V6.3 APAR] Additional Performance Toolkit Enhancements

With the PTF for APAR VM65357, the Performance Toolkit supports the following new screens and associated commands:
- FCX288, Multiprocessor User Activity Log Screen – USRMPLOG
- FCX315, Multiprocessor User Wait States Log Screen – USTMPLOG

With the PTF for APAR VM65399, the Performance Toolkit supports the following new screens and associated commands:

- FCX310, PCI Function Menu Screen – PCIMENU
- FCX311, PCI Function Configuration Screen – PCICONF
- FCX312, PCI Function Activity Screen – PCIACT
- FCX313, PCI Function Activity Log Screen – PCILOG
- FCX314, Pinned Storage Log Screen – PINLOG
- FCX316, Subpool Storage Log Screen – SUBPLOG

Support for the following new CP monitor records has been added:
- D1R27 - MRMTRPCI – PCI function Configuration Data
- D6R36 - MRIODPAT – Attach PCI Function
- D6R37 - MRIODPDT – Detach PCI Function
- D6R38 - MRIODPEN – Guest Enables a PCI Function
- D6R39 - MRIODPAC – PCI Activity
- D6R40 - MRIODPDS – Guest Disables a PCI Function
- D6R41 - MRIODPER – PCI function error
- D6R42 - MRIODPAD – PCI function added to the system
- D6R43 - MRIODPDL – PCI function deleted from the system
- D6R44 - MRIODPMD – PCI function program controls modified
- D6R45 - MRIODPON – Real PCI function varied on
- D6R46 - MRIODPOF – Real PCI function varied offline

With the PTF for APAR VM65527, the Performance Toolkit provides support for new and updated reports for the z13 server:
- FCX310, PCI Function Menu Screen – PCIMENU
- FCX311, PCI Function Configuration Screen – PCICONF
- FCX312, PCI Function Activity Screen, Format 0 – PCIACT
- FCX313, PCI Function Activity Log Screen, Format 0 – PCILOG
- FCX318, PCI Function Activity Screen, Format 1 – PCIACT ROCE
- FCX319, PCI Function Activity Log Screen, Format 1 – PCILOG
- FCX320, PCI Function Activity Screen, Format 2 – PCIACT ZEDC
- FCX321, PCI Function Activity Log Screen, Format 2 – PCILOG

With the PTF for APAR VM65528, the Performance Toolkit provides new and updated reports to support Multi-VSwitch Link Aggregation:
- FCX266, General Virtual Switch Description Screen – GVSWITCH
- FCX317, Global networking object activity – GLONACT

With the PTF for APAR VM65529, the Performance Toolkit provides new and updated reports to support simultaneous multithreading (SMT):
- FCX154, System Settings Screen – SYSSET
- FCX155, Monitor Data Statistics Screen – MONDATA
- FCX180, System Configuration Screen – SYSCONF
- FCX234, Processor Configuration Log Screen – PROCCONF
- FCX239, Processor Summary Log Screen – PROCSUM
- FCX265, Spin Lock Log Screen – LOCKLOG
- FCX268, General Virtual Network Device Description Screen – GVNIC
- FCX269, Virtual Network Device Activity Screen – VNIC
- FCX287, System Topology Machine Organization Screen – TOPOLOG
- FCX298, Logical PU Organization Log Screen – PUORGLOG
- FCX302, Real CPU Utilization Log Screen – PHYSLOG
- FCX303, DSVBK Steals per Processor Log Screen – DSVSLOG
- FCX304, Processor Log Screen – PRCLOG
- FCX306, Logical Partition Share Screen – LSHARACT

With the PTF for APAR VM65656, the Performance Toolkit supports using CMS pipelines as monitor data input to the MONSCAN/BATCH operation. This enhancement will allow Performance Toolkit to process multiple mondata files at one time, and the files may be packed.

With the PTF for APAR VM65697, the Performance Toolkit provides new and updated reports to support CPU pooling:
- FCX124, Performance Data Selection Menu Screen – MONITOR, MENU
- FCX126, LPAR Load Screen – LPAR
- FCX202, LPAR Load Log Screen – LPARLOG
- FCX226, User Configuration Screen – UCONF
- FCX306, Logical Partition Share Screen – LSHARACT
- FCX308, CPU Pool Configuration Screen – CPLCONF
- FCX309, CPU Pool Activity Screen – CPLACT
- FCX324, CPU Pool Menu Screen – CPLMENU

With the PTF for APAR VM65698, the Performance Toolkit provides new and updated reports to support guest usage of virtual PCIe functions:
- FCX310, PCI Function Menu Screen – PCIMENU
- FCX322, PCI Function Activity, Format 3 – PCIACT ISM
- FCX323, PCI Function Activity Log, Format 3 – PCILOG *pcifunc*

With the PTF for APAR VM65699, the Performance Toolkit implements improvements to the following reports:
- FCX103, Storage Utilization Screen – STORAGE
- FCX114, User Wait States Screen – USTAT / USTATG
- FCX126, LPAR Load Screen – LPAR
- FCX135, User Wait State Log Screen – USTLOG
- FCX179, System Facilities Log Screen – SYSLOG
- FCX180, System Configuration Screen – SYSCONF
- FCX202, LPAR Load Log Screen – LPARLOG
- FCX215, FICON Channel Load Screen – FCHANNEL
- FCX234, Processor Configuration Log Screen – PROCCONF
- FCX287, System Topology Machine Organization Screen – TOPOLOG
- FCX292, User Page Utilization Data Screen – UPGUTL
- FCX298, Logical PU Organization Log Screen – PUORGLOG
- FCX302, Real CPU Utilization Log Screen – PHYSLOG
- FCX306, Logical Partition Share Screen – LSHARACT
- FCX307, Logical Partition Logs Menu Screen – LPARLOGM
- FCX315, Multiprocessor User Wait States Log Screen – USTMPLOG *userid*

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

[V6.3 APAR] Support for IBM Cloud Manager with OpenStack for System z, V4.1

The PTFs for APARs VM65593, VM65614, VM65625, VM65628, VM65631, and VM65634 provide support for running IBM Cloud Manager with OpenStack for System z®, V4.1, as an appliance on z/VM V6.3. IBM Cloud Manager with OpenStack is designed to be an easy-to-deploy, simple-to-use cloud management offering that is built on OpenStack to help deliver new and improved private cloud and Service Provider solutions on IBM Z. New capabilities include the installation and operation of IBM Cloud Manager with OpenStack on z/VM, independent of your licensed Linux distribution, to manage the IBM Z environment.

IBM Cloud Manager with OpenStack for System z, V4.1, delivers:

- Full access to OpenStack Icehouse APIs, backed with IBM support
- An easy-to-deploy, simple-to-use cloud-management solution featuring a self-service portal for workload provisioning and virtualized image management
- Management of z/VM systems from the z/VM environment
- Support for deployment, resize and capture, backup and restore, image management, approvals, expiration, billing, and metering

For more information about enabling this appliance, see:

- *z/VM: Enabling z/VM for OpenStack (Support for OpenStack Icehouse Release)*
- *z/VM: Systems Management Application Programming*

[V6.3 APAR] Support for IBM Cloud Manager with OpenStack for z Systems, V4.2

The PTF for APAR VM65669 provides support for running IBM Cloud Manager with OpenStack for z Systems, V4.2, as an appliance on z/VM V6.3.

IBM Cloud Manager with OpenStack for z Systems, V4.2, delivers:

- Support for the OpenStack Juno release.
- Support for managing from IBM Z for Systems p/x/flex.
- Support for managing multiple regions from IBM Z.
- Support for multipath and NPIV Cinder volumes.
- Support for the IBM Platform Resource Scheduler.
- Support for OpenStack Heat.
- Support for OpenStack Horizon.
- Chef server and client enhancements to install IBM Cloud Manager with OpenStack for z Systems.
- "Manage to" support allows multiple compute nodes on the same virtual machine (not available in the Cloud Manager Appliance).

Additional enhancements to the Cloud Manager Appliance for z/VM include:

- Boot from volume support
- Enhancements to the z/VM OpenStack/xCAT Installation Verification Scripts

xCAT enhancements include:

- xCAT and ZHCP user exits space recovery on start up of the xCAT and ZHCP servers
- xCAT GUI support to add/remove a volume to/from the z/VM system configuration file

For more information, see:

- *z/VM: Enabling z/VM for OpenStack (Support for OpenStack Juno Release)*
- *z/VM: Systems Management Application Programming*

[V6.3 APAR] xCAT Enhancements

The PTF for APAR VM65667 provides support for the following enhancements to the Extreme Cloud Administration Toolkit (xCAT) appliance shipped with z/VM:

- Exploitation of multiple OSA ports
- Management of workloads on multiple LAN segments

- Exploitation of VLAN-aware clients

These benefits then accrue to the z/VM OpenStack support.

For more information, see:

- *z/VM: Enabling z/VM for OpenStack (Support for OpenStack Icehouse Release)* or *z/VM: Enabling z/VM for OpenStack (Support for OpenStack Juno Release)*
- *z/VM: Systems Management Application Programming*

[V6.3] IBM Cloud Manager with OpenStack, V4.3

IBM Cloud Manager with OpenStack, V4.3, supports the management of z/VM V6.3 systems from Power Systems™ and x86.

Note: z/VM V6.3 systems to be managed from Power® or x86 should ensure that the latest xCAT service is installed as described at IBM: z/VM xCAT Maintenance (www.ibm.com/vm/sysman/xcmntlvl.html).

For more information, see:

- *z/VM: Enabling z/VM for OpenStack (Support for OpenStack Kilo Release)*
- *z/VM: Systems Management Application Programming*

[V6.3] z/VM V6.3 Achieves Common Criteria Certification

All certification activities for z/VM 6.3 are complete. The certifying body issued its certification on March 30, 2015. z/VM 6.3, 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+). For more information about this announcement, see IBM z/VM Security and Integrity resources (www.ibm.com/vm/security/).

[V6.3 APAR] Virtual Machine Logoff Time Reduction

The PTF for APAR VM64770 reduces the time to perform a virtual machine logoff or other system-reset-clear operation in a storage overcommitted environment by providing a read ahead function to batch PGMBK reads from DASD asynchronously from the main storage releasing thread. This parallelism makes the algorithm more efficient, although the observed benefit is still dependent on system constraints, specifically:

- The availability of contiguous pairs of frames for reading in the PGMBKs.
- Any paging backlog. If the read queue has built up on the target volume(s), it will inhibit the effectiveness of the read ahead task.

[V6.3 APAR] RACF Security Policy and Password Encryption Enhancements

With the PTF for APAR VM65719, the z/VM RACF Security Server feature supports the following security policy and password encryption enhancements:

- Security policy enhancements
  - Minimum password or password phrase change interval

    The SETROPTS PASSWORD command has the capability of setting a minimum number of days that must pass between a user's password or password phrase changes.
  - ALTUSER Revoke/Resume date management

New NOREVOKE and NORESUME keywords for the ALTUSER and CONNECT commands will provide better revoke/resume date management. The NOREVOKE and NORESUME keywords can be used to clear the revoke and resume dates rather than clearing them with the REVOKE and RESUME keywords. Also, LISTUSER and LISTGRP are changed to always display the revoke and resume dates if present, and LISTGRP is changed to display the group's creation date, which is already present in the group profile, but not currently shown.

– RACUT200 Database Verification utility to perform a reserve/release

  The RACUT200 Database Verification utility has been modified to perform a reserve on the RACF database minidisk prior to executing any of its functions and to perform a release on the RACF database minidisk prior to exiting.

– New information from DIAGNOSE X'A0' subcode X'50'

  New RACF configuration information has been added for retrieval from subcode X'50':
  - SMF CPUID
  - SMF current disk
  - SMF current disk owner
  - SMF data file name

  Note: This enhancement is provided by the PTF for APAR VM65498.

– Help desk support

  The following new FACILITY class resources have been added to allow delegation of authority to non-security administrator users, such as help desk personnel, to reset passwords and password phrases and list user information.
  - IRR.PASSWORD.RESET
  - IRR.PWRESET.OWNER
  - IRR.PWRESET.TREE
  - IRR.PWRESET.EXCLUDE
  - IRR.LISTUSER
  - IRR.LU.OWNER
  - IRR.LU.TREE
  - IRR.LU.EXCLUDE

• Password encryption enhancements

– New password encryption algorithm

  A new KDFAES (key derivation function with advanced encryption standard) password encryption algorithm can be specified on the SETROPTS command. The KDFAES algorithm is more secure than the default DES algorithm, but is computationally more intensive, by design. If KDFAES is active, a new PWCONVERT keyword on the ALTUSER command can be used to convert legacy-format passwords and password history entries to KDFAES format.

– Special characters in passwords

  A new SPECIALCHARS option on the SETROPTS command allows the use of 14 additional special characters in passwords and password phrases. New password syntax rule content-keywords SPECIAL and MIXEDALL have been added to the SETROPTS command, and the meaning of the NOVOWEL content-keyword has been extended to include special characters. MIXEDALL can be used to force the use of special characters in passwords. The additional password characters are:

  ```
  . < + | & ! * - % _ > ? : =
  ```

– Expire password support

The ALTUSER command can be used to mark a user's existing password or password phrase as expired without specifying a new password or phrase. This can be useful to force users to change passwords as soon as possible after a change in password policy, such as the implementation of a new syntax rule. It can also be useful to evenly distribute password change dates across a user population.

– Password history cleanup

A new PWCLEAN keyword on the ALTUSER command can be used to clean up residual entries in a user's password and password phrase history resulting from lowering the SETROPTS PASSWORD(HISTORY(nn)) value.

The following z/VM RACF commands have been updated:
- ALTUSER (Alter User Profile)
- CONNECT (Connect User to Group)
- LISTGRP (List Group Profile)
- LISTUSR (List User Profile)
- PASSWORD or PHRASE (Special User Password or Password Phrase)
- SETROPTS (Set RACF Options)

For more information, see:
- *z/VM: RACF Security Server Command Language Reference*
- *z/VM: RACF Security Server Macros and Interfaces*
- *z/VM: RACF Security Server Security Administrator's Guide*
- *z/VM: RACF Security Server System Programmer's Guide*

[V6.3 APAR] xCAT/OpenStack Discovery Tool and Cloud Manager Appliance Enhancements

With the PTFs for APARs VM65723 and VM65724, support is added to allow OpenStack and xCAT to discover Linux on z Systems servers on a z/VM hypervisor and then to add those servers to their respective databases. At the end of this process the servers will be unlocked and ready to be managed.

This support includes a new z/VM systems management API:
- Virtual_Network_Adapter_Query_Extended

The xCAT xdsh command is updated with new options, and the following xCAT commands are added:
- nodediscoverdef
- nodediscoverls
- nodediscoverstart
- nodediscoverstatus
- nodediscoverstop

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

The PTF for APAR VM65724 also provides several fixes and enhancements for the z/VM Cloud Manager Appliance (CMA) 4.2.0.2. These changes include updates to OpenStack to respond to the discovery commands coming from xCAT, as well as enhancements to simplify installation. For more information on the updated CMA 4.2.0.2 installation procedure, see CMA42 FILE on the MAINT 400 disk.

[V6.3 APAR] Enhanced z/VM OpenStack Juno Support

With the PTF for APAR VM65753, z/VM support for OpenStack Juno is enhanced to include the following functions:

- Support for reconfiguring the CMA to make z/VM manageable by an external cross-platform OpenStack controller

  You can use the ibm-openstack-zvm-driver-compute-anywhere cookbook to reconfigure the CMA so that the CMA compute role and the CMA compute_mn role are managed by an external cross-platform OpenStack controller, which contains OpenStack controller components such as glance, cinder, keystone, neutron-server, and so on. The CMA compute_mn role runs the nova-compute service, the neutron-zvm-agent service, and xCAT.
- Remote keystone support

  Remote keystone support allows you to use the ibm-openstack-zvm-appliance-external-keystone cookbook to configure the CMA to use a remote (external) keystone server for the authentication and authorization of OpenStack services.
- Default network support

  The default network feature allows you to define a default network on the CMA, and thus reduce the number of manual steps involved in this process.
- Multipath support

  Multipath support improves the reliability of persistent disks by allowing you to define two paths to a disk. If one of the paths to the disk fails, the host will access the disk using the second path. Applications running on the host are not interrupted.

For more information, see *z/VM: Enabling z/VM for OpenStack (Support for OpenStack Juno Release)*.

[V6.3 APAR] Relocating the SSI PDR without a Planned Outage

The PTF for APAR VM65712 adds support for the new PDRVOLUME operand on the SET SSI command, which allows you to relocate the z/VM single system image (SSI) persistent data record (PDR) to another DASD volume without a planned outage. The FORMSSI utility has been updated to display the replacement volume ID.

The following CP interface has been added:

- Monitor record Domain 11 Record 8 - MRSSIPDR - SSI PDR volume change

The following CP interfaces have been updated:

- SET SSI command
- FORMSSI utility

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Planning and Administration*

[V6.3 APAR] z/VM support for OpenStack Liberty

With the PTFs for APARs VM65780 and VM65807, z/VM provides support for OpenStack Liberty, which includes the following support:

- Support for the Liberty release of OpenStack

- Integration of xCAT function into the z/VM Cloud Manager Appliance (CMA), which allows running a fully functional z/VM OpenStack solution in a single virtual server, so separate ZHCP servers are not required

  For information to help you migrate from using separate z/VM user IDs for the xCAT and ZHCP servers to using one z/VM user ID for both, see Appendix B, "[V6.3] Migrating to an Integrated xCAT MN and ZHCP Server in the Same CMA," on page 461.
- Support for provisioning Red Hat RHEL 7 and SUSE SLES 12 servers (previously only servers up to RHEL 6.5 and SLES 11 were supported)
- Support for distributed keystone, so identity and security credentials do not have to be in a single place
- Support for Ceilometer, which can collect measurements to be used by performance monitoring tools
- Improved boot from volume (an alternative OpenStack way to boot servers)
- Install Verification Program enhancements

For more information, see:

- *z/VM: Enabling z/VM for OpenStack (Support for OpenStack Liberty Release)*
- *z/VM: Systems Management Application Programming*

[V6.4] DirMaint Support

The Directory Maintenance Facility (DirMaint), function level 640, includes the following changes:

- DIRM GLOBALOPTS, MACHINE, and SETMACH commands have been updated to allow the definition of a z/Architecture-only virtual machine (MACHINE Z).
- New exits and configuration data specifications are added to the interface between DirMaint and RACF/VM for link and NICDEF creation and deletion, allowing for further coordination of CP and RACF privileges, and discrete resource profile creation and deletion.
- The XSTORE operand of the DIRMAINT command has been removed due to the removal of Expanded Storage (XSTORE) support.
- DirMaint messages DVH3261I and DVH3324E have been removed due to the removal of Expanded Storage (XSTORE) support.

For more information, see:

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

[V6.4] 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.

A new FLASH operand is added to the following:

- SET EDEVICE command
- EDVICE configuration file statement

In addition:

- The QUERY EDVICE command has been enhanced to display information about emulated FLASH devices.
- The systems management System_EQID_Query API has been updated.

The following monitor record has been updated:

- Domain 1 Record 6 - MRMTRDEV - Device Configuration Data

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Planning and Administration*
- *z/VM: Systems Management Application Programming*.

[V6.4] z/VM System SSL and Utilities z/OS V2.2 Equivalency

The z/VM V6.4 System SSL cryptographic library has been upgraded to z/OS V2.2 equivalence for the TLS/SSL and LDAP servers. This includes AES_GCM support previously provided via PTF to V6.3. Additionally, GSKKYMAN can now support RFC 5280 certificates. For more information, see *z/VM: TCP/IP User's Guide*.

The TLS/SSL Server defaults have also been updated so that TLS 1.1 and TLS 1.2 are enabled by default, and all other protocol levels are disabled. These changes, along with a shift in which cipher suites are enabled by default, represent a more aggressive cryptographic security policy for IBM z/VM. Be advised that the change prevents communication between V6.4 and V6.3 unless TLS/SSL Server configuration changes are made on one or both sides of the connection.

Additionally, the following new features were added to the TLS/SSL Server:

- Cipher management: an ENABLE operand to restore cipher suites which are disabled by default.
- Compatibility: A PKCS #12 file of type .P12 can now be used in place of a certificate database file (.KDB) in configuring certificate support.

For more information about the new defaults and new features, see *z/VM: TCP/IP Planning and Customization*.

[V6.4] 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.

The following CP configuration statements have been updated:

- DEFINE VSWITCH
- MODIFY VSWITCH

The following CP user directory statement has been updated:

- NICDEF

The following CP commands have been updated:

- COUPLE
- DEFINE NIC
- DEFINE VSWITCH
- QUERY CONTROLLER
- QUERY LAN

- QUERY OSA
- QUERY VIRTUAL NIC
- QUERY VMLAN
- QUERY VSWITCH
- SET NIC
- SET VMLAN
- SET VSWITCH

The following diagnose code has been updated:
- DIAGNOSE Code X'26C' - Access Certain System Information

The following have been changed to general purpose vswitch controllers:
- DTCENS1 - has been renamed to DTCVSW3
- DTCENS2 - has been renamed to DTCVSW4

All VSWITCH controllers now have share 3000 relative.

The following TCP/IP configuration statements have been updated:
- DEVICE and LINK statements for OSD devices
- PERMIT
- VSWITCH CONTROLLER

The following TCP/IP commands have been updated:
- IFCONFIG
- NETSTAT CONFIG
- NETSTAT DEVLINKS

The following CP monitor records have been updated:
- Domain 1 Record 4 - MRMTRSYS - System Configuration Data
- Domain 1 Record 19 - MRMTRQDC - QDIO Device Configuration
- Domain 6 Record 23 - MRIODVSR - Virtual Switch Recovery
- Domain 6 Record 25 - MRIODQDA - QDIO Device Activation Event
- Domain 6 Record 27 - MRIODQDD - QDIO Device Deactivation Event
- Domain 8 Record 1 - MRVNDSES - Virtual NIC Session Activity
- Domain 8 Record 2 - MRVNDLSU - Virtual NIC Guest Link State - Link Up
- Domain 8 Record 3 - MRVNDLSD - Virtual NIC Guest Link State - Link Down

For more information, see:
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Planning and Administration*
- *z/VM: CP Programming Services*
- *z/VM: TCP/IP Planning and Customization*
- *z/VM: TCP/IP User's Guide*

[V6.4] 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.

[V6.4] 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.

[V6.4] 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.

z/VM supports more than one thread per core for only IFLs. Although the command allows a request of two threads per core for the other CPU types, z/VM will not use more than one thread per core for those CPU types.

- The following new CP command was added for this support:
  - SET MULTITHREAD
- The following CP commands have been updated for this support:
  - INDICATE MULTITHREAD
  - QUERY MULTITHREAD
  - SET CPTRACE
  - VARY CORE
- The MULTITHREADING configuration statement has been updated.
- New CP trace code has been added:
  - Threading level change - 3614
- The following new monitor record has been added:
  - Domain 5 Record 21 - MRPRCSMT - SMT Configuration Change Event
- The following monitor records have been updated:
  - Domain 0 Record 2 - MRSYTPRP - Processor data (per processor)
  - Domain 1 Record 4 - MRMTRSYS - System Configuration Data
  - Domain 5 Record 1 - MRPRCVON - Vary On Processor
  - Domain 5 Record 2 - MRPRCVOF - Vary Off Processor
  - Domain 5 Record 20 - MRPRCMFM - MT CPUMF Counters

For more information, see

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Messages and Codes*
- *z/VM: CP Planning and Administration*
- *z/VM: Diagnosis Guide*

[V6.4] 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.

The following CP command has been added:
- QUERY SHUTDOWN

The following CP commands have been updated for this support:
- FORCE
- QUERY SHUTDOWNTIME
- QUERY SIGNALS
- SHUTDOWN
- SIGNAL

For more information, see:
- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Messages and Codes*.

[V6.4] Control Program Environment Variables

With this support, system programmers and Class B virtual machines can set system-wide environment variables. Any Class G virtual machine on the system can query the values of these 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.

The following CP commands have been added for this support:
- QUERY VARIABLE
- SET VARIABLE

The following CP command has been updated:
- SET IPLPARMS

The following CP utility has been updated:
- CPSYNTAX

The following CP configuration statement has been added:
- SET VARIABLE

The following IPL parameter has been added:
- IPLVAR=*value*

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Messages and Codes*
- *z/VM: CP Planning and Administration*
- *z/VM: System Operation*.

[V6.4] Virtual Processor Management Improvement

CP's virtual processor management has been improved so that no user stays in the eligible list more than an instant before being added to the dispatch list. Therefore some functions intended to improve performance by managing the eligible list are now less meaningful.

The following CP functions are affected:
- QUICKDSP user directory option
- INDICATE ACTIVE command
- INDICATE LOAD command
- INDICATE QUEUES command
- SET QUICKDSP command
- SET SRM command

[V6.4] 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 setting. This is managed independently for each processor type (General Purpose, IFL, zIIP, and so on) across the virtual machines.

The following CP monitor records have been updated:
- Domain 2 Record 4 - MRSCLADL - Add User to Dispatch List
- Domain 2 Record 5 - MRSCLDDL - Drop User from Dispatch List
- Domain 2 Record 6 - MRSCLAEL - Add User to Eligible List
- Domain 2 Record 13 - MRSCLALL - Add VMDBK to the Limit List
- Domain 2 Record 14 - MRSCLDLL - Drop User from the Limit List

[V6.4] 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.

For more information, see:
- *z/VM: RACF Security Server Command Language Reference*
- *z/VM: RACF Security Server Security Administrator's Guide*

[V6.4] 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 CPOWNED 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

Attention: 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 following CP commands have been added for this support:
- QUERY PAGING
- SET PAGING

The following CP commands have been updated:
- ATTACH
- QUERY AGELIST
- QUERY CU
- SET AGELIST
- SET CU
- SET IPLPARMS

The following CP configuration statements have been updated:
- CU
- FEATURES
- STORAGE AGELIST

The following monitor records are changed:
- D0R23 - MRSYTLCK - Formal Spin Lock Data
- D1R7 - MRMTRMEM - Memory Configuration Data
- D1R20 - MRMTRHPP - HyperPAV Pool Definition
- D3R4 - MRSTOASP - Auxiliary Storage Management
- D3R11 - MRSTOASS - Auxiliary Shared Storage Management
- D6R3 - MRIODDEV - Device Activity
- D6R28 - MRIODHPP - HyperPAV Pool Activity
- D6R32 - MRIODHPF - Indicates an HPF Feature Change

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Messages and Codes*
- *z/VM: CP Planning and Administration*
- *z/VM: System Operation*.

[V6.4] Performance Toolkit Enhancements

Performance Toolkit for VM is enhanced to exploit z/Architecture 64-bit addressing mode and its expanded set of instructions. Consequently, the PERFSVM virtual machine must run on z/Architecture CMS (z/CMS). Any virtual machine running the Performance Toolkit module must also be using z/CMS.

Performance Toolkit also exploits virtual storage above the 2 GB line, or High Memory Area (HMA), for its temporary data areas, especially for print buffers, thus freeing the conventional virtual storage that resides below the 2 GB line.

The following command is new:
- FCONTROL HMA

The following command is updated:
- FCONTROL LIMIT

Various messages have been added or updated.

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

[V6.4] Performance Toolkit Changes in How Spin Lock Activity Is Reported

Starting in z/VM V6.3, the number of spin locks reported in monitor record D0R23 MRSYTLCK began growing. This caused large LOCKLOG reports and also increased the memory Performance Toolkit needed to do a reduction. To provide relief, IBM replaced the log-style LOCKLOG report with an activity-style report called LOCKACT. The LOCKACT report shows mean lock behavior. If interval-to-interval lock statistics are needed, you can generate an INTERIM LOCKACT report.

When z/VM V6.4 Performance Toolkit reduces data from a z/VM V6.4 or later system, it can produce only the LOCKACT and INTERIM LOCKACT reports. When z/VM V6.4 Performance Toolkit reduces data from a pre-V6.4 system, it can produce LOCKLOG, LOCKACT, and INTERIM LOCKACT reports.

The following commands are new:
- CPUMENU
- LOCKACT

The following command is updated:
- LOCKLOG

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

[V6.4] Nondefault Tape Mode Support

The z/VM Tape Products Architecture (TPA) provides additional options to further specify MODE settings for IBM System Storage 3592 TS1120 (E05), TS1130 (E06), and TS1140 (E07) Tape Drives. This support provides flexibility in setting the Write format for a tape device. This extends the use and interoperability of cartridges between different tape hardware subsystems.

The following CP command has been updated for this support:
- ATTACH

The following CP utility has been updated for this support:

- DDR

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Messages and Codes*

[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.

In general, applications created in earlier z/VM releases that use CMS Pipelines as specified will continue to run without any change to the application. For additional migration information, see *z/VM: CMS Pipelines User's Guide and Reference*. This new CMS Pipelines publication is based on *CMS/TSO Pipelines: Author's Edition*, SL26-0018, and replaces both of the previous z/VM CMS Pipelines publications.

All of the CMS Pipelines HELP files have been replaced with files that correspond to the new publication. CMS Pipelines messages and HELP files are no longer translated into Japanese (KANJI).

[V6.4] RACF ICHRCX02 Exit Disabled by Default

IBM recommends that the ICHRCX02 RACROUTE REQUEST=AUTH postprocessing exit be disabled, however prior to this release it has been shipped enabled. From this release onwards the exit will be shipped disabled.

[V6.4] 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. This eliminates the risk of receiving HCD message CBDD824 CONFIGURATION CHANGE TOO LARGE FOR CP TO HANDLE, in most cases. 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.

[V6.4] IBM Cloud Manager V4.2 on System z

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.

[V6.4] z/VM V6.4 Support for OpenStack Liberty

Attention

z/VM V6.4 does not support OpenStack releases prior to Liberty.

If you are running a Liberty Cloud Management Appliance (CMA) on your z/VM V6.3 system, you need to migrate that CMA to z/VM V6.4. For instructions, see Appendix C, "[V6.4] Migrating an OpenStack Liberty CMA from z/VM V6.3 to z/VM V6.4," on page 465.

z/VM V6.4 support for OpenStack Liberty includes:

- Support for alternate deployment provisioning (cloning).
- z/VM CMA contained in the OPNCLOUD virtual machine.

For more information, see:

- *z/VM: Enabling z/VM for OpenStack (Support for OpenStack Liberty Release)*
- *z/VM: Systems Management Application Programming*

Also see "[V6.4] z/VM OpenStack Documentation" on page 277.

[V6.4] 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, since the data being exported is a copy of the logical volume, the logical volume data remains accessible by the production host systems.

For more information, see *z/VM: DFSMS/VM Removable Media Services*.

[V6.4 APAR] z/VM V6.4 Support for OpenStack Newton

With the PTF for APAR VM65893, the OpenStack Cloud Manager Applicance (CMA) packaged with z/VM V6.4 is updated to support the Newton level of OpenStack. Enhancements include:

- A new configuration wizard to configure the CMA.
- Support for provisioning Ubuntu 16.04 servers, in addition to RHEL 6 and 7 and SLES 11 and 12.
- The Inter-User Communications Vehicle (IUCV) replaces ssh as the default choice for all communications between the CMA and deployed guests.
- The Chef client has been removed and replaced with scripts to implement configuration changes.
- Logging and serviceability have been enhanced.
- The installation verification program (IVP) now runs automatically at configurable intervals, and it will also send reports to a z/VM user ID.

For more information, see:

- *z/VM: Enabling z/VM for OpenStack (Support for OpenStack Newton Release)*
- *z/VM: Systems Management Application Programming*
- *z/VM: CMS and REXX/VM Messages and Codes*

With the PTF for APAR VM65996, the CMA support for Newton is enhanced:

- Installation verification program (IVP) updates to clean up messages, disable unneeded tests, improve usability of the IVP user interface, support disabling of scheduled IVPs.
- Boot from volume is now supported on all types of guests that can be deployed from OpenStack (Red Hat, SLES and Ubuntu).
- Fix to rotate and clean up httpd logs to match how other CMA log files are rotated.
- Fix to allow a deploy instance whose name is less than 8 characters.

You can find more information about these enhancements at https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/W21ed5ba0f4a9_46f4_9626_24cbbb86fbb9/page/CMA Newton Documentation Updates.

Deprecated Interfaces

IBM has deprecated the following interfaces starting with the Newton level:

- The `openstack_xcat_mgt_ip` property in the DMSSICMO COPY file.
- The `openstack_xcat_mgt_mask` property in the DMSSICMO COPY file.

IBM fully supports interfaces first deprecated in this OpenStack level and intends to fully support them in any fix packs for this level, but IBM components may ignore them in any future level, or require that you remove them as part of upgrading to any future level. Whenever you must take action to stop using deprecated interfaces, you can find the planning and implementation information for doing so in Appendix D, "[V6.4 APAR] xCAT Management Network Changes in the OpenStack Newton CMA," on page 469 on the Web at http://www.vm.ibm.com/sysman/osmntlvl.html.

[V6.4 APAR] RACF Security Policy Enhancements

With the PTFs for APARs VM65930 and VM65982, the z/VM RACF Security Server feature supports the following security policy enhancements:

- Read-Only Auditor (ROAUDIT)

  This new user role allows many of the common auditing tasks to be performed without the ability to modify settings or manipulate audit logs.

- XAUTOLOG..ON control

  This enhancement introduces new security policy requirements for the ON operand of the CP XAUTOLOG command. Currently, the ON operand acts similarly to a LOGON..THERE function, whereby a privileged user can initiate a connection between an existing rdev (terminal window) and a virtual machine. The security controls will disallow this by default, unless appropriate PERMITs have been granted to the console's resource profile.

  Attention: The introduction of the XAUTOLOG..ON control will cause the CP XAUTOLOG ON command to fail under all circumstances in installations where RACF/VM is enabled, until some RACF configuration is performed. Customers desiring XAUTOLOG..ON to continue to work as it currently does need to enable generic command processing and create a generic VMCMD resource profile named XAUTOLOG.ON.** that has universal READ access. For example:

  ```
  RDEFINE VMCMD XAUTOLOG.ON.** UACC(READ)
  ```

- List the current VMXEVENT profile

This enhancement updates the SETEVENT LIST command to provide an authorized user with the names of the VMXEVENT profiles activated and in use by RACF.

Note: The PTF for APAR VM65923 provides infrastructure support in z/VM V6.2, V6.3, and V6.4 and must be installed on all members of an SSI cluster before any V6.4 member is running with the PTF for APAR VM65930.

The following z/VM RACF commands have been updated:
- ADDUSER (Add User Profile)
- ALTUSER (Alter User Profile)
- LDIRECT (List SFS Directory Profile)
- LFILE (List SFS File Profile)
- LISTDSD (List Data Set Profile)
- LISTGRP (List Group Profile)
- LISTUSER (List User Profile)
- RLIST (List General Resource Profile)
- SEARCH (Search RACF Database)
- SETEVENT (Set z/VM Events)
- SETROPTS (Set RACF Options)
- SRDIR (Obtain a List of SFS Directory Profiles)
- SRFILE (Obtain a List of SFS File Profiles)

For more information, see:
- *z/VM: RACF Security Server Auditor's Guide*
- *z/VM: RACF Security Server Command Language Reference*
- *z/VM: RACF Security Server Macros and Interfaces*
- *z/VM: RACF Security Server Security Administrator's Guide*
- *z/VM: RACF Security Server System Programmer's Guide*
- *z/VM: Security Server RACROUTE Macro Reference*

The following CP command has been updated:
- XAUTOLOG

For more information, see:
- *z/VM: CP Commands and Utilities Reference*

[V6.4 APAR] z/VM HELP Quality Improvement

With the PTFs for APARs VM66028 and VM66029, along with the PTF for VM66053 (a VMSES/E required if-req), new HELP components are defined and the help files for some CP and CMS commands are renamed, so the HELP commands used to display those help files are more logical and consistent with the commands used to display similar help files.

The major operands (subcommands) of some CP and CMS commands are documented as separate commands with two-word names, such as DEFINE ALIAS. The HELP command syntax supports only single-word command names, so HELP components are defined for these subcommands. A HELP component is a group of help files with a common HELP*xxxx* file type that uses the first four letters of the HELP component name, which is usually the name of the parent command, and the file name is the subcommand name. For example, the HELP component for the DEFINE subcommands is DEFINE, and the file ID of the help file for the DEFINE ALIAS command is ALIAS HELPDEFI.

To display the help file for a subcommand, enter:

```
help HELP_component_name command_name
```

This syntax allows you to enter a help request for a subcommand that is logical because it simulates the two-word command name, for example:

```
help define alias
```

The help files for some CP and CMS subcommands are defined in the general CP and CMS HELP components, so the file type is HELPCP or HELPCMS, which does not identify the command, and the subcommand is identified only by the file name. To display the help file for one of these subcommands, you need to know what the file name is. For example, the file ID of the help file for the ASSOCIATE EXIT subcommand is ASSOEX HELPCP, so the command to display the help file is:

```
help cp assoex
```

This HELP enhancement removes those subcommand help files from the CP and CMS HELP components and defines new components for them. New (renamed) help files are provided that enable you to issue more logical HELP commands to display them, like the commands used to display the help files for other CP and CMS subcommands. For example, the command to display the new help file for the ASSOCIATE EXIT subcommand is:

```
help associate exit
```

New and updated menus and abbreviation files are also included.

Table 6 shows the affected commands. The table shows the new HELP components (same name as the parent command), the file IDs of the new (renamed) help files for the subcommands within those HELP components, and the resulting improved (more logical) HELP commands used to display the help files.

Notes:

1. The HELP command is not changed.
2. The file IDs of the help files for the parent commands are not changed. Those help files are still included within the CP and CMS components. For example, the help file for the CP ASSOCIATE command is ASSOCIAT HELPCP and can be displayed with the following command:

   ```
   help cp associate
   ```
3. Because a help file on the system cannot be renamed via service, a new (renamed) help file and a revised version of the old help file are provided. The old file is provided for compatibility, and has been revised to contain only a statement that the file is obsolete and a link to the new file. The old files are deprecated and will be removed from the help database in the next z/VM release.

Table 6. Renamed Help Files for Some CP and CMS Subcommands

Command [Component]	**Subcommand**	**Help File ID**	**Command Used to Display Help File**
CP ASSOCIATE [ASSOCIATE]	EXIT	Old: ASSOEX HELPCP New: EXIT HELPASSO	Old: `help cp assoex` New: `help associate exit`
	MESSAGES	Old: ASSOMSGS HELPCP New: MESSAGES HELPASSO	Old: `help cp assomsgs` New: `help associate messages`

Table 6. Renamed Help Files for Some CP and CMS Subcommands (continued)

Command [Component]	Subcommand	Help File ID	Command Used to Display Help File
CP DEACTIVE [DEACTIVE]	CONV	Old: CONV HELPCP New: CONV HELPDEAC	Old: `help cp conv` New: `help deactive conv`
	ISLINK	Old: ISLINK HELPCP New: ISLINK HELPDEAC	Old: `help cp islink` New: `help deactive islink`
CP DRAIN [DRAIN]	(Disk)	Old: DRAINDIS HELPCP New: DISK HELPDRAI	Old: `help cp draindis` New: `help drain disk`
	UR	Old: DRAINUR HELPCP New: UR HELPDRAI	Old: `help cp drainur` New: `help drain ur`
CP FLASHCOPY [FLASHCOPY]	BACKGNDCOPY	Old: FLASHBND HELPCP New: BACKGNDC HELPFLAS	Old: `help cp flashbnd` New: `help flashcopy backgndcopy`
	ESTABLISH	Old: FLASHEST HELPCP New: ESTABLIS HELPFLAS	Old: `help cp flashest` New: `help flashcopy establish`
	RESYNC	Old: FLASHRES HELPCP New: RESYNC HELPFLAS	Old: `help cp flashres` New: `help flashcopy resync`
	TGTWRITE	Old: FLASHTGT HELPCP New: TGTWRITE HELPFLAS	Old: `help cp flashtgt` New: `help flashcopy tgtwrite`
	WITHDRAW	Old: FLASHWIT HELPCP New: WITHDRAW HELPFLAS	Old: `help cp flashwit` New: `help flashcopy withdraw`
CP FREE [FREE]	LOGON	Old: FREELOGN HELPCP New: LOGON HELPFREE	Old: `help cp freelogn` New: `help free logon`
	PRINTER	Old: FREEPRIN HELPCP New: PRINTER HELPFREE	Old: `help cp freeprin` New: `help free printer`
CP GIVE [GIVE]	(Real Device)	Old: GIVRDEV HELPCP New: RDEV HELPGIVE	Old: `help cp givrdev` New: `help give real`
	(Virtual Device)	Old: GIVVDEV HELPCP New: VDEV HELPGIVE	Old: `help cp givvdev` New: `help give virtual`
CP HOLD [HOLD]	LOGON	Old: HOLDLOGN HELPCP New: LOGON HELPHOLD	Old: `help cp holdlogn` New: `help hold logon`
	PRINTER	Old: HOLDPRT HELPCP New: PRINTER HELPHOLD	Old: `help cp holdprt` New: `help hold printer`
CP START [START]	(Disk)	Old: STARTDIS HELPCP New: DISK HELPSTAR	Old: `help cp startdis` New: `help start disk`
	UR	Old: STARTUR HELPCP New: UR HELPSTAR	Old: `help cp startur` New: `help start ur`

Table 6. Renamed Help Files for Some CP and CMS Subcommands (continued)

Command [Component]	Subcommand	Help File ID	Command Used to Display Help File
CP VARY [VARY]	CHPID	Old: VARYCHPI HELPCP New: CHPID HELPVARY	Old: `help cp varychpi` New: `help vary chpid`
	CORE	Old: VARYCORE HELPCP New: CORE HELPVARY	Old: `help cp varycore` New: `help vary core`
	PATH	Old: VARYPATH HELPCP New: PATH HELPVARY	Old: `help cp varypath` New: `help vary path`
	PCIFUNCTION	Old: VARYPCIF HELPCP New: PCIFUNCT HELPVARY	Old: `help cp varypcif` New: `help vary pcifunction`
	PROCESSOR	Old: VARYPROC HELPCP New: PROCESSO HELPVARY	Old: `help cp varyproc` New: `help vary processor`
	(Real Device)	Old: VARYRDEV HELPCP New: RDEV HELPVARY	Old: `help cp varyrdev` New: `help vary real`
	SUBCHANNEL	Old: VARYSUBC HELPCP New: SUBCHANN HELPVARY	Old: `help cp varysubc` New: `help vary subchannel`
CMS CREATE [CREATE]	ALIAS	Old: ALIAS HELPCMS New: ALIAS HELPCREA	Old: `help cms alias` New: `help create alias`
	DIRECTORY	Old: DIRECTOR HELPCMS New: DIRECTOR HELPCREA	Old: `help cms directory` New: `help create directory`
	FILE	Old: FILE HELPCMS New: FILE HELPCREA	Old: `help cms file` New: `help create file`
	LOCK	Old: LOCK HELPCMS New: LOCK HELPCREA	Old: `help cms lock` New: `help create lock`
	NAMEDEF	Old: NAMEDEF HELPCMS New: NAMEDEF HELPCREA	Old: `help cms namedef` New: `help create namedef`
CMS DELETE [DELETE[1]]	LOCK	Old: DLOCK HELPCMS New: LOCK HELPDELE	Old: `help cms dlock` New: `help delete lock`
	NAMEDEF	Old: DNAMEDEF HELPCMS New: NAMEDEF HELPDELE	Old: `help cms dnamedef` New: `help delete namedef`
CMS SEGMENT [SEGMENT]	ASSIGN	Old: SASSIGN HELPCMS New: ASSIGN HELPSEGM	Old: `help cms sassign` New: `help segment assign`
	LOAD	Old: SLOAD HELPCMS New: LOAD HELPSEGM	Old: `help cms sload` New: `help segment load`
	PURGE	Old: SPURGE HELPCMS New: PURGE HELPSEGM	Old: `help cms spurge` New: `help segment purge`
	RELEASE	Old: SRELEASE HELPCMS New: RELEASE HELPSEGM	Old: `help cms srelease` New: `help segment release`
	RESERVE	Old: SRESERVE HELPCMS New: RESERVE HELPSEGM	Old: `help cms sreserve` New: `help segment reserve`

[1] The CMS DELETE subcommands have been added to the existing DELETE component (formerly CP only). Help files for CP DELETE subcommands and CMS DELETE subcommands can be displayed by specifying the following command:

`help delete` *`command_name`*

The following lists identify all of the help files that have been added or updated for this support.

- New CP help files

- For the ASSOCIATE subcommands:
  - ASSOCIAT HELPMENU
  - ASSOCIAT HELPABBR
  - EXIT HELPASSO
  - MESSAGES HELPASSO
- For the DEACTIVE subcommands:
  - DEACTIVE HELPMENU
  - DEACTIVE HELPABBR
  - CONV HELPDEAC
  - ISLINK HELPDEAC
- For the DRAIN subcommands:
  - DRAIN HELPMENU
  - DRAIN HELPABBR
  - DISK HELPDRAI
  - UR HELPDRAI
- For the FLASHCOPY subcommands:
  - FLASHCOP HELPMENU
  - FLASHCOP HELPABBR
  - BACKGNDC HELPFLAS
  - ESTABLIS HELPFLAS
  - RESYNC HELPFLAS
  - TGTWRITE HELPFLAS
  - WITHDRAW HELPFLAS
- For the FREE subcommands:
  - FREE HELPMENU
  - FREE HELPABBR
  - LOGON HELPFREE
  - PRINTER HELPFREE
- For the GIVE subcommands:
  - GIVE HELPMENU
  - GIVE HELPABBR
  - RDEV HELPGIVE
  - VDEV HELPGIVE
- For the HOLD subcommands:
  - HOLD HELPMENU
  - HOLD HELPABBR
  - LOGON HELPHOLD
  - PRINTER HELPHOLD
- For the MESSAGE commands:
  - MESSAGE HELPMENU
- For the START subcommands:
  - START HELPMENU
  - START HELPABBR
  - DISK HELPSTAR
  - UR HELPSTAR
- For the VARY subcommands:
  - VARY HELPMENU
  - VARY HELPABBR
  - CHPID HELPVARY
  - CORE HELPVARY
  - PATH HELPVARY
  - PCIFUNCT HELPVARY
  - PROCESSO HELPVARY
  - RDEV HELPVARY

- SUBCHANN HELPVARY

- Updated CP help files
  - ASSOCIAT HELPCP
  - CP HELPMENU
  - CP HELPABBR
  - DEACTIVE HELPCP
  - DELETE HELPCP
  - DELETE HELPMENU
  - DELETE HELPABBR
  - DRAIN HELPCP
  - FLASHCOP HELPCP
  - FREE HELPCP
  - GIVE HELPCP
  - HOLD HELPCP
  - MESSAGE HELPCP
  - START HELPCP
  - VARY HELPCP
- Updated (deprecated) CP help files
  - ASSOEX HELPCP
  - ASSOMSGS HELPCP
  - CONV HELPCP
  - DRAINDIS HELPCP
  - DRAINUR HELPCP
  - FLASHBND HELPCP
  - FLASHEST HELPCP
  - FLASHRES HELPCP
  - FLASHTGT HELPCP
  - FLASHWIT HELPCP
  - FREELOGN HELPCP
  - FREEPRIN HELPCP
  - GIVRDEV HELPCP
  - GIVVDEV HELPCP
  - HOLDLOGN HELPCP
  - HOLDPRT HELPCP
  - ISLINK HELPCP
  - STARTDIS HELPCP
  - STARTUR HELPCP
  - VARYCHPI HELPCP
  - VARYCORE HELPCP
  - VARYPATH HELPCP
  - VARYPCIF HELPCP
  - VARYPROC HELPCP
  - VARYRDEV HELPCP
  - VARYSUBC HELPCP
- New CMS help files
  - For the CREATE subcommands:
    - CREATE HELPMENU
    - CREATE HELPABBR
    - ALIAS HELPCREA
    - DIRECTOR HELPCREA
    - FILE HELPCREA
    - LOCK HELPCREA
    - NAMEDEF HELPCREA
  - For the DELETE subcommands:
    - LOCK HELPDELE

    - NAMEDEF HELPDELE
  – For the SEGMENT subcommands:
    - SEGMENT HELPMENU
    - SEGMENT HELPABBR
    - ASSIGN HELPSEGM
    - LOAD HELPSEGM
    - PURGE HELPSEGM
    - RELEASE HELPSEGM
    - RESERVE HELPSEGM
- Updated CMS help files
  – CMS HELPMENU
  – CREATE HELPCMS
  – DELETE HELPCMS
  – HELP HELPCMS
  – MENUS HELPMENU
  – SEGMENT HELPCMS
- Updated (deprecated) CMS help files
  – ALIAS HELPCMS
  – DIRECTOR HELPCMS
  – DLOCK HELPCMS
  – DNAMEDEF HELPCMS
  – FILE HELPCMS
  – LOCK HELPCMS
  – NAMEDEF HELPCMS
  – SASSIGN HELPCMS
  – SLOAD HELPCMS
  – SPURGE HELPCMS
  – SRELEASE HELPCMS
  – SRESERVE HELPCMS

Application Development and Deployment

These topics describe changes that can affect how you develop and deploy applications on z/VM.

[V2.2] REXX Sockets

The REXX Sockets API allows you to write socket applications in REXX for the TCP/IP environment. This allow you to use REXX to implement and test TCP/IP applications. A new REXX/VM external function, SOCKET, is provided. The SOCKET function uses the TCP/IP IUCV API to access the TCP/IP internet socket interface. The REXX socket functions are similar to socket calls in the C programming language. For more information, see *z/VM: REXX/VM Reference*.

[V2.2] Querying the CMS Level

The CMSPROG field of NUCON, returned in register 1 by the CMS command QUERY CMSLEVEL when the command is invoked internally by a program, has been frozen at X'0F', the value for CMS level 12 (VM/ESA V2.1). The CMSLEVEL assembler macro will not map CMS levels beyond CMS level 12. To determine the CMS level, use the DMSQEFL CSL routine or the DMSQEFL assembler macro.

[V2.2] Pseudo Timer Extended

DIAGNOSE code X'270' (Pseudo Time Extended) can be used to replace DIAGNOSE code X'0C' (Pseudo Timer). DIAGNOSE code X'0C' returns the time

only in SHORTDATE format. DIAGNOSE code X'270' returns the same information as DIAGNOSE code X'0C' plus two additional fields containing the time in FULLDATE format and ISODATE format.

[V2.2] Calling CSL Routines That Have Dates as Output

The following CSL routines, which provide dates in their output, have been enhanced with a new parameter for specifying the format in which the dates are to be returned:

- DMSEXIDI
- DMSEXIFI
- DMSOPBLK
- DMSGETDA
- DMSGETDF
- DMSGETDS
- DMSGETDX

The date formats are:

Keyword
Format

SHORTDATE
yy/mm/dd

FULLDATE
yyyy/mm/dd

ISODATE
yyyy-mm-dd

If no keyword is specified, the default is SHORTDATE, which is the format compatible with prior VM releases. The output field in which the date is returned is either 8 or 10 characters in length, depending on the format requested.

In REXX, the date field returned is always 10 characters. If SHORTDATE is specified or allowed to default, the 8-character date is padded on the right with two blanks. An incompatibility is that old REXX programs which now get the date returned as an 8-character string will start getting the date returned as a 10-character string.

If you call one of these routines from a higher level language, the date field is not padded. For SHORTDATE, an 8-character field is returned. For FULLDATE or ISODATE, a 10-character field is returned. Therefore, if you specify FULLDATE or ISODATE, you must be sure to also specify a 10-character output field. Otherwise, you could get storage overlays.

[V2.2] OS Simulation Support for Tape Library Dataserver

CMS OS simulation has been enhanced to call DFSMS/VM Removable Media Services (RMS) CSL routines to mount and demount tapes on Tape Library Dataserver machines. The new LIBSRV option on the CMS FILEDEF command is used to indicate that a tape is under the control of a Tape Library Dataserver. OS simulation can also determine that a tape is under Dataserver control if the user has issued the DFSMS/VM MOUNT command to premount the tape before the CMS tape processing function is invoked.

When a tape is under Dataserver control:

- OS simulation calls the RMS FSMRMDMT (Demount) and FSMRMMNT (Mount) routines to get subsequent multivolume tapes mounted for the user through the native DMSTVS mounting service and the CMS native rewind and unload tape processing functions.
- The RUN (rewind and unload) function of the CMS TAPE and VMFPLC2 commands and the TAPECTL macro calls the RMS FSMRMDMT routine to demount the tape.

See the following documents for more information:

- *z/VM: CMS Application Development Guide for Assembler*
- *z/VM: CMS Commands and Utilities Reference*
- *z/VM: DFSMS/VM Removable Media Services*

[V2.2] GCS Pathlength Reduced

GCS processor usage has been reduced, which improves the overall performance of GCS. This can also improve the performance of applications and program products that use GCS.

[V2.2] Full MP Capability for VMCF

The Virtual Machine Communications Facility (VMCF) has been changed so that it no longer has to run on the master processor. This can improve performance and capacity for systems that run applications (such as TCP/IP, OV/VM, and RACF) that use VMCF.

[V2.3] Language Environment

Language Environment is now included with VM. Language Environment provides a common run-time environment for programs generated with IBM C for VM/ESA and other high-level languages.

[V2.3] Dynamic Link Libraries

Dynamic link library (DLL) support is available for applications running under VM and Language Environment.

A DLL is a collection of one or more functions or variables gathered in a load module and executable or accessible from a separate application load module. The term derives from the fact that the connection or link between the application that uses the DLL and the DLL functions or variables is made dynamically while the application is executing rather than statically when the application is built.

For more information, see:

- *z/VM: CMS Application Development Guide*
- *z/VM: OpenExtensions Callable Services Reference*
- *C for VM/ESA: Library Reference*, SC23–3908

[V2.3] Non-Relocatable Modules Loaded between 16-20 MB

Non-relocatable modules that are loaded in the virtual machine between 16-20 MB must be regenerated. The CMS nucleus now spans from 15-20 MB, which will prevent these modules from being loaded. You must either regenerate the non-relocatable modules at a different storage location or generate the modules as relocatable modules using the RLDSAVE option on the CMS LOAD command. See *z/VM: CMS Commands and Utilities Reference* for information on generating modules using the LOAD and GENMOD commands.

[V2.3] CMS Migration Utility Feature No Longer Available

Beginning with VM/ESA V2.1, which became generally available in December 1995, CMS no longer runs in 370-mode virtual machines. To simplify the migration of older applications and to facilitate the running of 370-mode-only CMS applications in non-370-mode virtual machines (that is, in XA or XC mode) IBM provides the 370 Accommodation Facility. This function originally shipped with VM/ESA V1.2.1 (available July 1993) and has since been enhanced as the result of customer experience and input. 370 Accommodation handles the vast majority of inconsistencies between 370 mode and either XA or XC mode, eliminating the need to change these applications to exploit the new architectures.

IBM also provided the 370-capable level of CMS which shipped with VM/ESA V1.2.2 (known as CMS 11) as a no-charge feature for customers with specialized applications which 370 Accommodation could not support. Few VM customers have needed this CMS Migration Utility Feature and IBM has not updated it. Most importantly, the CMS Migration Utility Feature has not been made Year 2000 Ready. IBM recommends that customers who are using the CMS Migration Utility should modify their applications to run on a Year 2000-Ready level of CMS (VM/ESA V2.2 or later) and stop using the CMS Migration Utility as an application environment. Customers should not rely upon the CMS Migration Utility as a production environment. Also, IBM has announced the discontinuance of service for VM/ESA V1.2.2 (including CMS 11) and the CMS Migration Utility Feature as of April 30, 1999.

Consistent with IBM's policy of delivering only Year 2000-Ready products to customers in 1998 and beyond, the CMS Migration Utility Feature will no longer be available as part of VM. Customers who find that their applications cannot run under the current level of CMS should modify their applications if possible, or use the 370 Accommodation function. If problems persist, customers should call IBM Service for assistance.

[V2.3] POSIX Support Enhancements

Enhancements to VM POSIX support include additional C library routines, new callable services and mapping macros, new OPENVM commands for fork processing, and new shell commands for data compression.

Additional C/VM Library Routines

Over 150 C library routines have been added to the C/VM™ run-time library. For more information about these routines, see *C for VM/ESA: Library Reference*, SC23–3908.

New Callable Services and Macros

The following callable services have been added:

Service
: **Purpose**

DLL_delete (BPX1DEL)
: Delete a previously-loaded program from storage

DLL_load (BPX1LOD)
: Load a program into storage

fork (BPX1FRK)
: Create a new process

msgctl (BPX1QCT)
: Control message queues

msgget (BPX1QGT)
Create or find a message queue

msgrcv (BPX1QRC)
Receive a message from a message queue

msgsnd (BPX1QSN)
Send a message to a message queue

semctl (BPX1SCT)
Control semaphores

semget (BPX1SGT)
Create or find a set of semaphores

semop (BPX1SOP)
Perform semaphore operations atomically

shmat (BPX1MAT)
Attach a shared memory segment

shmctl (BPX1MCT)
Control shared memory segments

shmdt (BPX1MDT)
Detach a shared memory segment

shmget (BPX1MGT)
Create or find a shared memory segment

wait-extension (BPX1WTE)
Obtain status information about child processes

w_getipc (BPX1GET)
Query interprocess communications

The following mapping macros have been added:

Macro	Purpose
BPXYIPCP	Map interprocess communications permissions
BPXYIPCQ	Map the data structure used by w_getipc (BPX1GET)
BPXYMSG	Map interprocess communications message queues
BPXYSEM	Map interprocess communications semaphores
BPXYSHM	Map interprocess communications shared memory segments
BPXYSINF	Map the Siginfo_t structure used by wait-extension (BPX1WTE)

For additional information about these routines and macros, see *z/VM: OpenExtensions Callable Services Reference*.

OPENVM Commands for fork (BPX1FRK) Processing

Because the VM implementation of the fork (BPX1FRK) service does not meet all POSIX.1 requirements, you must explicitly turn fork (BPX1FRK) processing ON

with the CMS OPENVM SET FORK command before running your program. You can use the OPENVM QUERY FORK command to determine the current setting.

For information about restrictions in the VM implementation of fork (BPX1FRK), see *z/VM: OpenExtensions Callable Services Reference*. For more information about OPENVM commands, see *z/VM: OpenExtensions Commands Reference*.

Shell Commands for Data Compression

The following shell commands have been added to the Shell and Utilities:

- **compress**
- **uncompress**
- **zcat**

The **compress** and **uncompress** commands use Lempel-Ziv compression techniques to compress and uncompress data in files or from the standard input. The **zcat** command calls **uncompress** to uncompress data from one or more files or from the standard input and writes it to the standard output.

For more information about these commands, see *z/VM: OpenExtensions Commands Reference*.

[V2.4] OS Simulation Enhancements

CMS OS Simulation support has been enhanced as follows:

- The CMS FILEDEF command has been changed to allow record length (LRECL) definitions up to 65535 bytes for variable spanned OS records and non-OS CMS files, and block size (BLOCK or BLKSIZE) definitions up to 65535 bytes for non-OS CMS files.
- The CMS LKED command has been changed to allow larger default work area sizes (SIZE option): '400K' for *value1* and '100K' for *value2*.
- The CMS MOVEFILE command has been changed to allow the processing of QSAM variable spanned records up to 65535 bytes in length under the Extended Logical Record Interface (XLRI). MOVEFILE can also process non-OS CMS files up to 65535 bytes.
- The CMS QUERY FILEDEF command has a new optional operand, ATTRIBUT, which allows you to display the RECFM, LRECL, and block size attributes associated with the current FILEDEF.
- OS Simulation DCB macro processing has been changed to allow it to describe and pass both LRI and XLRI conventions for QSAM variable spanned long records, up to 65535 bytes in length, for subsequent OPEN, CLOSE, GET, or PUT processing. DCB can now also describe non-OS CMS files up to 65535 bytes in length.
- The OS Simulation OPEN, CLOSE, GET, and PUT macros, and the SVC 19 and SVC 20 supervisor calls, have been changed to allow the processing of OS formatted variable spanned QSAM records and non-OS CMS files up to 65535 bytes in length.
- The OS Simulation tape processing routines, such as DMSTVS, will issue a new message, DMS2139I, if SENSE data from a tape mount indicates that the mounted tape cartridge may be incorrect for the tape device in use. These tape processing routines may be invoked by MOVEFILE processing and by the following OS Simulation macros: OPEN, CLOSE, GET, PUT, READ, WRITE, and FEOV.

For more information, see:

- *z/VM: CMS Application Development Guide for Assembler*

- *z/VM: CMS Commands and Utilities Reference*

[V2.4] POSIX Support Enhancements

New realpath (BPX1RPH) Callable Service

The realpath (BPX1RPH) service finds the absolute path name for a specified relative path name. Any dot (.) or dot dot (..) components, symbolic links, or mount external links included in the relative path name input are resolved in the absolute path name output.

New setopen (BPX1VM6) Callable Service

The setopen (BPX1VM6) service sets certain flags specific to the platform without creating a new POSIX process in the virtual machine. The function codes used in the setopen service are mapped by the BPXYVM6 macro.

[V3.1] POSIX Support Now Called OpenExtensions

The support for POSIX and XPG standards in z/VM is now called OpenExtensions. Support services and the Shell and Utilities are now included in CMS.

[V3.1] Open Files Limit Increased

Under OpenExtensions, the limit on open files has been increased from 1024 to 65536.

[V3.1] BFS Root Not Case Sensitive

The fully qualified VM byte file system root (**/../VMBFS:***filepoolid***:***filespaceid***/**) is no longer case sensitive.

[V3.1] Binder/Loader Support

CMS hosts the OS/390 DFSMS Program Management binder and provides a native implementation of the corresponding program object loader functions. Included with this support is a new CMS command, BIND, to invoke the services of the binder to bind and store an executable file containing a program object on DASD (minidisk, SFS directory, or BFS directory). An executable file produced by the binder may be used on a CMS system containing the program object loader in exactly the same way as a conventional MODULE file generated by the CMS GENMOD command.

Also included in this support is an enhancement to the **c89** command to use the new BIND command in place of the combination of the Language Environment prelinker and the old LOAD, INCLUDE, and GENMOD commands. In addition, the complete binder API is available to application programs that need to invoke the services of the binder dynamically.

CMS support of the binder and loader allows ready transportability of executable files between z/VM and OS/390 or z/OS systems and provides the infrastructure to permit CMS to exploit new and emerging programming language technology being developed for z/OS.

For more information, see *z/VM: Program Management Binder for CMS*.

[V3.1] CMS OS Simulation Enhancement

The new SET TAPENEVR command allows you to control CMS OS Simulation tape label date checking for 'Unexpired Files'. The new QUERY TAPENEVR command allows you to display the current setting.

For more information, see:

- *z/VM: CMS Application Development Guide for Assembler*
- *z/VM: CMS Commands and Utilities Reference*

[V3.1] Tape FICON and RAS Support

Extended function in DIAGNOSE code X'210' allows application programs to find out the underlying real hardware that is associated with tape devices operating in emulation mode. For more information, see *z/VM: CP Programming Services*.

[V4.2] C Socket Application Programming Interface

A new z/VM C sockets library within Language Environment provides C socket support for new or existing applications that use Language Environment services. This support includes:

- New C socket-related function calls. These calls are documented in *z/VM: Language Environment 1.8 C Run-Time Library Reference*, SC24–6038. (This document replaces *C for VM/ESA: Library Reference*.)
- New OpenExtensions callable services (BPX routines), which correspond to the basic z/VM C socket calls. These services are documented in *z/VM: OpenExtensions Callable Services Reference*.
- A version of the TCP/IP Remote Procedure Call (RPC) function library, VMRPC, to be used in conjuction with the z/VM C sockets. Information about RPC is in *z/VM: TCP/IP Programmer's Reference*.

The z/VM C sockets are intended as replacements for the existing TCP/IP C sockets and OE sockets. Although those socket libraries are still supported for compatibility, the z/VM C socket API is preferred. Existing applications may need to recompile to use the new z/VM C socket functions, but no source code changes are required. For more information about using the z/VM C socket API, see *z/VM: TCP/IP Programmer's Reference*.

[V4.3] Increased 64-bit Addressing Support for DIAGNOSE Codes

The following DIAGNOSE codes now support the 64-bit addressing mode:

- DIAGNOSE code X'08' – Virtual Console Function
- DIAGNOSE code X'44' – Voluntary Time Slice End
- DIAGNOSE code X'4C' – Generate Accounting Records
- DIAGNOSE code X'DC' – Control Application Monitor Record Collection
- DIAGNOSE code X'258' – Page-Reference Services

For more information, see *z/VM: CP Programming Services*.

[V4.4] C/C++ Compiler Support

IBM IBM C/C++ for z/VM, Version 1 (5654-A22) Release 1.0, is a z/VM-enabled version of the C/C++ compiler for version 1 release 2.0 of z/OS. This new C/C++ compiler, which runs on CMS, allows C/C++ programs to be compiled and executed on CMS and creates portability between z/VM and z/OS C/C++ programs. The new OpenExtensions command **cxx** invokes the C/C++ for z/VM

compiler. C/C++ source files can be read from a CMS minidisk, the SFS, or the byte file system (BFS), and output can be written to any of these file systems. C/C++ for z/VM runs only on z/VM V4.4 and can be licensed only for standard processor engines. However, applications compiled with this C/C++ compiler will run on standard processor engines or Integrated Facility for Linux (IFL) processor features. For more information, see:

- *z/VM: OpenExtensions Commands Reference*
- *XL C/C++ for z/VM: User's Guide*, SC09-7625

[V4.4] Language Environment Upgrade

Language Environment for z/VM has been upgraded to provide C/C++, COBOL, and PL/I run-time libraries at the same level shipped with version 1 release 4.0 of z/OS. In addition, this new level of Language Environment has been integrated into z/VM as a component.

Attention: Do not migrate Language Environment 1.8 or earlier to z/VM V4.4. The only level of Language Environment supported on z/VM V4.4 is the new Language Environment component. The files for the z/VM Language Environment component are installed on the MAINT 19E disk. Depending on how you plan to migrate other files from your old 19E disk to the new system, you may need to remove the old Language Environment files. See "Removing the Old Level of Language Environment" on page 7.

Some of the new and enhanced features of the new level of Language Environment are:

XPLINK
Reduces program size and improves performance for small programs that frequently call each other.

Large File Support
Improves the porting capabilities of C/C++ applications accessing BFS and NFS files larger than 2 GB.

For more information, see *z/VM: Language Environment User's Guide*.

[V4.4] Systems Management APIs

The systems management APIs provide a basic set of functions that can be called by applications to allocate and manage resources for guests running in z/VM virtual machines (virtual images). Although these APIs are primarily intended for managing Linux virtual images, they can be used for managing any type of virtual image running in a z/VM virtual machine. The APIs are designed so that applications can be written by customers or solution providers to help administrators, especially those who lack in-depth VM knowledge, manage large numbers of virtual images running in a single z/VM system.

Functions supported by the systems management APIs allow you to:

- Create and delete virtual images
- Add and remove resources in an image's static (directory) or dynamic (active) configuration
- Activate and deactivate virtual images (individual or lists)
- Manage connectivity between virtual images, including exploitation of the new virtual switch support

- Query the status of "long-running" asynchronous operations initiated through the APIs
- Create, delete, replace, and query shared storage segments

The systems management APIs require a directory manager. The IBM Directory Maintenance Facility (DirMaint) optional feature of z/VM has been enhanced with support for these APIs through a service upgrade. See *z/VM: General Information*.

Client applications invoke the systems management APIs through a set of IBM-supplied Remote Procedure Calls (RPCs) to a z/VM server machine (either remotely or from within the z/VM system). The server then calls a set of CMS Callable Services Library (CSL) routines, which perform the requested functions. The CSL routines are replaceable to allow customized solutions for the systems management API functions.

[V5.1] Java and NetRexx Support Removed

Java and NetRexx support has been removed from z/VM. If you want to use Java in a z/VM environment, you should consider using Java for Linux running in a Linux guest.

[V5.1] Enhanced Systems Management APIs

All enhancements to the systems management APIs in z/VM V5.1 have been implemented using Version 2 (V2) of the RPC server. Previous APIs implemented using V1 of the RPC server can also function with the new V2 server. A new server security identification procedure is provided for special authorized clients, which can remove the need to log on with a password, simplifying the logon process. In addition to some usability enhancements, new functions include:

- DASD volume management for virtual images
- Virtual Machine Resource Manager (VMRM) configuration file management
- Query VMRM measurement data
- Query status of active images

Parameter lists for some V1 APIs have been changed to create enhanced APIs for V2. These enhancements are accessible only from a V2 client. The same RPC program number is used, but a new version number (2) is registered with the portmapper during server initialization. A client is required to log in as either a V1 or V2 client. During a client-server session, the client has access only to the functionality for the specified version.

The following V1 APIs have been enhanced:

- AUTHORIZATION_LIST_REMOVE has been enhanced to optionally remove multiple entries for a user ID from the VSMSERVE AUTHLIST file.
- LOGIN has been enhanced to recognize certain V2 clients as trusted users that are not required to supply a password when logging in to the server.
- SHARED_STORAGE_QUERY has been enhanced to optionally return information on all shared segments.
- VIRTUAL_NETWORK_LAN_CONNECT and VIRTUAL_NETWORK_VSWITCH_CONNECT have been enhanced with a new parameter to support Layer 2 LAN, which provides an alternative transport mechanism for guest LANs and virtual switches.
- VLANID parameter of VIRTUAL_NETWORK_VSWITCH_SET has been enhanced to support changes to virtual switch operation.

The IBM z/VM Directory Maintenance Facility (DirMaint) feature, Function Level 510, has been enhanced to support the new API functions.

[V5.1] Dynamic Virtual Machine Timeout

A new programming service is provided by an emulated DIAGNOSE instruction that helps enable a guest operating system to specify an action to be taken by CP when the guest becomes unresponsive. The DIAGNOSE instruction also allows each guest to define what "unresponsive" means for itself. A time interval and action are specified by the guest. If the guest fails to reissue the DIAGNOSE instruction within the specified time interval, CP performs the action. The guest may define the action as any set of CP commands, although it is envisioned that this function will be used to terminate the virtual machine in which the guest is running if it becomes unresponsive. This can help ensure that a malfunctioning guest would no longer have access to shared system resources and would relinquish system resources to other guests.

For more information, see *z/VM: CP Programming Services*.

[V5.1] Relocation of REXX Work Area for Applications

In previous releases, the REXX work area that CMS used to pass parameters to applications was always allocated from below the 16 MB line. Now this work area is allocated from above 16 MB if the application has identified itself as either AMODE 31 or AMODE ANY.

[V5.2] 64-Bit DIAGNOSE Code X'250' Support

Support for 64-bit addressing has been added to DIAGNOSE code X'250' to allow users to use storage above 2 GB for the DIAGNOSE code X'250' - Block I/O parameter list (BIOPL), DIAGNOSE code X'250' Block I/O entries, and I/O buffers. This allows users of DIAGNOSE code X'250' to be in 64-bit addressing mode.

The 31-bit interface and the existing users of the 31-bit interface will not be affected.

For more information, see *z/VM: CP Programming Services*.

[V5.2] CMS Binder Upgraded

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

[V5.2] Enhanced Systems Management APIs

All enhancements to the systems management APIs in z/VM V5.2 have been implemented using Version 3 (V3) of the RPC server. Previous APIs implemented using V1 or V2 of the RPC server can also function with the new V3 server. A new server security identification procedure is provided for special authorized clients, which can remove the need to log on with a password, simplifying the logon process. In addition to some usability enhancements, new functions include:

- V3 APIs to support the DirMaint DEFINESTAG, SETSTAG, and SCAN commands:
  - Directory_Manager_Local_Tag_Define
  - Directory_Manager_Local_Tag_Set
  - Directory_Manager_Local_Tag_Query

  - Directory_Manager_Search
- V3 APIs to allow creation and updating of the LOADDEV directory statement for a virtual image, as well as queries of a virtual image's LOADDEV settings:
  - Image_SCSI_Characteristics_Define
  - Image_SCSI_Characteristics_Query
- A V3 API, Image_Query_Activate_Time, to query the time when a virtual image was activated.

In addition to these new V3 APIs, new V2 APIs are added to support DirMaint's ability to manage subscriptions to directory updates, with the ability to add new subscriptions, delete old subscriptions, and query existing subscriptions:

- Asynchronous_Notify_Enable
- Asynchronous_Notify_Disable
- Asynchronous_Notify_Query

Parameter lists for some V1 and V2 APIs have been changed to create enhanced APIs for V3. These enhancements are accessible only from a V3 client. The same RPC program number is used, but a new version number (3) is registered with the portmapper during server initialization. A client is required to log in as either a V1, V2, or V3 client. During a client-server session, the client has access only to the functionality for the specified version.

Enhancements to existing APIs supporting Vswitch functions are as follows:

- *VLan* and *PortType* parameters are added to Virtual_Network_Vswitch_Connect and Virtual_Network_Vswitch_Set APIs. These parameters will support the corresponding operands on the DEFINE VSWITCH and SET VSWITCH commands.
- Capability is added to define persistent Vswitch definitions.

The RPC and CSL interfaces for the following APIs are updated to allow up to 10 digits in the block size parameters:

- Image_Disk_Create (*ImageDiskAllocationType* and *ImageDiskSize* parameters)
- Image_Disk_Copy (*ImageDiskAllocationType* parameter)
- Image_Volume_Space_Define (*Size* parameter)

An optional *DeviceType* parameter is added to the Image_Volume_Space_Define API to allow space (regions and groups) to be defined for an image volume before the volume exists on the system.

The DirMaint optional feature, function level 510, has been enhanced to support the new API functions.

[V5.2] LT MACRO Renamed

Because the IBM High Level Assembler, version 1 release 5, supports a hardware opcode called LT, that function is used before the CMS LT MACRO is called. Therefore, a new DMSLT MACRO has been created (identical to the LT MACRO), and references to the LT MACRO in CMS, GCS, and REXX/VM have been changed to use the DMSLT MACRO instead. If you currently use the LT MACRO, IBM recommends that you also change to using the DMSLT MACRO. Although the LT MACRO is still supplied in the DMSGPI MACLIB, it could be withdrawn in a future release.

[V5.2] SRPI Support Discontinued

Support for the Server-Requester Programming Interface (SRPI) has been discontinued. The following CMS functions are no longer supported:

- ADDENTRY macro
- ADDENTRY subcommand
- CMSSERV command
- CPRB macro
- CSMRETCD macro
- DELENTRY macro
- DELENTRY subcommand
- GETREQ subcommand
- SENDREQ macro
- SENDREQ subcommand
- SETREPLY subcommand

[V5.2] VSE/VSAM Support

IBM VSE/VSAM, V6.1 (5686-081), was withdrawn from marketing on September 30, 2005, and was removed from the z/VM V5.2 System Delivery Offering on January 1, 2006. Refer to Withdrawal Announcement 905-121 (US), dated June 7, 2005. Although z/VM CMS DOS/BAM will continue to provide its current level of function in support of DOS simulation and VSE/VSAM, customers should consider looking for alternatives to VSE/VSAM and moving their data off of VSAM packs on z/VM.

[V5.3] Sockets-Based Systems Management APIs

With z/VM V5.3, a new sockets-based server supports the z/VM virtual systems management API. The sockets-based server is multitasking-capable and supports both AF_INET and AF_IUCV socket requests.

In addition to the new server, enhancements provided in z/VM V5.3 include:

- New functions to:
  - Create, delete, and query the IPL statement in a virtual image's directory entry
  - Create and delete virtual switches and guest LANs
  - Obtain processor, memory, and device information for active virtual images
  - Check the validity of a given user ID/password (or passphrase) combination
- Enhancements to existing functions to:
  - Provide values of specific attributes for selected query functions, rather than return a buffer containing QUERY command output
  - Exploit the new Asynchronous CP Command function available in z/VM V5.3
  - Accept passphrases and forward them to the external security manager to be set, changed, or deleted
  - Provide a list of active virtual images

The new and enhanced API functions for z/VM V5.3 have been implemented using the new sockets-based server. Functions provided in earlier releases of z/VM can also be invoked through the new server. The socket interface and API functions are described in a new edition of *z/VM: Systems Management Application Programming*.

The sockets-based server replaces the Remote Procedure Call (RPC) server and CSL routines that were used to call the virtual systems management API in previous releases of z/VM. The RPC server is still available in z/VM V5.3, with all of the functions that were available in z/VM V5.2. However, the enhancements provided in z/VM V5.3 are not available through the RPC server, for which no future enhancements are planned. Documentation on the use of the API with the RPC server and CSL routines will not be updated and will not be included in the z/VM V5.3 bookshelf. IBM intends to remove the RPC server from a future z/VM release.

[V5.3] TCP/IP and CMS Dependency

Due to changes in CMS resolver code, there is a requirement that some TCP/IP modules such as FTP must be at the z/VM V5.3 level (TCP/IP FL530) when the IPLed CMS is z/VM V5.3 (CMS level 23). For the same reason, customer-built modules that include COMMTXT in the GLOBAL TXTLIB command in order to use routines such as SayIntAd must be rebuilt using the z/VM V5.3 level of COMMTXT. In addition, some TCP/IP modules require Language Environment to include APAR VM64055. (For information about Language Environment program requirements, see *z/VM: General Information*.)

The following message will appear if either a TCP/IP routine or Language Environment is back-level:

```
DMSZER2571E A resolver request failed due to missing LE support or incorrect
TCP/IP module levels.
```

To work around this problem, upgrade TCP/IP and Language Environment and rebuild programs that use COMMTXT, or IPL an earlier level of CMS.

[V5.4] Enhanced Systems Management APIs

In z/VM V5.4, communication between request servers and worker servers is improved, and many functions will now see gains in performance and multitasking capabilities. In addition, new APIs have been added to perform the following functions:

- Create, delete, replace, and query directory profiles.
- Manage and query LAN access lists (granting and revoking access to specific userids).
- Define, delete, and query virtual CPUs, both within an active virtual image and in a virtual image's directory entry. Also set a maximum number of virtual processors that can be defined in a virtual image's directory entry.

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

[V5.4] Program Management Binder Upgrade

The program management binder for CMS has been upgraded to the z/OS V1.9 level.

The binder supports autocall from OpenExtensions archive libraries. These archive libraries may contain members that are object files or members that are side files (of IMPORT control statements) or other files of control statements.

The following BIND command options have been added for this support:

- COMPRESS
- INFO
- MODMAP
- STRIPCL

- STRIPSEC

The LIBRARY control statement now provides a more robust search of autocall libraries during final autocall processing.

The following binder control statements have been updated to add the -IMMED suboption:
- CHANGE
- REPLACE

For more information, see *z/VM: Program Management Binder for CMS*.

[V5.4] Language Environment Upgrade

This upgrade brings the Language Environment runtime libraries to the same level shipped with z/OS 1.9, providing a base for the new IBM XL C/C++ for z/VM compiler as well as other service and general improvements. For more information, see:

- *z/VM: Language Environment User's Guide*
- *XL C/C++ for z/VM: Runtime Library Reference*, SC09-7624

[V5.4] XL C/C++ Compiler Support

IBM XL C/C++ for z/VM, V1.2, is a z/VM-enabled version of the z/OS V1.9 XL C/C++ compiler. This new level of the compiler for z/VM offers new and enhanced compiler features. For more information, see:

- *XL C/C++ for z/VM: User's Guide*, SC09-7625
- *z/VM: OpenExtensions Commands Reference*

[V5.4] DIAGNOSE Code X'2A8' – Network Diagnose

DIAGNOSE code X'2A8' provides an interface to establish, activate, send and receive layer 2 Ethernet data using a z/VM virtual NIC device coupled to an Ethernet VSWITCH. The Diagnose instruction manages the control and data devices required by a network connection on behalf of the guest. This allows a virtual machine to send and receive data without any knowledge of the underlying hardware architecture. For more information, see *z/VM: CP Programming Services*.

[V6.1] RPC-Based Systems Management APIs Not Supported

With z/VM V6.1, IBM has withdrawn support for the RPC-based systems management APIs. Only the sockets-based systems management APIs are supported. For information about the sockets-based APIs, see *z/VM: Systems Management Application Programming*.

[V6.1] LT MACRO Removed

The LT MACRO, which was replaced by the DMSLT MACRO in z/VM V5.2, has been removed from DMSGPI MACLIB.

[V6.1] Enhanced Systems Management APIs

The z/VM systems management APIs are enhanced to enable the IBM z Unified Resource Manager to manage the z/VM hypervisor and the virtual servers it supports. The enhancements provide new and updated APIs and function in the following functional areas:

- System configuration file management

- User directory management
- Disk management
- Network management
- Discontiguous saved segment (DCSS) management
- Guest facilities management
- Directory parsing
- API function level interrogation

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

[V6.1] Program Management Binder Upgrade

The Program Management Binder for CMS has been upgraded to the level shipped with z/OS V1.11 MVS Program Management. The following functions have changed:

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

For more information, see *z/VM: Program Management Binder for CMS*.

[V6.2] Enhanced Systems Management APIs

The z/VM systems management APIs are enhanced to support single system image (SSI) clusters and other new functionality in z/VM V6.2, and to support Unified Resource Manager enhancements provided with z196 Drivers D86 and D93. The z/VM systems management enhancements include new and updated APIs in the following functional areas:

- Single System Image Management
- Platform Performance Management
- Availability Management
- Energy Management
- Network Virtualization Management
- Hypervisor Virtualization Management
- Virtual Server Management
- Storage Virtualization Management
- Directory Management

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

[V6.2] Language Environment Upgrade

The z/VM Language Environment runtime libraries have been upgraded to z/OS V1.12 equivalency. See *z/VM: Language Environment User's Guide* for more information.

[V6.2] Program Management Binder Upgrade

The program management binder for CMS has been upgraded to the z/OS V1.12 level. The following functions have changed:

- The COMPAT option of the BIND command now includes a new suboption: ZOSV1R12.

- The RMODE option of the BIND command now includes several new suboptions.

For more information, see *z/VM: Program Management Binder for CMS*.

[V6.2] Storage Location X'C8' Now Used by CMS Initialization

CMS initialization is now using the STFL instruction, which stores at location X'C8' a list of bits detailing hardware facilities. If a user application assumes the space at X'C8' is available and attempts to use it, that application might fail now that CMS is storing data there.

[V6.2] XL C/C++ Compiler Upgrade

IBM XL C/C++ for z/VM, V1.3, is a z/VM-enabled version of the z/OS V1.12 XL C/C++ compiler. For more information, see *XL C/C++ for z/VM: User's Guide*, SC09-7625.

[V6.3] Enhanced Systems Management APIs

In z/VM V6.3, new APIs are added to perform the following functions:

- Obtain DASD read and write byte counts for SCSI EDEV and ECKD volumes owned by z/VM, and for which the control units have information.
- Obtain a list of the system devices assigned a device equivalency ID.
- Query the status of directory manager locks in effect for a specific virtual image.
- Query the links to an image's MDISK.
- Create, change, remove and obtain information about a network interface configuration for the z/VM TCP/IP stack.
- Query the status of directory manager locks in effect for a specific profile.
- Capture data to assist with identification and resolution of a problem with the SMAPI servers. (Note that both an API and a stand-alone EXEC are provided for this purpose.)
- Obtain information about a CP instance, including time, storage, system levels, IPL time, system generation time, language, CPU ID, and CPU capability information, and more.
- Obtain information about the z/VM paging space defined on the system.
- Gather hypervisor performance data, including available/used, processor number, total processor percentages, and optional detailed CPU information for all visible LPARs on the CEC, and query, set and stop the monitor rate and interval values.
- Query the status of an APAR, PTF, or RSU for a zVM component.
- Systematically end all system function.
- Obtain information about the z/VM spool space defined on the system.

Additional enhancements include:

- The LOHCOST Data Base, used to cache directory and meta data, is now available for use in all SMAPI installations.
- A new First Failure Data Capture tool (both an API version and a stand-alone version) are added for gathering information to assist in problem determination.

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

[V6.3] Program Management Binder Upgrade

The Program Management Binder for CMS has been upgraded to z/OS V1.13 equivalency. The following functions have changed:

- COMPAT option of the BIND command includes a new ZOSV1R13 parameter.
- STRIPSEC option of the BIND command includes a new PRIV parameter.
- Support for RLD conditional sequential resolution.

[V6.3] Unified Resource Manager Support Withdrawn

The IBM z Unified Resource Manager, first supported in z/VM V6.1, is no longer supported in z/VM V6.3. Therefore z/VM V6.3 cannot participate as a member of an ensemble.

[V6.3 APAR] Support for OpenStack Exploitation

With the PTF for APAR VM65535, z/VM provides support for OpenStack exploitation, which will enable the OpenStack API interface for other IBM deliverables.

Administrators who wish to use OpenStack to help manage their systems should see the new publication, *z/VM: Enabling z/VM for OpenStack (Support for OpenStack Icehouse Release)*, for more information.

[V6.3 APAR] Systems Management API Support for Multi-VSwitch Link Aggregation

With the PTF for APAR VM65670, the z/VM systems management API provides support for Multi-VSwitch Link Aggregation. The following APIs have been updated:

- Virtual_Network_Vswitch_Create_Extended
- Virtual_Network_Vswitch_Query_Extended
- Virtual_Network_Vswitch_Set_Extended

For more information about Multi-VSwitch Link Aggregation, see "[V6.3 APAR] Multi-VSwitch Link Aggregation Support" on page 155. For information about the updated APIs, see *z/VM: Systems Management Application Programming*.

[V6.4] z/VM Program Management Binder z/OS V2.2 Equivalency

The Program Management Binder for CMS is upgraded to a level equivalent to z/OS V2.2. 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 reliability, availability, and serviceability (RAS) items

The following functions have changed:

- COMPAT option of the BIND command includes a new ZOSV2R1 parameter.
- SYMTRACE option of the BIND command is new.
- ALIGNT control statement is new.

For more information, see *z/VM: Program Management Binder for CMS*.

[V6.4 APAR] VMUDQ Macro Identifies SUBCONFIGs That Contain MDISKs

With the PTF for APAR VM65877, the VMUDQ macro is enhanced to support identifying the IDs of SUBCONFIGs that contain MDISK definitions in an SSI-enabled or SSI-ready directory.

System Diagnosis

These topics describe changes that can affect how you diagnose problems on z/VM.

[V2.2] FST and ADT Macro Changes to Support 4-digit Years

File Status Table

When you access a disk or SFS directory, a file directory is stored in your virtual machine. The entries in the file directory for each CMS file are called the file status table (FST). The FST describes the attributes of the file. One of the attributes of a file is date/time of last update. This is currently stored in 6 bytes (*yy mm dd hh mm ss*), where each byte holds two decimal digits. This is the date and time that the accessed file was last updated.

In VM/ESA V2.2, support was added for 4-digit years by adding a century flag FSTCNTRY (X'08') in the FST flag byte (FSTFLAGS) for both FST forms (see table below). This is bit 4, which identifies the century in which the file was last written or updated. If bit 4 is off, the year is in the 1900s. If bit 4 is on, the year is in the 2000s. Therefore, the range of years supported is from 1900 to 2099.

If you are interested in the existing flag settings for the FSTFLAGS field, see the usage notes for the FSSTATE macro in *z/VM: CMS Macros and Functions Reference*.

If you are using the FST to retrieve the date and time of last update in your application, you can add support to use an appropriate CSL routine (for example, DMSGETDI, DMSEXIST, or DMSERP) or you can use the FSSTATE macro. For information about CSL routines, see *z/VM: CMS Callable Services Reference*. For information about FSSTATE, see *z/VM: CMS Macros and Functions Reference*.

Table 7. Base versus Extended FST Format. Note that not all field names are shown.

Hex Disp	Dec Disp	Size	Base	Ext	Field Name	Field Description
0	0	16	B	E	FSTDFNFT	File name File type
		8	B	E	FSTFNAME	File name
		8	B	E	FSTFTYPE	File type
10	16	2	B		FSTDATEW	Date (*mm yy*) last written
1F	31	1	B	E	FSTFLAGS	FST flag byte
		Bit 4	B	E	FSTCNTRY	X'08' Century last written (0 - 19*nn*, 1 - 20*nn*)
26	38	2	B		FSTYEARW	Year (*yy*) last written
36	54	6		E	FSTADATI	Alternate date/time (*yy mm dd hh mm ss*)
4E	78	1		E	FSTFB3	FST flag byte 3
		Bit 4		E	FSTCDOLR	X'08' Century date of last reference (0 - 19*nn*, 1 - 20*nn*)
54	84	3		E	FSTDOLR	Date of last reference

Table 7. Base versus Extended FST Format (continued). Note that not all field names are shown.

Hex Disp	Dec Disp	Size	Base	Ext	Field Name	Field Description
Note: FSTFB3, FSTCDOLR, and FSTDOLR apply only to an SFS FST.						

Also, for an SFS FST, a new flag FSTCDOLR (X'08') is added in the FSTFB3 flag byte 3. This is the flag that indicates the Century for Date of Last Reference (0 indicates the year is in the 1900s, 1 indicates the year is in the 2000s) that corresponds to FSTDOLR.

Active Disk Table - Disk Label

The ADTSECT maps information in the active disk table (ADT). It also contains information about the disk label. One of the fields contained in the disk label is ADTDCRED, which is the creation date and time of the minidisk. Following the ADTDCRED field is a new flag byte for the volume label called ADTFLGL, which contains the new ADTCNTRY flag. The ADTCNTRY flag is X'01', and corresponds to the ADTDCRED field. If the value of this flag is 0, it indicates that the creation year is in the 1900s. If the value is 1, it indicates the creation year is in the 2000s.

[V2.3] Viewing and Printing CP and CMS Control Blocks

CP and CMS control block documentation is now available at IBM: z/VM data areas, control blocks, and monitor records (www.ibm.com/vm/pubs/ctlblk.html).

You can view or print an entire control block or select general sections.

In addition, VM provides an unsupported tool, DACBGEN EXEC, that you can use against CP and CMS control blocks on your system to get the most current data. You can also use DACBGEN against user blocks if they follow the prescribed formula. You can tailor DACBGEN to your own environment.

[V3.1] Dump Viewing Facility No Longer Supports CP Dumps

The Dump Viewing Facility no longer supports CP dumps; it supports only VM (virtual machine) dumps. The VM Dump Tool is the dump viewing program to use for z/VM CP dumps.

[V3.1] VM Dump Tool

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

[V4.3] VM Dump Tool Enhancement

The structure of the VM Dump Tool has been changed to use 31-bit addressing instead of 24-bit addressing. While this did not directly affect any external interfaces, it should reduce storage requirements in the virtual machine below the 16 MB line.

[V4.4] VM Dump Tool Enhancements

The following support is added to the VM Dump Tool:

- A new EXTRACT function for easy access from a macro environment to selected information in the dump

- Support for a macro file type of VMDT.
- Improved support for non-CP dumps, including new subcommands/macros for migration from the Dump Viewing Facility environment (DVFSTACK, FINDSTRG, INIT, and READSTRG)
- You can now halt long-running commands without losing unsaved dump session information by using the existing HI IMMEDIATE command. This CMS command is now supported by the following VM Dump Tool subcommands: CHAIN, FRAMES, LOCATE, TRACE, and VMDBK.
- The new XEDITPRE option of the VM Dump Tool SET subcommand allows you to disable or change the XEDIT escape value (the function that indicates when a command should be forwarded directly to XEDIT). You can reset it to either a single character or a string. This option replaces the XEDIT subcommand. Also, the new XEDITPRE option of the VM Dump Tool QUERY subcommand allows you to query the current setting for the XEDIT escape value.
- Output from the DISPLAY subcommand, when issued from a macro, is now the same as DISPLAY subcommand output when the subcommand is issued from the command line. Any macro which depends on the output of DISPLAY will have to be modified to handle the new output.

For more information, see *z/VM: VM Dump Tool*.

[V5.1] VM Dump Tool Enhancements

The following support is added to the VM Dump Tool:

- The VM Dump Tool can be used to look at all dumps from Version 3, Release 1.0 through Version 5, Release 1.0.
- The MAP command and related infrastructure which creates a VMDTMAP file has been rewritten. For a CP dump the externals are not changed. The new code obtains the information for the VMDTMAP file from the CP symbol table if possible, which avoids the time required to scan storage. If the CP symbol table is not available (which is the case for older dumps), then the VM Dump Tool scans storage looking for CP module headers as it has in the past.
- New MAPA and MAPN options of the EXTRACT subcommand make it easier for a macro to find out the name or address of a module. EXTRACT MAPA allows a macro to find the name of a module or entry point from an address. EXTRACT MAPN allows a macro to find the address of a module or entry point name.
- The new sample VMDTNCPM macro helps you build a VMDTMAP file for non-CP dumps.
- The new SET DEBUG function provides new support to help you debug macros written for the VM Dump Tool. With SET DEBUG ON, non-zero return codes associated with a macro are reported to the virtual machine console. The new matching QUERY DEBUG function displays your DEBUG setting (ON or OFF).
- The preferred file type for a VM Dump Tool macro is now VMDT rather than XEDIT. In a future release all IBM-supplied VM Dump Tool macros will be renamed from their current file type of XEDIT to a file type of VMDT. Also, the default environment when a VM Dump Tool macro is entered is now VMDUMPTL.
- To be consistent with the SETVAR function, the EXTRACT DFIR/DFIZ and EXTRACT SYMPTOM functions now require a decimal length value rather than hexadecimal. If you have any customer-written VM Dump Tool macros which use these functions, you will have to change the macros to use a decimal length value. Also, the maximum size for EXTRACT DFIR has been increased to 36,864 bytes.

- The new EXTRACT DISPL function allows a macro to obtain the displacement of a specified field. Only selected fields of PFXPG and SYSCM are supported at this time.
- The INDQ is no longer supported. It now returns only a response of "complete."
- The CPEBK subcommand and the CALLERS macro now both take up to 17 characters of input. Both still have the restriction of allowing up to only eight significant digits in the address operand.
- The TRACE FROM and TRACE TO functions no longer check the frame type to be sure it is a trace frame. Because of this these functions can now be used when the frame table is not available. (However, these functions can cause unpredictable output when the page is not really part of the trace table.)
- The output from the RDEVBK, RSCH, VDEVBK and VSCH subcommands has been changed to produce the addresses of significant related control blocks instead of displaying the storage of the control blocks themselves.

For more information, see *z/VM: VM Dump Tool*.

[V5.2] 64-Bit Dump Support

The 64-Bit dump support changes z/VM's various dump producers to create dumps that include relevant storage above 2 GB if that storage is defined. Prior to this release, all dump producers (Stand-Alone Dump, Hard Abend Dump, Soft Abend Dump, and VMDUMP) dumped storage only below 2 GB.

There are now four different dump formats:

- Little Storage Dump Format

  This format is produced by Stand-Alone Dump and the CP VMDUMP command when the processor or virtual machine is in ESA/390 mode. This dump format is limited to 2 GB of storage.
- Big Storage Dump Format

  This format is new and also shows storage above 2 GB, if it exists. This format is produced by Stand-Alone Dump and the CP VMDUMP command when the processor or virtual machine is in z/Architecture mode.
- CP Hard Abend Dump Format

  This format is new and is created when CP takes a Hard Abend or when the CP SNAPDUMP command is used. Dumps in this format contain only CP pages. The ALL operand has been removed from the SET DUMP command.
- CP Soft Abend Dump Format

  This format is new and is created for Soft Abend dumps. At the time of a Soft Abend dump, CP is still running. This format allows CP to dump out a few pages that will help debug the problem and then continue running.

DUMPLOAD has been updated to process all four dump formats.

Support for older tape devices has been dropped. For Stand-Alone Dump and Hard Abend Dump, only dumps to 3480 or 3590 tapes are supported. For Stand-Alone Dump, IPL off a CKD device is no longer supported. IPL off ECKD devices is still supported.

Support for writing a formatted Hard Abend dump to a printer has been eliminated.

The following CP functions have been updated:

- QUERY ABEND command

- QUERY DUMP command
- QUERY VMDUMP command
- SET ABEND command
- SET DUMP command
- SET UNDERSCORE command
- SNAPDUMP command
- VMDUMP command
- DIAGNOSE code X'94'

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: CP Planning and Administration*
- *z/VM: CP Programming Services*

[V5.2] High Level Assembler No Longer Required for Stand-Alone Dump Utility

The IBM High Level Assembler is no longer required to build the Stand-Alone Dump Utility. This has resulted in incompatible changes to a number of macros. Customers who have saved their HCPSDC ASSEMBLE source from previous releases must recreate it on z/VM V5.2 when prompted by the HCPSADMP EXEC.

[V5.2] VM Dump Tool Enhancements

The following support is added to the VM Dump Tool:

- Some new subcommands to establish the default addressing mode for the dump session.
- Many subcommands will now take a prefix of L/R/U on an address to indicate how an address should be interpreted.
- A number of subcommands now require an address to be in a certain form (that is, Logical or Real).
- Most existing syntax will continue to work without change.
- User-written macros may need to be updated to indicate that address references are real or logical.
- The new VMDTSET and VMDTQRY subcommands are introduced for SET and QUERY functions unique to the VM Dump Tool. The minimum abbreviations are DTS and DTQ.
  - Existing SET/QUERY functions will continue to work.
  - Existing SET/QUERY operands will work with the new vmDTSet/vmDTQry subcommands.
  - Some existing VM Dump Tool subcommands are also supported by VMDTSET and VMDTQRY. The existing commands continue to be supported for compatibility.
  - Any future parameters will be added to vmDTSet and vmDTQry, and not to SET and QUERY.
- The preferred file type for VM Dump Tool macros is changed to VMDT (it was XEDIT). Macros shipped with the VM Dump Tool will have a file type of VMDT, and must be invoked with ADDRESS VMDUMPTL to be invoked from a macro. Existing customer XEDIT macros can be invoked with ADDRESS XEDIT and will work without change.
- Many of the macro and subcommand examples have been updated.

For more information, see *z/VM: VM Dump Tool*.

[V5.3] VM Dump Tool Enhancements

The following support is added to the VM Dump Tool:

- **TRACE Subcommand**
  - All trace format output is now table driven.
  - The output line is no longer limited to 80 characters.
  - The user can add formatting information for new entries or override the format of any existing trace entry.
  - New options are provided to change the display of the standard 'header' part of each trace entry.
  - There are many new ways to display the time on trace entries:
    - both CPU and time
    - time in TOD format
    - delta time from last entry
    - time in microseconds
    - and others
  - TRACE by userid finds *all* VMDBKs for that user (was primary only).
  - TERM and TYPE handling has been changed to allow the user to change them or add new ones.
  - A supported trace macro interface is provided.
  - TRACE output is now more consistent, particularly VMDBK & USER information.
  - Old TRACE is still available at this time.
- **Appendix G. Using the Trace Format Definition Table**
- **Other Changes**
  - A number of related changes have also been made to the VMDTSET and VMDTQRY commands.
  - The preferred VM Dump Tool profile is now PROFILE VMDT. Older names of VMDTPROF VMDT and VMDTPROF XEDIT will still be handled, in that order, if PROFILE VMDT is not found.

For more information, see *z/VM: VM Dump Tool*.

[V5.4] VM Dump Tool Enhancements

The following support is added to the VM Dump Tool:

- The mnemonic table (INSTR subcommand), abend table (DESCRIBE subcommand) and trace definition table (TRACE subcommand) are now on a per-release basis instead of one of each table across all releases. The name of the table being used is available from extensions to the VMDTQRY subcommand. The VMDTSET subcommand can be used to set or override the name of any of these tables. VM Dump Tool initialization has been updated to recognize the release of the dump and default to the appropriate data files for the dump being processed.
- A number of new mnemonics and formats have been added to the INSTR subcommand.
- The HCQGDSPL macro is available to return the displacement of a field from the BLOCK data. A number of VM Dump Tool macros were updated to use this new function.
- A new VMDTSET SEPARATE subcommand has been added to add a blank line after every subcommand entered from the command line. A matching VMDTQRY option is also included.

- The SXSTE macro is now supported instead of being a sample/example.
- *Appendix F. Sample Programs Associated with the VM Dump Tool* contains new samples/examples for other VM Dump Tool macros and related pipes stages.

For more information, see *z/VM: VM Dump Tool*.

[V5.4] Multiple File Dump Support

z/VM V5.4 provides additional dump support:

- The new DUMPLD2 utility splits a single dump into multiple CMS files and can be used as an alternative to the DUMPLOAD command. Dumps that have been split into multiple smaller files can be stored across multiple minidisks or SFS directories instead of requiring a single larger disk. Segmenting a large dump into multiple files allows for easier handling of the dump by support teams. DUMPLD2 will handle all dump types currently supported by DUMPLOAD, except for multiple address space dumps. Large Storage Format, Small Storage Format, Hard Abend, Soft Abend, Standalone, and VMDUMP dumps will be handled.
- Support has been added for invoking the VM Dump Tool against one dump stored in multiple CMS files. The VM Dump Tool ascertains if the dump is a single file or multiple file by the file type – multiple files must have a file type of the form 'MDMPxxxx', whereas a single file dump can have any file type (the default file type for a single file dump is 'DUMP0001'). All multiple files must be available in order to view the dump, but they can span multiple file modes. One or more of the file modes may also be an SFS directory.
- The VMDTQRY subcommand has been updated with some additional output.

For more information, see:

- *z/VM: CP Commands and Utilities Reference*
- *z/VM: VM Dump Tool*

[V6.2] VM Dump Tool Enhancements

The following support is added to the VM Dump Tool:

- The TRACE subcommand will now support Function Related trace tables in addition to the standard CP trace table.
  - Related changes have also been made to the VMDTQRY command.
- A new macro, FRT2MAIN, has been added for Function Related trace tables.

For more information, see *z/VM: VM Dump Tool*.

[V6.3] Enhanced Dump Support

z/VM dump capabilities have been enhanced:

- The stand-alone dump utility has been rewritten. The new stand-alone dump creates a CP hard abend format dump, which is usually much smaller than a storage dump, and the dump is written to either ECKD or SCSI DASD.
- Larger memory sizes are supported, up to a maximum of 1 TB. The support includes stand-alone dump, hard abend dump, SNAPDUMP, DUMPLD2, and VM Dump Tool.
- Hard abend dump processing has been enhanced by writing out more than one page of the CP Frame Table per I/O. The CP Frame Table accounts for a significant portion of the dump on large storage systems. Compared with the

previous version, on a system with the same amount of real storage and using the same DASD devices, the amount of time to create a hard abend dump or SNAPDUMP will decrease.

There are new recommendations on how much space to allow for dumps. See *z/VM: CP Planning and Administration*.

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 following utility is added for this support:
- SDINST

The following commands are updated for this support:
- CP DEFINE CPOWNED command
- CP SET DUMP command
- VM Dump Tool FRAMES command

A new DASD operand is added:
- DUMPLD2

The following configuration statement is changed:
- CP_OWNED

The following monitor record is changed:
- D1R7 - MRMTRMEM - Monitor Domain - Memory Configuration data

[V6.4 APAR] DUMP Processing Enhancements

With the PTF for APAR VM65989, the amount of time it takes for z/VM to write a hard abend or snap dump to 3390 DASD may be reduced. The improvements were achieved via changes to the I/O channel program used to write central memory to z/VM spool space located on 3390 DASD.

The SNAPDUMP command has been changed to omit PGMBKs from the dump by default. PGMBKs are the CP structures used to define the virtual storage used by users on the system. In many cases, excluding PGMBKs results in a large reduction in dump size. An optional PGMBKS operand is added to allow PGMBKs to be included in the dump if desired.

For more information, see:
- *z/VM: CP Commands and Utilities Reference*

Product Documentation

These topics describe changes to the format, location, or availability of the z/VM product documentation. For information about which z/VM publications have been updated for the new release, see *z/VM: General Information*.

[V4.1] System Messages and Codes Book Divided

z/VM: System Messages and Codes has been divided into three documents:

- *z/VM: System Messages and Codes - CP*, GC24-6030
- *z/VM: System Messages and Codes - CMS*, GC24-6031
- *z/VM: System Messages and Codes - Other Components*, GC24-6032

[V4.1] Relocated Information

The following table shows information that has been relocated within the z/VM library.

Table 8. Information Relocated in z/VM V4.1

Topic	Old Location	New Location
CMS utilities	*VM/ESA: CMS Utilities Feature*	• *z/VM: CMS Commands and Utilities Reference* • *z/VM: CMS and REXX/VM Messages and Codes*
GCS nucleus build information	*z/VM: Installation Guide*	*z/VM: Group Control System*
CMS nucleus build information	*z/VM: Installation Guide*	*z/VM: Planning and Administration*

[V4.1] Deleted Publications

The following books are not included in the z/VM V4.1 library:

- *LAN Resource Extension and Services/VM: General Information*, GC24-5618
- *LAN Resource Extension and Services/VM: Guide and Reference*, SC24-5622
- *LAN Resource Extension and Services/VM: Licensed Program Specifications*, GC24-5617
- *DCE for VM/ESA: Administration Guide*, SC24-5730
- *DCE for VM/ESA: Administration Reference*, SC24-5731
- *DCE for VM/ESA: Application Development Guide*, SC24-5732
- *DCE for VM/ESA: Application Development Reference*, SC24-5733
- *DCE for VM/ESA: Configuring and Getting Started*, SC24-5734
- *DCE for VM/ESA: Introducing the Distributed Computing Environment*, SC24-5735
- *DCE for VM/ESA: Messages and Codes*, SC24-5736
- *DCE for VM/ESA: Planning*, SC24-5737
- *DCE for VM/ESA: User's Guide*, SC24-5738
- *VM/ESA: CMS Utilities Feature*, SC24-5535
- *VM/ESA: REXX/EXEC Migration Tool for VM/ESA*, GC24-5752

[V4.2] Monitor Records Provided on the Web

z/VM monitor records are now provided on the z/VM web site in the same HTML format as the data areas and control blocks.

[V4.2] Relocated Information

The following table shows information that has been relocated within the z/VM library.

Table 9. Information Relocated in z/VM V4.2

Topic	Old Location	New Location
C sockets API (reference)	• *z/VM: TCP/IP Programmer's Reference* • *OE for VM/ESA: Sockets Reference*	*z/VM: Language Environment 1.8 C Run-Time Library Reference*
Test the installation or service of the DirMaint servers	*Directory Maintenance Facility Program Directory*	*z/VM: Directory Maintenance Facility Tailoring and Administration Guide*

[V4.3] Planning Book Divided

z/VM: Planning and Administration has been divided into four documents:

- *z/VM: CMS Planning and Administration*, SC24-6042
- *z/VM: CP Planning and Administration*, SC24-6043
- *z/VM: Dynamic I/O Configuration Planning and Administration*, SC24-6044
- *z/VM: Saved Segments Planning and Administration*, SC24-6056

[V4.3] Relocated Information

The following table shows information that has been relocated within the z/VM library.

Table 10. Information Relocated in z/VM V4.3

Topics	Old Location	New Location
• Setting up OpenExtensions • DIRPOSIX utility	*z/VM: Planning and Administration*	*z/VM: OpenExtensions User's Guide*
• DFSMS/VM installation	*VM/ESA: DFSMS/VM Function Level 221 Installation and Customization* **Note:** This book has been renamed to *z/VM: DFSMS/VM Customization*.	*DFSMS/VM Function Level 221 Program Directory*
• DirMaint control files	*Directory Maintenance VM/ESA: Diagnosis Reference* (This book has not been updated for the Directory Maintenance Facility feature for z/VM.)	*z/VM: Directory Maintenance Facility Tailoring and Administration Guide*
• Tailorable and nontailorable DirMaint system files	*Directory Maintenance Facility Program Directory*	*z/VM: Directory Maintenance Facility Tailoring and Administration Guide*

[V4.4] Retitled Publications

The following table lists z/VM books that have been retitled (and might have new form numbers).

Table 11. Books Retitled in z/VM V4.4

Old Title and Number	New Title and Number
z/VM: Dynamic I/O Configuration Planning and Administration, SC24-6044	*z/VM: I/O Configuration*, form number unchanged
z/VM: Installation Guide, GC24-5992	*z/VM: Guide for Automated Installation and Service*, GC24-6064

[V4.4] Relocated Information

The following table shows information that has been relocated within the z/VM library.

Table 12. Information Relocated in z/VM V4.4

Topics	Old Location	New Location
• Installation-related CP utilities: INSTALL INSTDEF INSTIIS INSTPLAN INSTVM IPWIZARD MIGR51D MOVE2SFS	*z/VM: Installation Guide*	*z/VM: CP Commands and Utilities Reference*
• Messages for installation-related CP utilities	*z/VM: Installation Guide*	*z/VM: CP Messages and Codes*
• CP access control interface (ACI)	*z/VM: CP Planning and Administration*	*z/VM: CP Programming Services*

[V4.4] Deleted Publications

The following table lists books that are not included in the z/VM V4.4 library.

Table 13. Books Removed in z/VM V4.4

Title and Number	Reason for Removal
Debug Tool User's Guide and Reference, SC09-2137	This product does not support the new C/C++ compiler.
Language Environment for OS/390 & VM: Concepts Guide, GC28-1945	Replaced by *z/OS: Language Environment Concepts Guide*, SA22-7567-03
Language Environment for OS/390 & VM: Debugging Guide and Run-Time Messages, SC28-1942	Replaced by *z/OS: Language Environment Debugging Guide*, GA22-7560-03, and *z/OS: Language Environment Run-Time Messages*, SA22-7566-03
Language Environment for OS/390 & VM: Programming Guide, SC28-1939	Replaced by *z/OS: Language Environment Programming Guide*, SA22-7561-03
Language Environment for OS/390 & VM: Programming Reference, SC28-1940	Replaced by *z/OS: Language Environment Programming Reference*, SA22-7562-03
Language Environment for OS/390 & VM: Run-Time Migration Guide, SC28-1944	Not required for z/VM
Language Environment for OS/390 & VM: Writing Interlanguage Communication Applications, SC28-1943	Replaced by *z/OS: Language Environment Writing ILC Applications*, SA22-7563-01
z/VM: Language Environment 1.8 C Run-Time Library Reference, SC24-6038	Retitled to *C/C++ for z/VM: Run-Time Library Reference*, SC09-7625, and included in the C/C++ for z/VM bookshelf on the *z/VM Collection* CD-ROM.

[V5.1] Retitled Publications

The following table lists z/VM publications that have been retitled.

Table 14. Publications Retitled in z/VM V5.1

Old Title	New Title
VM/ESA: Connectivity Planning, Administration, and Operation	*z/VM: Connectivity*
z/VM: System Messages and Codes – CMS	*z/VM: System Messages and Codes – CMS and REXX/VM*

Table 14. Publications Retitled in z/VM V5.1 (continued)

Old Title	New Title
z/VM: System Messages and Codes – Other Components	*z/VM: System Messages and Codes – AVS, Dump Viewing Facility, GCS, TSAF, and VMSES/E*

[V5.1] Relocated Information

The following table shows information that has been relocated within the z/VM library.

Table 15. Information Relocated in z/VM V5.1

Topics	Old Location	New Location
• "Converting to System Configuration Files"	*z/VM: CP Planning and Administration*	*z/VM: Migration Guide*, "Converting from HCPRIO, HCPSYS, and HCPBOX to Configuration Files" on page 415
• "Linux under z/VM"	*z/VM: Running Guest Operating Systems*	*z/VM: Getting Started with Linux on z Systems*
• "Working with Virtual Networks"	*z/VM: Virtual Machine Operation*	*z/VM: Connectivity*

[V5.1] Deleted Publications

The following table lists publications that are not included in the z/VM V5.1 library.

Table 16. Publications Removed in z/VM V5.1

Title and Number	Reason for Removal
eServer zSeries: Open Systems Adapter-Express Customer's Guide and Reference, SA22-7476	Replaced by *eServer zSeries: Open Systems Adapter-Express Customer's Guide and Reference*, SA22-7935-02.
S/390: Open Systems Adapter-Express Customer's Guide and Reference, SA22-7403	This book is for S/390 servers (G5, G6), which are not supported by z/VM V5.1.
S/390: Planning for the S/390 Open Systems Adapter (OSA-1, OSA-2) Feature, GC23-3870	This book is for S/390 servers (G5, G6), which are not supported by z/VM V5.1.
VM/ESA: Open Systems Adapter Facility User's Guide for OSA-2, SC28-1992	This book is for S/390 servers (G5, G6), which are not supported by z/VM V5.1.
VM/ESA: Programmer's Guide to the Server-Requester Programming Interface for VM/ESA, SC24-5455	IBM intends to withdraw the Server-Requester Programming Interface (SRPI) in z/VM V5.2.
VM/ESA: REXX/VM Primer, SC24-5598	The same topics are covered in *z/VM: REXX/VM User's Guide*, SC24-6114
z/VM: Performance Reporting Facility Function Level 410, SC24-6027	The Performance Reporting Facility functions have been merged into the Performance Toolkit. The Performance Reporting Facility feature is not available with or supported by z/VM V5.1.
z/VM: RealTime Monitor Function Level 410, SC24-6027	The RealTime Monitor functions have been merged into the Performance Toolkit. The RealTime Monitor feature is not available with or supported by z/VM V5.1.
z/VM: System Administration Facility, SC24-6034	IBM intends to withdraw the System Administration Facility from a future release of z/VM.

[V5.2] Retitled Publications

The following table lists z/VM publications that have been retitled.

Table 17. Publications Retitled in z/VM V5.2

Old Title	New Title
z/VM: Getting Started with Linux on zSeries	*z/VM: Getting Started with Linux on System z9 and zSeries*
z/VM Summary for Automated Installation and Service (Tape/CD-ROM Installation)	*z/VM Summary for Automated Installation and Service (Tape Installation)*
z/VM: System Messages and Codes – AVS, Dump Viewing Facility, GCS, TSAF, and VMSES/E	*z/VM: Other Components Messages and Codes*
z/VM: System Messages and Codes – CMS and REXX/VM	*z/VM: CMS and REXX/VM Messages and Codes*
z/VM: System Messages and Codes – CP	*z/VM: CP Messages and Codes*

[V5.2] Deleted Publications

The following table lists publications that are not included in the z/VM V5.2 library.

Table 18. Publications Removed in z/VM V5.2

Title and Number	Reason for Removal
OS/390: DFSMS Program Management, SC27-0806	Replaced by the following publications: • *z/OS MVS Program Management: Advanced Facilities*, SA22-7644-04 • *z/OS MVS Program Management: User's Guide and Reference*, SA22-7643-03

[V5.3] Retitled Publications

The following table lists publications in the z/VM library that have been retitled.

Table 19. Publications Retitled in z/VM V5.3

Old Title	New Title
External Security Interface (RACROUTE) Macro Reference for MVS & VM, GC28-1366	*z/VM: Security Server RACROUTE Macro Reference*, SC24-6150
Resource Access Control Facility: Auditor's Guide, SC28-1342	*z/VM: RACF Security Server Auditor's Guide*, SC24-6143
Resource Access Control Facility: Command Language Reference, SC28-0733	*z/VM: RACF Security Server Command Language Reference*, SC24-6144
Resource Access Control Facility: Diagnosis Guide, GY28-1016	*z/VM: RACF Security Server Diagnosis Guide*, GC24-6145
Resource Access Control Facility: General User's Guide, SC28-1341	*z/VM: RACF Security Server General User's Guide*, SC24-6146
Resource Access Control Facility: Macros and Interfaces, SC28-1345	*z/VM: RACF Security Server Macros and Interfaces*, SC24-6147
Resource Access Control Facility: Messages and Codes, SC38-1014	*z/VM: RACF Security Server Messages and Codes*, GC24-6148
Resource Access Control Facility: Security Administrator's Guide, SC28-1340	*z/VM: RACF Security Server Security Administrator's Guide*, SC24-6142
Resource Access Control Facility: System Programmer's Guide, SC28-1343	*z/VM: RACF Security Server System Programmer's Guide*, SC24-6149
z/VM: Getting Started with Linux on System z9 and zSeries	*z/VM: Getting Started with Linux on System z*

[V5.3] Divided Publications

The following table lists z/VM publications that have been divided.

Table 20. Publications Divided in z/VM V5.3

Old	New
z/VM: Performance Toolkit	• *z/VM: Performance Toolkit Guide*, SC24-6156 • *z/VM: Performance Toolkit Reference*, SC24-6157

[V5.3] Added Publications

The following publications have been added to the z/VM library:

- *z/VM: TCP/IP LDAP Administration Guide*, SC24-6140
- *z/VM: RSCS Networking Diagnosis*, GC24-6151
- *z/VM: RSCS Networking Exit Customization*, SC24-6152
- *z/VM: RSCS Networking Messages and Codes*, GC24-6153
- *z/VM: RSCS Networking Operation and Use*, SC24-6154
- *z/VM: RSCS Networking Planning and Configuration*, SC24-6155

[V5.3] Deleted Publications

The following publications have been removed from the z/VM library:

- *z/VM: Quick Reference*, SC24-6111
- *Resource Access Control Facility: General Information*, GC28-0722
- *Resource Access Control Facility: Migration and Planning*, GC23-3054

[V5.4] z/VM Information Center

IBM now publishes the z/VM documentation in an information center using the IBM Eclipse Help System framework. Publishing in an information center is in addition to z/VM product documentation published in BookManager and PDF format and available through the z/VM Collection, the z/VM Internet Library, and the IBM Publications Center. The content of the z/VM Information Center (information for the z/VM base and optional features) is identical to that in the traditional IBM BookManager and Adobe PDF formats; however, the presentation might differ to some degree.

Advantages of information centers are:

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

The z/VM V5R4 Information Center is available from the following sources:

- z/VM V5R4 Information Center (publib.boulder.ibm.com/infocenter/zvm/v5r4/
- *z/VM V5R4 Information Center*, SK5T-7098-00. This DVD includes the IBM Eclipse Help System framework, which allows you to install the information center on a workstation or intranet. The DVD can be ordered (for a fee) from IBM Publications Center (www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss).

[V5.4] Added Publications

The following publications have been added to the z/VM library:

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

[V6.1] Added Publications

The following publications have been added to the z/VM library:

- *System z: Open Systems Adapter-Express Integrated Console Controller 3215 Support*, SA23-2247
- *System z10: Open Systems Adapter-Express3 Integrated Console Controller Dual-Port User's Guide*, SA23-2266

[V6.1] Deleted Publications

The following z/VM publications have been discontinued:

- *z/VM Summary for Automated Installation and Service (DVD Installation)*, GA76-0406
- *z/VM Summary for Automated Installation and Service (DVD Installation)*, GA76-0407
- *IBM Online Library: z/VM Collection* (CD-ROM version), SK2T-2067. This CD-ROM has been replaced by *IBM Online Library: z/VM Collection on DVD*, SK5T-7054.

[V6.1] Relocated Information

The GSKKYMAN and GSKTRACE commands have been moved from *z/VM: TCP/IP LDAP Administration Guide* to *z/VM: TCP/IP User's Guide*.

[V6.1] Media Changes

The following publications are no longer available as printed books:

- *z/VM: CMS Primer*
- *z/VM: Getting Started with Linux on System z*

These publications are still available in IBM BookManager and Adobe PDF formats.

[V6.2] Relocated Information

LDAP client reason codes and return codes have been moved from *z/VM: TCP/IP LDAP Administration Guide* to *z/VM: TCP/IP Messages and Codes*.

[V6.2] Media Changes

The program directories for the z/VM base product and its facilities, features, and prerequisite IBM products are no longer provided in printed form with the z/VM product. These publications are provided online in PDF files at IBM: z/VM Program Directories (www.ibm.com/vm/progdir/).

[V6.3] Media Changes

License information for z/VM V6.3 is provided on a DVD supplied with the z/VM product:

- *z/VM V6.3 Agreements and License Information*, LC27-5378-00

The DVD includes the following documents:

- International Program License Agreement (PDF)
- *z/VM: License Information* (PDF)
- International Agreement for the Acquisition of Software Maintenance (PDF)
- Third-party notices (notices.htm)

  Note: If these notices are updated between editions of the DVD, an updated notices.txt file is provided on the MAINT 193 disk.

[V6.3] Procedures for Changing the Volume Labels and System Name of a Non-SSI z/VM System

New procedures are provided for performing the following tasks:

- Using DDR or FlashCopy to create a new z/VM system from an existing non-SSI z/VM system.
- Changing the DASD volume labels of the new z/VM system.
- Changing the system name of the new z/VM system.

See *z/VM: CP Planning and Administration*.

[V6.3] SAPL Information Moved

Information about using the Stand-Alone Program Loader (SAPL) and passing IPL parameters has been moved from *z/VM: CP Planning and Administration* to *z/VM: System Operation*.

[V6.3] z/VM Product Documentation in 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.

[V6.4] New z/VM CMS Pipelines Documentation

To support the upgraded z/VM CMS Pipelines, the new *z/VM: CMS Pipelines User's Guide and Reference*, SC24-6252, replaces the following publications:

- *z/VM: CMS Pipelines Reference*, SC24-6169
- *z/VM: CMS Pipelines User's Guide*, SC24-6170

CMS Pipelines HELP files have been replaced with files derived from the new publication.

[V6.4] z/VM OpenStack Documentation

z/VM V6.4 does not support OpenStack releases prior to Liberty. The following publications have been removed from the z/VM V6.4 bibliography and are not included in the z/VM V6.4 documentation provided in IBM Knowledge Center:

- *z/VM: Enabling z/VM for OpenStack (Support for OpenStack Icehouse Release)*, SC24-6248
- *z/VM: Enabling z/VM for OpenStack (Support for OpenStack Juno Release)*, SC24-6249
- *z/VM: Enabling z/VM for OpenStack (Support for OpenStack Kilo Release)*, SC24-6250

[V6.4] z/VM Omnibus PKIT

The *IBM Online Library: z/VM Omnibus PKIT*, SK5T-9509, has replaced the *IBM Online Library: z/VM Collection*, SK5T-7054. Like the *z/VM Collection* (CKIT), the *z/VM Omnibus PKIT* contains product information libraries (PKITs) for current z/VM releases and current IBM licensed programs that run on z/VM. It is available as a zip file that can be downloaded from IBM Publications Center (www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss), and it will be updated quarterly as needed.

Note: The zip file is very large (over 2450 MB), and you need to allow sufficient time for the download.

Chapter 3. Changes to External Interfaces

These topics identify changes to specific external interfaces in the z/VM components (except Language Environment). External interfaces are commands, routines, macros, DIAGNOSE codes, directory control statements, and so on. If an external interface has changed, you might need to change the way you use it. Each change is identified as either ***upwardly compatible*** or ***incompatible***. For definitions of these terms, see "Compatibility Terms" on page 2.

Changes to external interfaces are identified in the following major topics:

Notes:

1. There are no changes to AVS or TSAF external interfaces.
2. For information about changes to Language Environment, see *z/VM: Language Environment User's Guide*.

CP Changes

Changes to CP interfaces are identified in the following topics:

For information about changes to CP monitor records, see IBM: z/VM data areas, control blocks, and monitor records (www.ibm.com/vm/pubs/ctlblk.html).

System Configuration Statements

Table 21 on page 280 lists system configuration statements that have changed. For additional information, see *z/VM: CP Planning and Administration*.

Table 21. Changes to System Configuration Statements

Statement	Changes
BEGIN / END	**Upwardly compatible:** • **[V6.2]** New statements. These statements are not valid on any prior release. When IPLing a prior release, the release displays a message that the BEGIN or END statement is not a valid system configuration file statement. Every statement inside the BEGIN block is then always processed as though it was not inside a BEGIN block. This processing could cause the system configuration to be different than what was intended. Do not use these new statements until you are sure that you will never revert to a release that does not support them.
CHARACTER_DEFAULTS	**INCOMPATIBLE:** • **[V2.3]** Letters A-Z, numbers 0-9, and bytes X'OE' (shift out) and X'0F' (shift in) cannot be defined as the system default line edit symbols (line-delete, character-delete, escape, line-end, and tab).
CP_OWNED	**INCOMPATIBLE:** • **[V6.3]** OWN and SHARED operands are ignored in all environments. CSE is no longer supported. The undocumented function of the SHARED operand to turn off minidisk caching for the specified volume is no longer supported. To turn off minidisk caching for a volume being shared through cross-system link (XLINK), which is still supported for non-SSI systems, specify SHARED YES on the RDEVICE configuration statement for the device. **Upwardly compatible:** • **[V6.2]** OWN and SHARED operands are ignored when configured for an SSI cluster. • **[V6.3]** New recommendations for allocating space for dumps. • **[V6.3]** New operand: RDEV *rdev*.
CPXLOAD	**INCOMPATIBLE:** • **[V5.1]** Operands removed: &TEXT, &TXTLIB. The dual-mode CP is no longer supported; CP runs only in 64-bit mode. The file type of all supplied CP text files is TEXT, not TEXT64. The file type of all supplied CP text libraries is TXTLIB, not T64LIB. **Upwardly compatible:** • **[V5.1]** LOCK and NOLOCK operands no longer have any effect. All symbols are considered resident, which means they cannot be locked or unlocked.
CRYPTO APVIRTUAL	**Upwardly compatible:** • **[V6.3 VM65577]** New statement for specifying APs and domains for shared crypto use. • **[V6.4 VM65942]** Supports new AP types for Crypto Express6S: CEX6A, CEX6C, and CEX6S.
CU	**Upwardly compatible:** • **[V6.3]** New operand: NOPPRCSN. • **[V6.4]** New operands, usage notes, and example: ALIAS, MDISK_SHARE, PAGING_SHARE.
DEFINE COMMAND / CMD	**Upwardly compatible:** • **[V4.4]** New operand: SILENT.

Table 21. Changes to System Configuration Statements (continued)

Statement	Changes
DEFINE LAN	**Upwardly compatible:** • **[V4.3]** New operands: TYPE HIPERSOCKETS, ACCOUNTING, GRANT *userlist*. • **[V5.1]** New operands: IP, ETHERNET.
DEFINE VSWITCH	**INCOMPATIBLE:** • **[V6.1]** New default value for NATIVE *natvid* option is 1. • **[V6.4]** Removed operands: IEDN, INMN. **Upwardly compatible:** • **[V5.1]** New operands: IP, ETHERNET, VLAN, PORTTYPE. • **[V5.2]** New operands: GVRP | NOGVRP. • **[V5.3]** New operands: GROUP | NOGROUP, NATIVE *natvid*. • **[V5.4]** Updated operand description: RDEV. • **[V6.1]** New operands: TYPE, NOUPLINK, VLAN AWARE, NATIVE NONE. • **[V6.2]** New operands: PORTBASED, USERBASED. • **[V6.2]** New operands: BRIDGEPORT, UPLINK. • **[V6.2]** New example for BRIDGEPORT. • **[V6.3 VM65583]** New operands: IVL, GLOBAL, LOCAL. • **[V6.4 VM65925]** Operational differences between USERBASED and PORTBASED VSwitches have been eliminated, although the specific designation might be significant for a relocation within an SSI cluster.
DEVICES	**Upwardly compatible:** • **[V6.1]** New operand: SENSED_BUT_OFFLINE.
DISTRIBUTE	**Upwardly compatible:** • **[V6.4 VM65872]** MAXIMUM can be specified with IUCV NO (meaningful only within an SSI cluster). • **[V6.4 VM65872]** Message HCP3012I (replacing HCP3012E) is issued at IPL if another member of the SSI cluster is configured with a different Distribute IUCV policy (or a different MAXIMUM buffer size) than the local node.
EDEVICE	**Upwardly compatible:** • **[V5.2]** New operands: 1750, PREFERRED, NOTPREFERRED. • **[V6.1]** New operand: XIV. • **[V6.2]** New operands: EQID, NOEQID. • **[V6.2]** New responses for device equivalency ID (EQID) operand. • **[V6.4]** New operand: FLASH.

Table 21. Changes to System Configuration Statements (continued)

Statement	Changes
FEATURES	**INCOMPATIBLE:** • **[V6.4]** Code for the PASSWORDS_ON_CMDS operand has been changed to implement the default "No" values documented in *z/VM: CP Planning and Administration*. The IBM-provided system configuration file will no longer include a statement specifying any PASSWORDS_ON_CMDS values as "Yes". **Upwardly compatible:** • **[V3.1]** New options for automatic start and restart, IPL message control, and forced disconnect timeout. • **[V5.4]** New VALIDATE_SHUTDOWN operand will allow installations to ENABLE or DISABLE the new SHUTDOWN *systemid* validation. • **[V6.1]** New options on the ENABLE and DISABLE operands: STP_TIMESTAMPING, STP_TIMEZONE, STP_TZ, XRC_OPTIONAL, XRC_TEST. • **[V6.3]** NOLIMIT operand now means that the system-defined limit of 99,999 logged on users is the maximum number allowed. • **[V6.4]** New ENABLE and DISABLE operands: PAGING_ALIAS, PAGING_HPF.
IODF	**Upwardly compatible:** • **[V5.4]** Updated operand description: *.
MODIFY COMMAND / CMD	**INCOMPATIBLE:** • **[V6.3]** UCR and OVERRIDE information has been removed. UCR spool files and the OVERRIDE utility are no longer supported. **Upwardly compatible:** • **[V4.4]** New operands: SILENT, NOTSILENT.
MODIFY DIAGNOSE	**INCOMPATIBLE:** • **[V6.3]** UCR and OVERRIDE information has been removed. UCR spool files and the OVERRIDE utility are no longer supported.
MODIFY LAN	**Upwardly compatible:** • **[V5.2]** New operands: PROMISCUOUS, NOPROMISCUOUS. • **[V6.1]** New operand: MACPROTECT.
MODIFY PORT	**Upwardly compatible:** • **[V5.3]** Use the MODIFY PORT statement to define or change the OSA-Express devices that make up a link aggregation group and to set the attributes of a link aggregation group. • **[V6.3 VM65583]** New operands: EXCLUSIVE, SHARED.
MODIFY PRIV_CLASSES	**INCOMPATIBLE:** • **[V6.3]** UCR and OVERRIDE information has been removed. UCR spool files and the OVERRIDE utility are no longer supported.

Table 21. Changes to System Configuration Statements (continued)

Statement	Changes
MODIFY VSWITCH	**INCOMPATIBLE:** • **[V5.1]** 'ANY' is not supported as a value for *vlanid*. It was formerly the default. The default VLAN for the user is now the VLAN ID specified on the DEFINE VSWITCH configuration statement or command. • **[V5.4]** 'DROP' and 'FORWARD' are not supported as a value for ISOLATION. The valid options are 'ON' or 'OFF'. • **[V6.4]** Removed operand: OSDSIM. **Upwardly compatible:** • **[V5.1]** New operand: PORTTYPE. • **[V5.2]** New operands: PROMISCUOUS, NOPROMISCUOUS. • **[V5.3]** New operand: RESTORE. • **[V5.3]** New operand: ISOLATION. • **[V5.4]** New operand: VLAN_counters. Updated message: HCP2846E. • **[V5.4]** Updated operand description: RDEV. • **[V6.1]** New operands: OSDSIM, UPLINK, NIC, MACPROTECT. • **[V6.2]** New operands: PORTNUMBER and VLANID. • **[V6.2]** New operands: PATHMTUDISCOVERY. BRIDGEPORT BUFFERS. • **[V6.3]** New operand: VEPA. • **[V6.4 VM65925]** Updated operands: GRANT, PORTTYPE, NIC, PORTNUMBER, VLANID. • **[V6.4 VM65925]** Updated usage notes.
MULTITHREADING	**Upwardly compatible:** • **[V6.3 VM65586, VM65696]** New statement for defining the multithreading characteristics of the system. • **[V6.4]** Updated usage note to include multithreading information.

OPERATOR_CONSOLES	**INCOMPATIBLE:**
	• **[V4.4]** New operand: SYSTEM_3270.
	This operand is not valid in the OPERATOR_CONSOLES statement for any prior VM release. If your only statement includes this operand, and you need to back out of the new release to your prior VM release, the exposure of a missing console could cause the initialization of the prior VM release to fail with a 1010 wait state. As a temporary measure during the migration to the new release, include two OPERATOR_CONSOLES statements, the first without the SYSTEM_3270 operand and the second with it. For example:
	`OPERATOR_CONSOLES 03E0 0009 001F 00078`
	`OPERATOR_CONSOLES SYSTEM_3270 03E0 0009 001F 00078`

RDEVICE	**Upwardly compatible:** • **[V5.4]** New device type: 3215.
RDEVICE (Advanced Function Printers)	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID.
RDEVICE (Card Punches)	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID.
RDEVICE (Card Readers)	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID.

Table 21. Changes to System Configuration Statements (continued)

Statement	Changes
RDEVICE (Communication Controllers)	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID.
RDEVICE (DASD)	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID.
RDEVICE (Graphic Display Devices)	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID.
RDEVICE (Impact Printers)	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID.
RDEVICE (Special Devices)	**INCOMPATIBLE:** • **[V6.1]** LAN_ADAPTER device type has no function; documentation removed. **Upwardly Compatible:** • **[V6.2]** New operands: EQID, NOEQID.
RDEVICE (Tape Units)	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID.
RDEVICE (Terminals)	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID.
RDEVICE (Unsupported Devices)	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID. • **[V6.3]** Operands DPS=Yes and RESERVE_RELEASE=Yes are the defaults.
RDEVICE (3800 Printers)	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID.
RELOCATION_DOMAIN	**Upwardly compatible:** • **[V6.2]** New statement. This statement is not valid on any prior release. When IPLing a prior release, the release displays a message that the RELOCATION_DOMAIN statement is not a valid system configuration file statement.
SET VARIABLE	**Upwardly compatible:** • **[V6.4]** New statement for defining and setting an environment variable that is accessible to every class G user on the system.
SRM	**Upwardly compatible:** • **[V6.3]** New statement for setting up the HiperDispatch environment. • **[V6.3 VM65586, VM65696]** Horizontal polarization and the rebalance algorithm are not compatible with multithreading, and multithreading will not be enabled if either is requested.
STORAGE	**INCOMPATIBLE:** • **[V3.1]** RIO370 operand removed. • **[V5.1]** Options removed: V=R, V=R_FREE. The V=R area and V=R recovery area are not supported. **Upwardly compatible:** • **[V6.3]** New operands: AGELIST, SIZE, EARLYWRITES, RESERVED SYSMAX, EDEVICE. • **[V6.3 VM65417]** New operands: IOAT, LOCKING. • **[V6.4]** New AGELIST operand: KEEPSLOT.

Table 21. Changes to System Configuration Statements (continued)

Statement	Changes
SYSTEM_IDENTIFIER	**Upwardly compatible:** • **[V6.2]** New operands: LPAR, &LPARNAME. These operands are not valid on any prior release. When IPLing a prior release, the LPAR keyword does not display any error message because the keyword is used as the processor model. Use of the &LPARNAME operand displays an error message indicating that &LPARNAME is not a valid system identifier. In either case, the system identifier is not set as expected. This situation might cause various problems, especially if the system identifier is used as a record qualifier or on an EQUATE statement, causing the system configuration to be different than what was intended. Do not use these new operands until you are sure that you will never revert to a release that does not support them.
SYSTEM_RESIDENCE	**Upwardly compatible:** • **[V3.1]** New VOLID options: &SYSRES, &SYSPARM.
USER_DEFAULTS	**Upwardly compatible:** • **[V3.1]** New operand: MESSAGE_LEVEL. • **[V4.1]** New operand: CPLANGUAGE. • **[V4.1]** Message HCP365I (new format) issued at system initialization.
USER_VOLUME_RDEV	**Upwardly compatible:** • **[V6.3]** New statement for specifying a user DASD volume at a specific real device number.
VMLAN	**Upwardly compatible:** • **[V4.3]** New operands: ACNT, ACCOUNTING. • **[V4.4]** New operand: MACPREFIX *macprefix* • **[V5.1]** New operand: MACIDRANGE. • **[V6.1]** New operands: MACPROTECT, USERPREFIX. • **[V6.4]** Message deleted: HCP3026E. • **[V6.4 VM65925]** New operand: DNA DISABLE | ENABLE.
XLINK_DEVICE_DEFAULTS	**INCOMPATIBLE:** • **[V5.1]** Operands removed: TYPE 3330, TYPE 3340, TYPE 3350, TYPE 3375, TYPE 9345. **Upwardly compatible:** • **[V4.2]** The Class 9 description of the TYPE 3390 operand was changed. Specifying TYPE3390 Class 9 supports the maximum number of cylinders on the 3390 device (10017 or greater). • **[V6.1]** The maximum number of cylinders on the 3390-9 device is 65520.
XLINK_SYSTEM _INCLUDE	**Upwardly compatible:** • **[V6.4]** New AVAILABLE operand to reserve a slot.
XLINK_VOLUME _INCLUDE	**Upwardly compatible:** • **[V4.2]** The description of the CYLINDER operand was changed. The CSE track can now be on a cylinder from 0 to the maximum number of cylinders that the 3390 supports. • **[V6.1]** The maximum number of cylinders on the 3390-9 device is 65520.
XSPOOL_SYSTEM	**INCOMPATIBLE:** • **[V6.3]** Removed. CSE is no longer supported.
XSPOOL_TRACE	**INCOMPATIBLE:** • **[V6.3]** Removed. CSE is no longer supported.

Table 21. Changes to System Configuration Statements (continued)

Statement	Changes
XSPOOL_XLIST_INPUT	**INCOMPATIBLE:** • **[V6.3]** Removed. CSE is no longer supported.
XSPOOL_XLIST_OUTPUT	**INCOMPATIBLE:** • **[V6.3]** Removed. CSE is no longer supported.

User Directory Statements

Table 22 lists user directory statements that have changed. For additional information, see *z/VM: CP Planning and Administration*.

Table 22. Changes to User Directory Statements

Statement	Changes
ACCOUNT	**Upwardly compatible:** • **[V6.2]** When using multiconfiguration virtual machines in an SSI cluster, the ACCOUNT statement(s) within the IDENTITY definition define which account numbers can be used by the multiconfiguration virtual machines. The ACCOUNT statement cannot be specified within a SUBCONFIG definition. To assign different account numbers to SUBCONFIG definitions, the COMMAND statement with the SET ACCOUNT command must be used.
CLASS	**INCOMPATIBLE:** • **[V6.3]** UCR and OVERRIDE information has been removed. UCR spool files and the OVERRIDE utility are no longer supported.
CONSOLE	**Upwardly compatible:** • **[V4.2]** New operand: OBSERVER.
CPU	**INCOMPATIBLE:** • **[V5.2]** Operands removed: VECTOR and NOVECTOR are not supported. • **[V6.1]** CRYPTO operand accepted for compatibility but provides no function; documentation removed. **Upwardly compatible:** • **[V5.3]** The CPU type will default to the primary CPU type.
CRYPTO	**INCOMPATIBLE:** • **[V6.1]** CSU, KEYENTRY, MODIFY, and SPECIAL operands accepted for compatibility but provide no function; documentation removed. **Upwardly compatible:** • **[V4.2]** New operand: APVIRT. • **[V5.1]** New operand: APDED. • **[V5.2]** Operands updated to clarify usage. • **[V5.3]** Information added to APDEDICATED operand. • **[V6.3 VM65577]** Information added to APDEDICATED and DOMAIN operands. • **[V6.4 VM65942]** Supports crypto domains defined on a Crypto Express6S.

Table 22. Changes to User Directory Statements (continued)

Statement	Changes
DEDICATE	**INCOMPATIBLE:** • **[V5.1]** Operand removed: NOIOASSIST. I/O assist is not supported. **Upwardly compatible:** • **[V4.3]** New operands: SINGLEUSER, MULTIUSER. • **[V5.1]** New operand: USERACCESSID *userid*.
DIRECTORY	**INCOMPATIBLE:** • **[V5.1]** Device types 3350, 3370, 3375, 9332, 9335, and 9345 are not supported. Only 3380, 3390, 9336, and FB-512 are supported. **Upwardly compatible:** • **[V6.2]** New option: SSI; if used, *altdev*, **EDIT**, and system affinity operands are not allowed. • **[V6.2]** Information added to explain how DIRECTXA handles tokens.
GLOBALOPTS	**Upwardly compatible:** • **[V6.2]** New operand: CHPIDVirtualization • **[V6.4]** New value for the MACHINE operand: Z.
IDENTITY	**INCOMPATIBLE:** • **[V6.3]** UCR and OVERRIDE information has been removed. UCR spool files and the OVERRIDE utility are no longer supported.
IUCV	**Upwardly compatible:** • **[V5.3]** New operand: *ASYNCMD (Asynchronous CP Command Response system service)
LOAD	**Upwardly compatible:** • **[V6.2]** New **I** (IDENTITY) and **S** (SUBCONFIG) operands. • **[V6.2]** Information added to explain how DIRECTXA handles tokens.
LOADDEV	**Upwardly compatible:** • **[V6.2]** Information added to explain how DIRECTXA handles tokens.
MACHINE	**Upwardly compatible:** • **[V5.1]** 370 operand removed. If a virtual machine defined as 370 mode is IPLed, it is logged on in XA mode with 370 Accommodation turned on. • **[V6.4]** New virtual machine type: Z. • **[V6.4 VM65942]** When logged on, an ESA or XA virtual machine may be put into ESA/390-compatibility mode instead of full ESA/390 mode.
MDISK	**Upwardly compatible:** • **[V3.1]** New volume identification option: &SYSRES. • **[V4.2]** New maximum number of cylinders for 3390 models A98, B9C supported. • **[V5.1]** Supports larger FBA disk size. • **[V6.2]** Information added to explain how DIRECTXA handles tokens.
NICDEF	**Upwardly compatible:** • **[V6.1]** New TYPE operands: IEDN, INMN. • **[V6.4]** IEDN and INMN operands accepted for compatability but IEDN and INMN nics will not be defined; documentation removed. • **[V6.4 VM65925]** New operands: PORTNUMBER, PORTTYPE, VLAN, PROMISCUOUS and NOPROMISCUOUS.

Table 22. Changes to User Directory Statements (continued)

Statement	Changes
OPTION	**INCOMPATIBLE:** • **[V5.1]** Operands removed: DEDICATE, NODEDICATE, VIRT=REAL, V=R, VIRT=FIXED, V=F. V=R and V=F virtual machines are not supported. • **[V6.2]** Many operands are not allowed in a subconfiguration stanza. **Upwardly compatible:** • **[V2.2]** Current LKFAC operand authorizes full-pack minidisks and devices for real MPLF use. For dedicated devices, MPLF channel commands might now succeed where they used to fail. For full-pack minidisks, the change is transparent until you issue the new SET LKFACR command. • **[V2.3]** Specifying the TODENABLE operand allows a user to change the virtual machine TOD clock with the new CP SET VTOD command. • **[V2.3]** New operands: CFVM, CFUSER. • **[V2.4]** New operand: DIAG88. • **[V4.3]** New operands: NETACCOUNTING, NETROUTER. • **[V5.2]** New operands: CRYMEASURE. • **[V5.3]** New operand: NOMEMASSIST. • **[V6.1]** New operand: LXAPP. • **[V6.2]** New operand: CHPIDVirtualization. • **[V6.4]** QUICKDSP function now less meaningful because virtual processor management improvements mean no users stay in the eligible list more than an instant.
SHARE	**Upwardly compatible:** • **[V6.1]** LIMITHARD description updated.
SPECIAL	**INCOMPATIBLE:** • **[V3.1]** Changed message: HCP2801E. **Upwardly compatible:** • **[V2.3]** New operand: MSGPROC. • **[V4.2]** New operands: SCTC, BCTC, CNC, *n*, HIPER, *devs*, *ownerid*, and *lanname*. • **[V4.4]** New operand: FCTC. • **[V6.2]** Information added to explain how DIRECTXA handles tokens.
SPOOL	**Upwardly compatible:** • **[V6.2]** Added clarification for how DIRECTXA handles tokens.
SYSAFFIN	**Upwardly compatible:** • **[V6.2]** SYSAFFIN specifications are not allowed when an SSI directory is being processed.
USER	**INCOMPATIBLE:** • **[V2.3]** Letters A-Z, numbers 0-9, and bytes X'OE' (shift out) and X'0F' (shift in) cannot be defined as logical line edit symbols (line-end, line-delete, character-delete, and escape). • **[V5.4]** Usage note updated for logical partition modes. • **[V6.2]** SYSTEMMP not allowed as a user ID. • **[V6.2]** The number for the *pri* option must be from 0 to 99. • **[V6.3]** UCR and OVERRIDE information has been removed. UCR spool files and the OVERRIDE utility are no longer supported.

Table 22. Changes to User Directory Statements (continued)

Statement	Changes
XSTORE	**Upwardly compatible:** • **[V6.4]** Supported for compatibility only; this statement has no function. Expanded storage (XSTORE) is not supported for host or guest use. If a directory entry includes the XSTORE statement, message HCP1401I will be issued when the guest logs on to state that XSTORE is not supported.

CP Commands

Table 23 lists CP commands that have changed. For additional information, see *z/VM: CP Commands and Utilities Reference*.

Table 23. Changes to CP Commands

Command	Changes
ACNT	**Upwardly compatible:** • **[V4.3]** Entering ACNT creates all available types of accounting records for each user specified.
ACTIVATE ISLINK	**Upwardly compatible:** • **[V6.2]** New operand: NODE *nodeid*. • **[V6.2]** Ignored operand: BUFFERS. Message HCP2722I is issued if the BUFFERS option is specified on the ACTIVATE ISLINK command. For compatibility reasons the option may be specified, but a default value of 16 is always used for the link device. If the link device is connecting to a system earlier than z/VM V6.4, change the ACTIVATE ISLINK command on that system to specify BUFFERS 16, or to use the default value of BUFFERS 16. • **[V6.2]** New message: HCP2745E. Updated messages: HCP2702I, HCP2704I, HCP2706I.
ASSOCIATE	**Upwardly compatible:** • **[V6.4 VM66029]** Renamed help files for subcommands. See "[V6.4 APAR] z/VM HELP Quality Improvement" on page 238.

Table 23. Changes to CP Commands (continued)

Command	Changes
ATTACH	**INCOMPATIBLE:** • **[V5.1]** NOIOASSIST operand removed. I/O assist is no longer supported, so there is no need to turn eligibility off. • **[V5.1]** Response for preferred virtual machine removed. Preferred virtual machines are no longer supported. • **[V5.1]** Messages removed: HCP811I, HCP1508I, HCP1511I, HCP1556I. • **[V5.2]** Message removed: HCP6864E. • **[V6.4]** XSTORE/XSTORAGE operands and associated responses removed. Expanded storage (XSTORE) is not supported. **Upwardly compatible:** • **[V3.1]** New messages: HCP6860E, HCP6863E, HCP6864E. • **[V4.3]** New operands: SINGLEUSER, MULTIUSER. • **[V4.3]** New message: HCP1128E. • **[V4.3]** New response for the TYPE operand of FCP to indicate a SCSI device adapter. • **[V4.4]** New operand: VOLID. • **[V5.1]** New option: USERACCESSID userid. • **[V5.1]** New message: HCP6275E. • **[V5.1]** Changed message: HCP2826E. • **[V5.2]** Added message variations: HCP6860E, HCP6863E. • **[V5.2]** New PAV responses. • **[V5.3]** Added HyperPAV to responses. • **[V5.3]** New option: SYSASCII and associated responses. • **[V5.3]** New option: KEY *keyalias*. • **[V5.4]** New usage note. New messages: HCP2462E, HCP2463E. • **[V6.1]** New message: HCP6430E. • **[V6.2]** New operand: EQID. • **[V6.2]** New messages: HCP048E, HCP125E, HCP130E, HCP1128E, HCP1156I, HCP6643E, HCP6644E, HCP6645E, HCP6863E. • **[V6.3 VM65417]** New operand: PCIFUNCTION. • **[V6.3 VM65417]** New message variations: HCP021E, HCP022E, HCP040E. • **[V6.3 VM65417]** Changed messages: HCP046E, HCP120E, HCP122E, HCP6790I. • **[V6.3 VM65583]** New message: HCP3173E. • **[V6.4]** New message: HCP1168E. • **[V6.4]** New option for dedicated tape device: CPMS • **[V6.4]** New message variation: HCP1128E.
AUTOLOG	**Upwardly compatible:** • **[V3.1]** New message: HCP093E.
BEGIN	**Upwardly compatible:** • **[V3.1]** Supports hexadecimal storage locations up to 16 digits. • **[V5.3]** This command is modified to prevent starting secondary processors that are not operational. • **[V5.3]** New message variation: HCP1455E.
COMMIT	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E.

Table 23. Changes to CP Commands (continued)

Command	Changes
COUPLE	**Upwardly compatible:** • **[V3.1]** New message: HCP6024E. • **[V4.2]** New operands: *ownerid*, *lanname*. • **[V4.2]** New response. • **[V4.2]** New messages: HCP2786E, HCP2787E, HCP2788E, HCP2789E, HCP2791E. • **[V4.2]** Changed message: HCP6011E. • **[V4.3]** Changed message: HCP6024E. • **[V4.4]** New operand: SYSTEM *switchnm*. • **[V5.1]** New messages: HCP6525E and HCP6528E. • **[V6.1]** Changed message: HCP2790E. New messages: HCP3018E, HCP3019I, and HCP3025I. • **[V6.2]** New operand: PORTNUMBER. • **[V6.2]** Changed messages: HCP2788E and HCP6011E. New messages: HCP3034E, HCP3041E, HCP3042E, HCP3043E, and HCP3044E. • **[V6.2]** New response. • **[V6.3 VM65583]** New message: HCP3171E. • **[V6.4]** Changed messages: HCP1982I, HCP2790E. • **[V6.4]** Removed messages: HCP3018E, HCP3019E, HCP3025E. • **[V6.4 VM65925]** New operands: *cfgowner* and *cfgname*. • **[V6.4 VM65925]** Updated operand: PORTNUMBER. • **[V6.4 VM65925]** Updated Usage Notes and Responses. • **[V6.4 VM65925]** Updated messages: HCP3034E, HCP3043E, HCP3046E, HCP6011E.
CPACCESS	**Upwardly compatible:** • **[V6.2]** New messages: HCP6640E, HCP6645E, HCP6648I. • **[V6.2]** Updated messages: HCP101E, HCP102E, HCP103E, HCP104E, HCP105E, HCP106E.
CPHX	**Upwardly compatible:** • **[V5.3]** Changed message: HCP6790I. • **[V6.2]** If executed with the AT command, *userid* is required. • **[V6.2]** Additional functions that CPHX can stop: VMDUMP and VMRELOCATE commands, DIAGNOSE code X'94'. • **[V6.2]** Changed message: HCP6790I. • **[V6.3]** Additional function that CPHX can stop: DELETE EDEVICE. • **[V6.3]** Changed message: HCP6790I.
CPLISTFILE	**Upwardly compatible:** • **[V2.2]** New operands to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** Responses include 4-digit years for FULLDATE and ISODATE operands.
CPXLOAD	**Upwardly compatible:** • **[V5.1]** LOCK and NOLOCK operands no longer have any effect. All symbols are considered resident, which means they cannot be locked or unlocked.

Table 23. Changes to CP Commands (continued)

Command	Changes
DEACTIVE (in general)	**Upwardly compatible:** • **[V6.4 VM66029]** Renamed help files for subcommands. See "[V6.4 APAR] z/VM HELP Quality Improvement" on page 238.
DEACTIVE CONV	**Upwardly compatible:** • **[V6.2]** Updated to allow DEACTIVATE as the command name.
DEACTIVE ISLINK	**Upwardly compatible:** • **[V6.2]** Extended to accept more than one device. • **[V6.2]** Updated to allow DEACTIVATE as the command name. • **[V6.2]** New operand: NODE *nodeid*. • **[V6.2]** New message: HCP2746E. Updated messages: HCP2706I, HCP2720E, HCP2738E.
DEDICATE	**INCOMPATIBLE:** • **[V5.1]** Default of V=R user removed. USER operand must be specified. V=R virtual machine is not supported. • **[V5.1]** Message removed: HCP893E. • **[V5.2]** Message removed: HCP664I. • **[V5.3]** Text for messages HCP892E and HCP896E has changed. **Upwardly compatible:** • **[V5.3]** New message: HCP889E. • **[V5.3]** This command is changed so only virtual CPUs that are the same CPU type as the primary real processor can be dedicated. • **[V6.1]** Message removed: HCP1708I.

Table 23. Changes to CP Commands (continued)

Command	Changes
DEFINE (in general)	**INCOMPATIBLE:** • **[V5.1]** Messages removed: HCP811I, HCP1450E. • **[V5.2]** Messages removed: HCP645E, HCP659E, HCP661E, HCP662I, HCP663E, HCP664I. • **[V5.2]** Operand removed: VECTOR. • **[V6.1]** CRYPTO operand has no function; documentation removed. • Also see specific DEFINE commands listed below. **Upwardly compatible:** • **[V2.3]** New operand: MSGPROC. • **[V2.3]** New messages: HCP045E, HCP260E, HCP1014E, HCP2800E, HCP2801E, HCP2802E, HCP2803E, HCP2804I, HCP2806E, HCP2811E. • **[V2.4]** New operand: EXIT. • **[V4.2]** New operands: CFLINK, LAN, NIC, SUBSTITUTE. • **[V4.2]** New messages: HCP2781E, HCP2782E, HCP2784E, HCP2792E, HCP2794E, HCP2795E, HCP2796E • **[V4.4]** New operand: VSWITCH. • **[V5.2]** New operands: LPAR, PAVALIAS. • **[V5.2]** New message: HCP6867E. • **[V5.2]** Changed messages: HCP864E, HCP1006E, HCP2846E, HCP6806E. • **[V5.3]** New message variations: HCP6867E. • **[V5.3]** New messages: HCP1462E, HCP14644I, HCP1466I, HCP1467E, HCP1468E. • **[V5.3]** New message variation: HCP1006E. • **[V5.4]** New message: HCP116I. • **[V6.1]** New message variation: HCP1006E. • **[V6.1]** New message: HCP6430E. • **[V6.2]** Changed messages: HCP101E, HCP102E, HCP103E, HCP104E, HCP105E, HCP106E. • **[V6.3 VM65418]** New operand: CPUPOOL. • **[V6.3 VM65417]** New operand: PCIFUNCTION. • **[V6.3 VM65417]** New message variation: HCP1006E. • **[V6.4 VM65942]** New message: HCP6296E. New message variation: HCP1006E. • Also see specific DEFINE commands listed below.

Table 23. Changes to CP Commands (continued)

Command	Changes
DEFINE CHPID / PATH	**Upwardly compatible:** • **[V2.2]** New operands: INTEGRATED_SYSTEM_DEVICE, ISD. • **[V2.4]** New operands: CLUSTER_BUS_SENDER_CHANNEL, CBS, FICON, FC, FICON_CONVERTER, FCV, OSA_DIRECT_EXPRESS, OSD, OSA_EXPRESS, OSE. • **[V2.4]** New return codes for message HCP6806E. • **[V3.1]** New operands: CLUSTER_BUS_PEER_CHANNEL, CBP, CLUSTER_BUS_RECEIVER_CHANNEL, CBR, COUPLING_FACILITY_RECEIVER_CHANNEL, CFR, DIRECT_SYSTEM_DEVICE, DSD, EMULATED_I/O, EIO, INTERNAL_COUPLING_PEER_CHANNEL, ICP, IOCLUSTER *iocname*, MANAGED, PEERPATH *nn*. • **[V4.3]** New operand: FCP. • **[V4.4]** New operands: PCHID *nnnn*, CSS *nn*, PEERCSS *nn*. • **[V5.2]** New operands: OSN, OSA_NCP. • **[V5.3]** New operands: COUPLING_OVER_INFINIBAND, CIB, AID, PORT, CSYSTEM. • **[V5.4]** New operands: OSA_CONSOLE_3215, OSA_3215. • **[V6.1]** New operands: OSA_DIRECT_EXPRESS_IEDN, OSX, OSA_DIRECT_EXPRESS_INMN, OSM. • **[V6.2]** New IQD options. • **[V6.3 VM65417]** New IQD and OSD option: NETID. • **[V6.4]** New CHPID type options: COUPLING_OVER_PCIe and CS5. • **[V6.4 VM65865]** New CHPID type options: COUPLING_OVER_ROCE and CL5.
DEFINE COMMAND / CMD	**Upwardly compatible:** • **[V4.4]** New operand: SILENT. • **[V5.1]** The SILENT operand can be specified as SILENTLY.
DEFINE CPOWNED	**INCOMPATIBLE:** • **[V6.3]** OWN and SHARED operands are ignored in all environments. CSE is no longer supported. The undocumented function of the SHARED operand to turn off minidisk caching for the specified volume is no longer supported. If OWN or SHARED is specified, message HCP6634I is issued. To turn off minidisk caching for a volume being shared through cross-system link (XLINK), which is still supported for non-SSI systems, specify SHARED YES on the RDEVICE configuration statement for the device, or specify RDEV OFF for the device on the SET MDCACHE command. **Upwardly compatible:** • **[V4.2]** New operand: RESERVED. • **[V6.2]** OWN and SHARED operands ignored when configured for an SSI cluster. • **[V6.3]** New usage note.

Table 23. Changes to CP Commands (continued)

Command	Changes
DEFINE CPU	**INCOMPATIBLE:** • **[V6.3 VM65577]** CPU type ZAAP not supported on z13. **Upwardly compatible:** • **[V5.3]** New operand: TYPE (CP, IFL ZAAP ZIIP). • **[V5.4]** New TYPE operand: ICF. • **[V6.3 VM65418]** New message: HCP6849E. • **[V6.4]** New VCONFIG mode GENERAL supports the same CPU types as ESA390 mode.
DEFINE CPUPOOL	**Upwardly compatible:** • **[V6.3 VM65680]** When multithreading is enabled, prorated core time is used in the limiting calculation for the CPU pool. When multithreading is not enabled, raw CPU time is used instead.
DEFINE CRYPTO	**INCOMPATIBLE:** • **[V6.1]** Command has no function; documentation removed.
DEFINE CTCA	**Upwardly compatible:** • **[V3.1]** New operands: SCTC, BCTC, CNC. • **[V4.4]** New operand: FCTC.
DEFINE CU / CNTLUNIT	**Upwardly compatible:** • **[V2.4]** CFS operand has been changed to CF because this operand now indicates that the control unit is connected to a coupling facility sender channel path or a cluster bus sender channel path, or both. However, the CFS form of the operand is still supported for compatibility. • **[V2.4]** Range of permitted values for CU_LOGICAL_ADDRESS has been increased. • **[V3.1]** New option for TYPE: *. • **[V3.1]** New operand: MANAGED_PATHS *n*. • **[V4.2]** New operand: FICON_CTC. • **[V4.4]** Changed operand: LINK_address. • **[V4.4]** New operand: CSS *nn*.
DEFINE DEVICE / IODEVICE	**Upwardly compatible:** • **[V2.4]** CFS operand has been changed to CF because this operand now indicates that the control unit is connected to a coupling facility sender channel path or a cluster bus sender channel path, or both. However, the CFS form of the operand is still supported for compatibility. • **[V4.4]** New operand: CSS *nn*. • **[V5.2]** New operands: UNIT 3745, UNIT OSN. • **[V6.4 VM65942]** New operand: SS *nn*. New response.
DEFINE LAN	**Upwardly compatible:** • **[V4.3]** New ACCOUNTING ON and ACCOUNTING OFF operands allow Class B users to control whether accounting records are to be created for the LAN being defined. • **[V5.1]** New operands: IP, ETHERNET. • **[V6.2]** This command will not automatically grant access to SYSTEM for a SYSTEM owned restricted LAN. • **[V6.2]** If executed with the AT command, the OWNERID *ownerid* or OWNERID SYSTEM operands are required.

Table 23. Changes to CP Commands (continued)

Command	Changes
DEFINE MDISK	**Upwardly compatible:** • **[V5.4]** New usage note. New response.
DEFINE MSGPROC	**INCOMPATIBLE:** • **[V3.1]** Changed message: HCP2801E. **Upwardly compatible:** • **[V4.2]** New operand: DEVICES.
DEFINE NIC	**INCOMPATIBLE:** • **[V6.4]** Removed operands: IEDN, INMN. **Upwardly compatible:** • **[V4.4]** Updated information for the HIPERSOCKETS and QDIO operands to include virtual switch information. • **[V6.1]** New TYPE operands: IEDN, INMN. Updated message: HCP2797E. • **[V6.2]** New variations for message HCP2797E. • **[V6.4]** Changed message: HCP2797E.
DEFINE PCIFUNCTION	**Upwardly compatible:** • **[V6.3 VM65577]** New operand: TYPE. • **[V6.3 VM65577]** New message: HCP6873E. • **[V6.3 VM65716]** New operand: ISM. • **[V6.4]** New operand: UID. • **[V6.4 VM65942]** New TYPE operands: RCE, HYL PORT, ROC2 PORT. Updated operand: VFN.
DEFINE STORAGE	**Upwardly compatible:** • **[V3.1]** New operand: CONFIGURATION. • **[V5.4]** New operands: RESERVED, STANDBY. • **[V5.4]** Response contains new fields to indicate the sizes of standby and reserved virtual machine storage. • **[V5.4]** New messages: HCP098E, HCP099E. • **[V5.4]** New message text variations for HCP093E and HCP094E.
DEFINE (Temporary Disk)	**INCOMPATIBLE:** • **[V5.1]** Operands removed: T3350, T3375, T9345, T3370, T9332, T9335.
DEFINE TIMEZONE	**Upwardly compatible:** • **[V6.1]** New message: HCP987E.
DEFINE VECTOR	**INCOMPATIBLE:** • **[V5.2]** Command removed. ESA/390 Vector Facility is not supported.

Table 23. Changes to CP Commands (continued)

Command	Changes
DEFINE VSWITCH	**INCOMPATIBLE:** • **[V6.1]** New default value for NATIVE *natvid* option is 1. • **[V6.4]** Removed operands: IEDN, INMN. **Upwardly compatible:** • **[V5.1]** New operands: IP, ETHERNET, VLAN, PORTTYPE. • **[V5.1]** New message: HCP2846E. • **[V5.2]** New operands: GVRP | NOGVRP. • **[V5.3]** New operands: GROUP | NOGROUP, NATIVE *natvid*. • **[V5.4]** New format of the RDEV operand. • **[V5.4]** Updated operand description: RDEV. • **[V6.1]** New operands: TYPE, NOUPLINK, VLAN AWARE, NATIVE NONE. • **[V6.1]** New message: HCP3024E. Updated messages: HCP2832E, HCP2838I. • **[V6.2]** New operands: PORTBASED, USERBASED. • **[V6.2]** New operands: BRIDGEPORT, UPLINK. • **[V6.2]** New BRIDGEPORT example. • **[V6.2]** New messages: HCP2838I, HCP3045E. • **[V6.3 VM65583]** New operands: IVL, GLOBAL, LOCAL. • **[V6.3 VM65583]** New messages: HCP3160E, HCP3161E, HCP3167E, HCP3170E, HCP3173E, HCP3186E, HCP3196E, HCP3221E. Updated message: HCP2782E. • **[V6.4]** Removed message: HCP3024E. • **[V6.4 VM65925]** Updated operands: PORTBASED, USERBASED.
DEFSEG	**Upwardly compatible:** • **[V5.4]** Page ranges for a DCSS can include pages above 2047 MB (up to 512 GB). • **[V5.4]** New message: HCP340E. • **[V6.3]** A new DCSS with the same name as an existing DCSS does not inherit a reserved storage setting for the old version.
DEFSYS	**INCOMPATIBLE:** • **[V5.1]** Value of '370' not supported with MACHMODE option. 370 virtual machines are not supported. • **[V6.3]** A new NSS with the same name as an existing NSS does not inherit a reserved storage setting for the old version. • **[V6.4]** New MACHMODE value Z.
DELETE (in general)	**INCOMPATIBLE:** • **[V5.2]** Operand removed: VECTOR. **Upwardly compatible:** • **[V4.4]** New operand: VSWITCH. • **[V5.2]** New operand: LPAR. • **[V6.3 VM65418]** New operand: CPUPOOL. • **[V6.3 VM65417]** New operand: PCIFUNCTION.
DELETE CHPID / PATH	**Upwardly compatible:** • **[V2.4]** New return codes for message HCP6806E. • **[V4.4]** New operand: CSS *nn*.

Table 23. Changes to CP Commands (continued)

Command	Changes
DELETE CU / CNTLUNIT	**Upwardly compatible:** • **[V2.4]** CFS operand has been changed to CF because this operand now indicates that the control unit is connected to a coupling facility sender channel path or a cluster bus sender channel path, or both. However, the CFS form of the operand is still supported for compatibility. • **[V4.4]** New operand: CSS *nn*. • **[V4.4]** New message: HCP1006E.
DELETE DEVICE / IODEVICE	**Upwardly compatible:** • **[V2.4]** CFS operand has been changed to CF because this operand now indicates that the control unit is connected to a coupling facility sender channel path or a cluster bus sender channel path, or both. However, the CFS form of the operand is still supported for compatibility. • **[V3.1]** Additional format for message HCP6818E. • **[V4.4]** New operand: CSS *nn*. • **[V4.4]** New message: HCP1006E. • **[V6.4 VM65942]** New operand: SS *nn*. Updated response. New message variation: HCP1006E.
DELETE RDEVICE	**Upwardly compatible:** • **[V3.1]** New message: HCP6862E.
DESTAGE	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E.

Table 23. Changes to CP Commands (continued)

Command	Changes
DETACH (in general)	**INCOMPATIBLE:** • **[V5.1]** Messages removed: HCP811I, HCP893E. • **[V5.2]** Operand removed: VECTOR. • **[V5.2]** Message removed: HCP664I. • **[V6.1]** CRYPTO operand has no function; documentation removed. • **[V6.2]** Message removed: HCP660E. • **[V6.4]** Operand removed: XSTORE. Expanded storage (XSTORE) is not supported. • **[V6.4]** Message removed: HCP1402I. • Also see specific DETACH commands listed below. **Upwardly compatible:** • **[V2.3]** New operand: MSGPROC. • **[V2.3]** Additional messages: HCP260E, HCP2805E, HCP2807E. • **[V4.2]** New operands: CFLINK, LAN, NIC. • **[V4.2]** Additional messages: HCP006E, HCP2783E, HCP2793E, HCP2795E. • **[V4.3]** New message: HCP1119E. • **[V4.3]** New response for the TYPE operand of FCP to indicate a SCSI device adapter. • **[V4.4]** New operand: VSWITCH. • **[V5.1]** New message: HCP6324I. • **[V5.2]** Changed message: HCP660E. • **[V5.2]** New message: HCP6866E. • **[V5.3]** Modified message: HCP892E. • **[V5.3]** New operand: SYSASCII and associated responses. • **[V6.2]** New messages: HCP2793E, HCP3049E. • **[V6.3 VM65417]** New operand: PCIFUNCTION. • **[V6.4]** Modified message: HCP1401I. • Also see specific DETACH commands listed below.
DETACH CFLINK	**INCOMPATIBLE:** • **[V6.2]** Class A authorization removed. Class B and G authorizations are valid.
DETACH CRYPTO	**INCOMPATIBLE:** • **[V6.1]** Command has no function; documentation removed.
DETACH (Real Device)	**Upwardly compatible:** • **[V4.3]** New operand: ALL. • **[V4.4]** New operand: VOLID. • **[V5.2]** New PAV responses. • **[V5.3]** Added HyperPAV to responses. • **[V6.2]** New operand: EQID. • **[V6.2]** New messages: HCP048E, HCP135E.
DETACH LAN	**Upwardly compatible:** • **[V6.2]** If executed with the AT command, the OWNERID *ownerid* or OWNERID SYSTEM operands are required.
DETACH SYSASCII	**Upwardly compatible:** • **[V6.2]** If executed with the AT command, the *userid* operand is required.

Table 23. Changes to CP Commands (continued)

Command	Changes
DETACH (Virtual Device)	**Upwardly compatible:** • **[V6.2]** New message: HCP1809E.
DETACH VECTOR	**INCOMPATIBLE:** • **[V5.2]** Command removed. ESA/390 Vector Facility is not supported.
DETACH VSWITCH	**Upwardly compatible:** • **[V6.3 VM65583]** New message: HCP3166E.
DETACH XSTORE	**INCOMPATIBLE:** • **[V6.4]** Command removed. Expanded storage (XSTORE) is not supported.
DIAL	**INCOMPATIBLE:** • **[V3.1]** The fields of an IP address included in the response no longer contain leading zeros.
DISCARD PINNED	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E.
DISCONNECT	**Upwardly compatible:** • **[V3.1]** Time the system waits between disconnect and automatic logoff can be set on the FEATURES system configuration statement.
DISPLAY (in general)	**INCOMPATIBLE:** • **[V5.2]** Messages removed: HCP645E, HCP785E, HCP1701E. • **[V5.2]** Vector function removed. • **[V6.1]** CDX operand has no function; documentation removed. • Also see specific DISPLAY commands listed below. **Upwardly compatible:** • **[V3.1]** New operand: PSWG. • **[V3.1]** Message HCP6150E might specify z/Architecture mode. • **[V6.3 VM65733]** Messages removed: HCP6153E, HCP6154E. • Also see specific DISPLAY commands listed below.
DISPLAY CDX	**INCOMPATIBLE:** • **[V6.1]** Command has no function; documentation removed.
DISPLAY ESA/XC Storage	See DISPLAY (Guest Storage - ESA/XC).
DISPLAY (Guest Storage - ESA/390)	**INCOMPATIBLE:** • **[V5.2]** In responses that contain a trailer identifying the second-level storage address, the trailer has changed from **L***hexloc* to **R***hexloc*. • **[V5.4]** In the response for the I option, all of the *mnemonic* fields are now 7 characters wide, in which the mnemonic value is left-justified. **Upwardly compatible:** • **[V3.1]** New operands for indirect addressing and displaying data in ASCII format. • **[V3.1]** Supports 64-bit addressing. • **[V3.1]** New z/Architecture variation. • **[V5.3]** Response to the K operand includes the block-usage state and the block-content state.

Table 23. Changes to CP Commands (continued)

Command	Changes
DISPLAY (Guest Storage - ESA/XC)	**INCOMPATIBLE:** • **[V5.4]** In the response for the I option, all of the *mnemonic* fields are now 7 characters wide, in which the mnemonic value is left-justified. **Upwardly compatible:** • **[V3.1]** New operands for indirect addressing and displaying data in ASCII format. • **[V3.1]** Supports 64-bit addressing. • **[V5.3]** Response to the K operand includes the block-usage state and the block-content state.
DISPLAY (Guest Storage - z/Architecture)	**INCOMPATIBLE:** • **[V5.2]** In responses that contain a trailer identifying the second-level storage address, the trailer has changed from **L***hexloc* to **R***hexloc*. • **[V5.4]** In the response for the I option, all of the *mnemonic* fields are now 7 characters wide, in which the mnemonic value is left-justified. **Upwardly compatible:** • **[V5.3]** Response to the K operand includes the block-usage state and the block-content state.
DISPLAY (Host Storage)	**INCOMPATIBLE:** • **[V5.2]** New L and R operands to specify host logical storage or host real storage. L is the default. • **[V5.2]** Response contains additional character, L or R, to indicate contents are from host logical storage or host real storage. • **[V5.2]** New messages: HCP160E, HCP164E. • **[V5.4]** In the response for the I option, all of the *mnemonic* fields are now 7 characters wide, in which the mnemonic value is left-justified. **Upwardly compatible:** • **[V3.1]** New operands for indirect addressing and displaying data in ASCII format. • **[V3.1]** Supports 64-bit addressing. • **[V5.2]** Response might contain a trailer identifying the host real storage address.
DISPLAY LKS	**Upwardly compatible:** • **[V2.2]** Response indicates the called-space identification (CSID) if the linkage-stack entry type is a program-call state entry with a called-space ID. • **[V3.1]** New responses for 64-bit mode.
DISPLAY PSW	**Upwardly compatible:** • **[V3.1]** For a z/Architecture guest, translates a z/Architecture format PSW into an ESA/390 format PSW (if the SET PSWTRANS ALL or SET PSWTRANS STORE command was previously issued).

Table 23. Changes to CP Commands (continued)

Command	Changes
DISPLAY (Registers)	**Upwardly compatible:** • **[V2.4]** New operand: FPC. • **[V2.4]** Yreg operands accept register numbers 0-15 if the IEEE Floating Point hardware feature is installed on the processor. • **[V2.4]** New response when FPC operand is used. • **[V2.4]** New messages: HCP6153E, HCP6154E. • **[V3.1]** New operands: GG*reg1*, GH*reg1*, XG*reg1*, XH*reg1*. • **[V5.2]** New operand: BEAR. • **[V6.3 VM65733]** Updated operands: Y*reg1*, FPC, and END. The Vector Facility for z/Architecture support enables access to the additional floating-point (AFP) registers without requiring prior guest program use.
DISPLAY Vector	**INCOMPATIBLE:** • **[V5.2]** Command removed. ESA/390 Vector Facility is not supported.
DRAIN	**Upwardly compatible:** • **[V6.4 VM66029]** Renamed help files for subcommands. See "[V6.4 APAR] z/VM HELP Quality Improvement" on page 238.
DUMP (in general)	**INCOMPATIBLE:** • **[V5.2]** Messages removed: HCP645E, HCP785E, HCP1701E. • **[V5.2]** Vector function removed. • **[V6.1]** CDX operand has no function; documentation removed. • Also see specific DUMP commands listed below. **Upwardly compatible:** • **[V3.1]** Message HCP6150E might specify z/Architecture mode. • **[V3.1]** New operand: PSWG. • Also see specific DUMP commands listed below.
DUMP CDX	**INCOMPATIBLE:** • **[V6.1]** Command has no function; documentation removed.
DUMP ESA/XC Storage	See DUMP (Guest Storage - ESA/XC)
DUMP (Guest Storage - ESA/390)	**INCOMPATIBLE:** • **[V5.2]** In responses that contain a trailer identifying the second-level storage address, the trailer has changed from **L***hexloc* to **R***hexloc*. • **[V5.4]** In the response for the I option, all of the *mnemonic* fields are now 7 characters wide, in which the mnemonic value is left-justified. **Upwardly compatible:** • **[V3.1]** New operands for indirect addressing and dumping data in ASCII format. • **[V3.1]** Supports 64-bit addressing. • **[V3.1]** New z/Architecture version of this command.
DUMP (Guest Storage - ESA/XC)	**INCOMPATIBLE:** • **[V5.4]** In the response for the I option, all of the *mnemonic* fields are now 7 characters wide, in which the mnemonic value is left-justified. **Upwardly compatible:** • **[V3.1]** New operands for indirect addressing and dumping data in ASCII format. • **[V3.1]** Supports 64-bit addressing.

Table 23. Changes to CP Commands (continued)

Command	Changes
DUMP (Guest Storage - z/Architecture)	**INCOMPATIBLE:** • **[V5.2]** In responses that contain a trailer identifying the second-level storage address, the trailer has changed from **L***hexloc* to **R***hexloc*. • **[V5.4]** In the response for the I option, all of the *mnemonic* fields are now 7 characters wide, in which the mnemonic value is left-justified.
DUMP (Host Storage)	**INCOMPATIBLE:** • **[V5.2]** New L and R operands to specify host logical storage or host real storage. L is the default. • **[V5.2]** Response contains additional character, L or R, to indicate contents are from host logical storage or host real storage. • **[V5.2]** New messages: HCP160E, HCP164E. • **[V5.4]** In the response for the I option, all of the *mnemonic* fields are now 7 characters wide, in which the mnemonic value is left-justified. **Upwardly compatible:** • **[V3.1]** New operands for indirect addressing and dumping data in ASCII format. • **[V3.1]** Supports 64-bit addressing. • **[V5.2]** Response might contain a trailer identifying the host real storage address.
DUMP LKS	**Upwardly compatible:** • **[V2.2]** Response indicates the called-space identification (CSID) if the linkage-stack entry type is a program-call state entry with a called-space ID. • **[V3.1]** New responses for 64-bit mode.
DUMP (Registers)	**Upwardly compatible:** • **[V3.1]** New operands: GG, GH, XG, XH. • **[V6.3 VM65733]** Updated operands: Y*reg1*, FPC, and END. The Vector Facility for z/Architecture support enables access to the additional floating-point (AFP) registers without requiring prior guest program use.
DUMP Vector	**INCOMPATIBLE:** • **[V5.2]** Command removed. ESA/390 Vector Facility is not supported.
DUPLEX	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E.

Table 23. Changes to CP Commands (continued)

Command	Changes
FLASHCOPY	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E. **Upwardly compatible:** • **[V5.3]** New operand: SYNCHRONOUS. • **[V5.3]** New messages: HCP293E, HCP6790I. • **[V5.4]** New operands: SAVELABEL, LABEL *volser*. New usage notes. Updated message: HCP296E. New messages: HCP116I, HCP2450E, HCP2451E, HCP2452E, HCP2453E, HCP2454E, HCP2455E, HCP2456E, HCP2458E, HCP2459E, HCP2460E, HCP2461I, HCP2462E, HCP2463E, HCP2464E, HCP2465I, HCP2466I, HCP2467E, HCP2468E, HCP2469E, HCP3201E. • **[V6.2]** Message variation change: HCP3201E. • **[V6.2]** New message variation: HCP3201E. • **[V6.4 VM66029]** Renamed help files for subcommands. See "[V6.4 APAR] z/VM HELP Quality Improvement" on page 238.
FORCE	**INCOMPATIBLE:** • **[V4.3]** Message HCP6704E is now returned instead of HCP020E if FORCE is issued with no operands. **Upwardly compatible:** • **[V4.3]** New operands for LOGOFF which will send a shutdown signal to the user. • **[V6.2]** New message HCP1801E. • **[V6.4]** New usage note. New message: HCP2118I.
FORWARD	**INCOMPATIBLE:** • **[V5.3]** The minimum abbreviation has been changed to 4 characters (FORW) from 2 characters (FO) to avoid collision with the FOR command.
FREE (in general)	**Upwardly compatible:** • **[V6.4 VM66029]** Renamed help files for subcommands. See "[V6.4 APAR] z/VM HELP Quality Improvement" on page 238.
FREE LOGON	**Upwardly compatible:** • **[V6.2]** When issued on a member of an SSI cluster for a user ID that is defined as a single-configuration virtual machine, the command is propagated to all the other members of the cluster. If the user ID is defined as a multiconfiguration virtual machine, the command is not propagated across the SSI cluster.
GIVE	**Upwardly compatible:** • **[V4.3]** New message: HCP1118E. • **[V6.2]** New messages: HCP1128E and HCP1156I. • **[V6.4 VM66029]** Renamed help files for subcommands. See "[V6.4 APAR] z/VM HELP Quality Improvement" on page 238.
HALT	**Upwardly compatible:** • **[V4.3]** New response for the TYPE operand of FCP to indicate a SCSI device adapter.
HOLD (in general)	**Upwardly compatible:** • **[V6.4 VM66029]** Renamed help files for subcommands. See "[V6.4 APAR] z/VM HELP Quality Improvement" on page 238.

Table 23. Changes to CP Commands (continued)

Command	Changes
HOLD LOGON	**Upwardly compatible:** • **[V6.2]** When issued on a member of an SSI cluster for a user ID that is defined as a single-configuration virtual machine, the command is propagated to all the other members of the cluster. If the user ID is defined as a multiconfiguration virtual machine, the command is not propagated across the SSI cluster.
HYPERSWAP	**Upwardly compatible:** • **[V6.1]** New operands: INCLUDE, EXCLUDE. • **[V6.1]** New messages: HCP6428I, HCP6429E. • **[V6.3]** Real device number description updated for a 5-digit device value in the alternate subchannel set. • **[V6.3]** New option: PDNUM *extent*. • **[V6.3]** New response for HYPERSWAP QUIESCE or RESUME. • **[V6.3 VM65544]** Multi-Target PPRC devices must be specified in a single subchannel set. • **[V6.3 VM65674]** New SWAP options NONE and OFFSEC for varying off additional Multi-Target PPRC secondary devices.
INDICATE (in general)	**INCOMPATIBLE:** • **[V5.2]** Messages removed: HCP645E, HCP785E, HCP1701E. • **[V5.2]** Operand removed: VECTOR. • Also see specific INDICATE commands listed below. **Upwardly compatible:** • **[V6.2]** New messages. • **[V6.3 VM65586, VM65696]** New MULTITHREAD operand to display core utilization information when multithreading is enabled. • Also see specific INDICATE commands listed below.
INDICATE ACTIVE	**Upwardly compatible:** • **[V6.4]** Virtual processor management improvements mean no users stay in the eligible list more than an instant, so users are rarely shown to be in the eligible list, and the count of users in the eligible list is usually 0.

Table 23. Changes to CP Commands (continued)

Command	Changes
INDICATE LOAD	**INCOMPATIBLE:** • **[V5.2]** 'STORAGE-*nnn*%' field remove from response. • **[V5.2]** 'PAGING-*nnnn*' field in response might contain 1 to 8 digits. • **[V5.2]** Vector Facility fields (AVGVEC-*nnn*%, VECTOR-*aaaaa*, VEC-*nnn*%) removed from response. ESA/390 Vector Facility is not supported. • **[V6.3]** The steal percentage is removed from the command response. **Upwardly compatible:** • **[V5.3]** Response now includes the real CPU type. • **[V5.4]** Response includes new CPU type: ICF. • **[V5.4]** Responses changed; several fields are now variable length depending on the value. • **[V6.2]** AT option not allowed when issued via the AT command. • **[V6.2]** New message: HCP1142E. • **[V6.3]** The Class E response is updated to include a polarization indication for each logical CPU. Additionally, the meaning of the CPU utilization for each logical CPU is updated to reflect the percentage of a real CPU that was consumed. Previously, its meaning was the percentage of resources made available to the logical CPU that were consumed. Also, the average CPU utilization value that is displayed in both the Class G and Class E responses is updated accordingly. This value now represents an average value of the portion of a real CPU that each logical CPU was able to consume. • **[V6.4]** Virtual processor management improvements mean no users stay in the eligible list more than an instant, so users are rarely shown to be in the eligible list, and the count of users in the eligible list is usually 0. • **[V6.4]** XSTORE information not included in response. Expanded storage (XSTORE) is not supported.
INDICATE MULTITHREAD	**Upwardly compatible:** • **[V6.4]** New response to indicate configuration change.
INDICATE NSS	**Upwardly compatible:** • **[V5.4]** New system data file type DCSSG in response to indicate DCSS above 2047 MB. • **[V6.3]** The command response is updated to include the count of instantiated pages. • **[V6.4]** XSTORE always displayed as dashes in response. Expanded storage (XSTORE) is not supported.
INDICATE PAGING	**INCOMPATIBLE:** • **[V5.1]** Response no longer includes page residency data for the CP nucleus. All modules in the CP nucleus are now resident. **Upwardly compatible:** • **[V6.2]** Responses changed to expand the number of pages the user has in XSTORE and on DASD. • **[V6.4]** XSTORE always displayed as dashes in response. Expanded storage (XSTORE) is not supported.

Table 23. Changes to CP Commands (continued)

Command	Changes
INDICATE QUEUES	**INCOMPATIBLE:** • **[V6.3]** Response is changed from two users per line to one user per line. **Upwardly compatible:** • **[V3.1]** Responses changed to support 64-bit architecture. • **[V6.3]** The affinity indication for a VMDBK in the INDICATE QUEUES EXP response is no longer displayed. • **[V6.4]** Virtual processor management improvements mean no users stay in the eligible list more than an instant, so users are rarely shown to be in the eligible list.
INDICATE SPACES	**Upwardly compatible:** • **[V3.1]** Responses changed to support 64-bit architecture. • **[V6.1]** Usage note added about user counts greater than 999999. • **[V6.2]** If executed with the AT command, either a user ID or an address space ID (that includes a specific user ID) is required. • **[V6.3]** Response indicates the number of instantiated pages for the address space. • **[V6.4]** XSTORE always displayed as dashes in response. Expanded storage (XSTORE) is not supported.
INDICATE USER	**INCOMPATIBLE:** • **[V6.1]** Usage note added about user counts greater than 999999. • **[V5.1]** IO field in the response no longer includes a plus sign (+) to indicate operation with SIE assist. SIE assist is not supported. • **[V5.2]** In standard response, LOCK field changed to LOCKEDREAL. • **[V5.2]** Response for EXPANDED operand contains new fields. • **[V5.2]** Responses changed to remove Vector Facility information (VVECTIME, TVECTIME). ESA/390 Vector Facility is not supported. **Upwardly compatible:** • **[V3.1]** Responses changed to support 64-bit architecture. • **[V5.3]** Response includes the virtual CPU type and whether CPUAFFINITY is in effect, suppressed, or off. • **[V5.4]** Responses indicate the virtual configuration modes and include the new virtual CPU type: ICF. • **[V6.2]** If executed with the AT command, either SYSTEM or *userid* is required. • **[V6.3]** The command response is updated to include the count of instantiated pages. • **[V6.3 VM65586, VM65696]** If multithreading is enabled, virtual processor time and total processor time are reported as MT-1 equivalent time. • **[V6.4]** Response updated to identify MACH=Z; expanded response updated to identify VCONFIG: MODE=GENERAL. • **[V6.4]** XSTORE always displayed as dashes in response. Expanded storage (XSTORE) is not supported.
INDICATE VECTOR	**INCOMPATIBLE:** • **[V5.2]** Command removed. ESA/390 Vector Facility is not supported.

Table 23. Changes to CP Commands (continued)

Command	Changes
IPL	**INCOMPATIBLE:** • **[V2.2]** New response is issued if tracing is active when a preferred guest is IPLed. Tracing must be turned off and the guest reIPLed. • **[V2.3]** Previously, when you IPLed CMS with the PARM operand, CMS initialization attached a fence of 8 bytes of X'FF's to the end of the PARM data before passing it to the SYSPROF EXEC. Now, no fence is attached, and only the actual PARM data (up to 64 characters) is passed. If you have tailored your SYSPROF EXEC to use the fence to determine the end of the PARM data, you must modify your SYSPROF EXEC to use a different method. • **[V5.1]** Text for message HCP6770E has changed. • **[V5.1]** Messages removed: HCP203E, HCP811I. **Upwardly compatible:** • **[V4.4]** New operand: *fcp-dev*, with STOP and ATTN options. • **[V5.1]** New messages: HCP260E, HCP838E. • **[V5.3]** New messages: HCP175E, HCP838E. • **[V5.3]** This command is not supported for secondary processors. • **[V5.4]** New message: HCP2813E, and updated message: HCP2824E. • **[V6.3]** New operand: NSSDATA. • **[V6.4]** If the virtual machine architecture mode on completion of the NSS IPL differs from that at SAVESYS, the PSW is converted to the appropriate format (that is, from a 64-bit ESA/390 or ESA/XC PSW to a 128-bit z/Architecture PSW, or vice versa).
LINK	**INCOMPATIBLE:** • **[V6.2]** If a requested link has extents that overlap or match an existing link, and either virtual reserve/release or simulated working allegiance is specified for one or both of the minidisks, the link request might now be rejected with message HCP6640E. For more information, see the usage notes on the LINK command in *z/VM: CP Commands and Utilities Reference*. **Upwardly compatible:** • **[V5.4]** New usage note. New response. New message: HCP116I. • **[V6.2]** New messages: HCP6640E, HCP6645E, HCP6648I.
LOCATE (in general)	**INCOMPATIBLE:** • See specific LOCATE commands listed below. **Upwardly compatible:** • **[V5.2]** New operand: SXSTE. • **[V6.3 VM65417]** New operands: RPCI, VPCI. • Also see specific LOCATE commands listed below.
LOCATE CMDBK	**INCOMPATIBLE:** • **[V5.2]** Displayed storage address is host logical. • **[V6.3]** UCR spool files and the OVERRIDE utility are no longer supported. **Upwardly compatible:** • **[V3.1]** New message: HCP026E.

Table 23. Changes to CP Commands (continued)

Command	Changes
LOCATE DGNBK	**INCOMPATIBLE:** • **[V5.2]** Displayed storage address is host logical. **Upwardly compatible:** • **[V3.1]** New message: HCP009E.
LOCATE FILID	**INCOMPATIBLE:** • **[V5.2]** Displayed storage addresses are host logical. **Upwardly compatible:** • **[V3.1]** New message: HCP026E.
LOCATE FRAMETBL	**INCOMPATIBLE:** • **[V5.2]** Displayed storage address of frame table entry is host logical. **Upwardly compatible:** • **[V3.1]** New message: HCP026E. • **[V3.1]** New response formats for 64-bit addresses.
LOCATE ICLBK	**INCOMPATIBLE:** • **[V5.2]** Displayed storage addresses are host logical.
LOCATE LDEV	**INCOMPATIBLE:** • **[V5.2]** Displayed storage addresses are host logical. **Upwardly compatible:** • **[V3.1]** New message: HCP026E.
LOCATE RDEV	**INCOMPATIBLE:** • **[V5.2]** Displayed storage addresses are host logical.
LOCATE SHPBK	**INCOMPATIBLE:** • **[V5.2]** Displayed storage address is host logical.
LOCATE SNABK	**INCOMPATIBLE:** • **[V5.2]** Displayed storage addresses are host logical.
LOCATE SPFBK	**INCOMPATIBLE:** • **[V5.2]** Displayed storage addresses are host logical.
LOCATE (Storage)	**INCOMPATIBLE:** • **[V5.2]** CPVIRTUAL operand removed. No longer supported. • **[V5.2]** New LOGICAL and REAL operands to specify host logical storage or host real storage. CPREAL operand is still supported and is a synonym for REAL. **Upwardly compatible:** • **[V3.1]** New operand: ASCII. • **[V3.1]** New messages: HCP002E, HCP003E, HCP004E, HCP009E, HCP6704E.
LOCATE SYMBOL	**INCOMPATIBLE:** • **[V5.1]** Response contains only one address for each symbol, as all symbols are now resident. • **[V5.2]** Displayed storage address is host logical.

Table 23. Changes to CP Commands (continued)

Command	Changes
LOCATE VDEV	**INCOMPATIBLE:** • **[V5.2]** Displayed storage addresses are host logical. **Upwardly compatible:** • **[V3.1]** New messages: HCP020E, HCP026E.
LOCATE VMDBK	**INCOMPATIBLE:** • **[V5.2]** Displayed storage addresses under VMDBK heading are host logical. New column in response for host real storage address. **Upwardly compatible:** • **[V3.1]** New messages: HCP020E, HCP026E. • **[V5.2]** Supports the SET UNDERSCORE command setting to control the separation of 16-digit output fields.
LOCATE VSMBK	**INCOMPATIBLE:** • **[V5.2]** Displayed storage address is host logical. **Upwardly compatible:** • **[V3.1]** New message: HCP026E.
LOCATE XITBK	**INCOMPATIBLE:** • **[V5.2]** Displayed storage address is host logical.
LOCATEVM	**Upwardly compatible:** • **[V3.1]** New operand: ASCII. • **[V3.1]** New messages: HCP002E, HCP003E, HCP004E, HCP009E, HCP6704E.
LOCK	**INCOMPATIBLE:** • **[V5.1]** SYSTEM operand removed. The CP nucleus is no longer pageable. Modules in the nucleus cannot be locked or unlocked. • **[V5.1]** Text for message HCP295E has changed. • **[V5.2]** New required operand, LOGICAL or REAL, to specify whether guest storage is to be locked to host logical storage or host real storage. Response is also changed. **Upwardly compatible:** • **[V3.1]** New response variations to support 64-bit architecture. • **[V5.1]** SYMBOL and *symbol* operands have no effect. All symbols are now considered resident (locked) and cannot be individually locked.
LOGOFF	**Upwardly compatible:** • **[V6.3 VM65586, VM65696]** If multithreading is enabled, virtual processor time and total processor time are reported as MT-1 equivalent time. • **[V6.4]** XSTORE information not displayed in response. Expanded storage (XSTORE) is not supported.

Table 23. Changes to CP Commands (continued)

Command	Changes
LOGON / LOGIN	**INCOMPATIBLE:** • **[V3.1]** The fields of an IP address included in the response no longer contain leading zeros. • **[V4.1]** 370 operand no longer supported. • **[V5.1]** Messages removed: HCP200E, HCP811I, HCP893E, HCP1106I, HCP1508I, HCP1511I, HCP1556I. • **[V5.2]** Messages removed: HCP662I, HCP663E, HCP664I. • **[V6.4]** New and changed variations for message HCP050E. Change in timing when message HCP053E is issued during logon processing. **Upwardly compatible:** • **[V2.3]** New message: HCP2808E. • **[V3.1]** Can specify larger amount of storage on the STORAGE operand. • **[V3.1]** The response contains a new field after the service level to identify the CP image that is running: '[32-bit]' or '[64-bit]'. • **[V3.1]** New message: HCP093E. • **[V4.3]** New message: HCP2094I. • **[V5.1]** The image field in the response always contains '[64-bit]'. • **[V5.1]** New message: HCP1128E. • **[V5.2]** In the response, the product level has changed. • **[V5.3]** Modified message: HCP892E. • **[V5.3]** New operand: CHANGE • **[V6.2]** New message variation: HCP054E. • **[V6.4]** New MACHINE mode: Z. • **[V6.4]** New message variation: HCP1512E. • **[V6.4]** XSTORE information not displayed in response. Expanded storage (XSTORE) is not supported. • **[V6.4]** Modified message: HCP1401I. • **[V6.4]** Messages removed: HCP1108E, HCP1400I. • **[V6.4 VM65925]** New message: HCP3224I. • **[V6.4 VM65942]** An ESA or XA virtual machine may be put into ESA/390-compatibility mode instead of full ESA/390 mode.
MESSAGE (in general)	**Upwardly compatible:** • **[V6.2]** AT option not allowed when issued via the AT command. • **[V6.2]** New message: HCP1142E. • **[V6.4 VM66029]** New menu file. See "[V6.4 APAR] z/VM HELP Quality Improvement" on page 238.
MESSAGE (User)	**Upwardly compatible:** • **[V6.2]** Responses updated.
MODIFY (in general)	**Upwardly compatible:** • **[V6.3 VM65417]** New operand: PCIFUNCTION. • Also see specific MODIFY commands listed below.

Table 23. Changes to CP Commands (continued)

Command	Changes
MODIFY CHPID / PATH	**Upwardly compatible:** • **[V2.4]** New return codes for message HCP6806E. • **[V4.4]** New operands: CSS *nn*, ADD_CSS_access, RCSS *nn*, INITIAL_access *pname*, DELETE_CSS_access. • **[V4.4]** New messages: HCP864E, HCP1006E. • **[V5.3]** New operand: CSYSTEM. • **[V5.3]** New message variation: HCP1006E. • **[V5.3]** New response that changes the name of a target system to which a CIB channel path is connected.
MODIFY COMMAND / CMD	**Upwardly compatible:** • **[V4.4]** New operands: SILENT, NOTSILENT. • **[V5.1]** The SILENT and NOTSILENT operands can be specified as SILENTLY and NOTSILENTLY. Also, they cannot be used when modifying a specific version of a command with the IBMCLASS *version* operands.
MODIFY CU / CNTLUNIT	**Upwardly compatible:** • **[V2.4]** New operands: TYPE CF. • **[V2.4]** Range of permitted values for CU_LOGICAL_ADDRESS has been increased. • **[V3.1]** New operand: MANAGED_PATHS *n*. • **[V4.4]** Changed operand: LINK_address. • **[V4.4]** New message for the LINK_address: HCP6536E. • **[V4.4]** New operands: CSS *nn*, ADD_CSS_access, RCSS *nn*, DELETE_CSS_access. • **[V4.4]** New message: HCP1006E.
MODIFY DEVICE / IODEVICE	**Upwardly compatible:** • **[V3.1]** New format for message HCP6818E. • **[V4.4]** New operand: CSS *nn*. • **[V4.4]** New messages: HCP864E, HCP1006E. • **[V6.4 VM65942]** New operand: SS *nn*. New message variation: HCP1006E.
MONITOR (in general)	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP6232E. **Upwardly compatible:** • **[V6.3 VM65417]** New class: PCIF. • **[V6.3 VM65417]** New message variations: HCP6240I, HCP6241I. • Also see specific MONITOR commands listed below.
MONITOR EVENT	**Upwardly compatible:** • **[V5.3]** New operand: NETwork. • **[V6.2]** New operand: ISFC. • **[V6.2]** New operand: SSI.

Table 23. Changes to CP Commands (continued)

Command	Changes
MONITOR SAMPLE	**Upwardly compatible:** • **[V5.3]** New operand: NETwork. • **[V6.2]** New operand: ISFC. • **[V6.2]** New operand: CPUMFC • **[V6.2]** If the basic counter set in the CPU-Measurement Facility is not authorized, message HCP6263I is issued for CP MONITOR SAMPLE ENABLE PROCESSOR and CP MONITOR SAMPLE ENABLE ALL. • **[V6.2]** Default size for sample configuration records when sample monitoring is started is increased from 241 pages to 4096 pages. • **[V6.2]** New operand: SSI. • **[V6.3 VM65586, VM65696]** CPUMFC operand does not control collection of CPU-Measurement counter sets for the Simultaneous Multi-Threading (SMT) facility. • **[V6.4]** No data collected for expanded storage. Expanded storage (XSTORE) is not supported.
MONWRITE	**Upwardly compatible:** • **[V5.3]** New operand: CLOSE. • **[V5.3]** New messages: HCP2400E, HCP2401I, HCP6272I, HCP6273A
MSGNOH	**Upwardly compatible:** • **[V6.2]** AT option not allowed when issued via the AT command. • **[V6.2]** New message: HCP1142E.
PURGE UCR	**INCOMPATIBLE:** • **[V6.3]** New usage note. UCR spool files and the OVERRIDE utility are no longer supported.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY (in general)	**INCOMPATIBLE:** • **[V5.1]** Operands removed: IOASSIST, VRFREE, V=R. • **[V5.2]** Operands removed: VECTOR, VIRTUAL VECTOR. • Also see specific QUERY commands listed below. **Upwardly compatible:** • **[V3.1]** New operands and associated responses: PAV, PSWTRANS, UNDERSCORE. • **[V4.2]** New operands and associated responses: CFLINKS, LAN, NIC, SUBSTITUTE, VMLAN, VTOD • **[V4.3]** New operand and associated response: IOPRIORITY • **[V4.4]** New operands and associated responses: CONTROLLER, LOADDEV, VSWITCH. • **[V5.1]** New operands and associated responses: EDEVICE, HYPERSWAP. • **[V5.2]** New operands and associated responses: SXSPAGES, SXSSTORAGE, VIRTUAL PAV. • **[V5.3]** New operands and associated responses: MEMASSIST, SYSASCII, VIRTUAL SYSASCII. • **[V5.3]** New operand and associated response: LSYSTEM. • **[V5.4]** New operand and associated responses: DUMPDEV. • **[V6.1]** New operand and associated response: REORDER. • **[V6.1]** New operand and associated response: STP. • **[V6.2]** New operands and associated responses: CHPIDV, EQID, ISFC, RELODOMAIN, SSI, VIRTUAL CHPID, VMRELOCATE. • **[V6.3]** New operands and associated responses: AGELIST, MSS. • **[V6.3 VM65418]** New operand: CPUPOOL. • **[V6.3 VM65417]** New operands: IOASSIST, PCIFUNCTION, VIRTUAL PCIFUNCTION. **Note:** The function of the new IOASSIST operand is similar to but different from the IOASSIST operand that was removed in V5.1. • **[V6.3 VM65586, VM65696]** New MULTITHREAD operand to show multithreading status and various thread information. • **[V6.4]** New operands and associated responses: CPSERVICE, PAGING, SHUTDOWN, VARIABLE. • Also see specific QUERY commands listed below.
QUERY ABEND	**Upwardly compatible:** • **[V5.2]** New operands: SOFT *aaannn*, OFF, OFF *aaannn*.
QUERY AGELIST	**Upwardly compatible:** • **[V6.4]** New response to indicate whether the auxiliary storage address (ASA) to which a page is written during frame replenishment remains allocated to that page.
QUERY ALL	**INCOMPATIBLE:** • **[V5.1]** Response for SPTAPE tape drive removed. SPTAPE is not supported. **Upwardly compatible:** • **[V4.3]** New responses to indicate tapes attached MULTIUSER, also to indicate a SCSI device adapter with a TYPE of FCP.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY ALLOC	**INCOMPATIBLE:** • **[V5.1]** Response has been changed to support large SCSI LUNs. The fields containing allocation values are expanded and reformatted to support the larger values.
QUERY BYUSER	**Upwardly compatible:** • **[V6.2]** If executed with the AT command, there is no default and *userid* is required.
QUERY CACHE	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E. **Upwardly compatible:** • **[V5.3]** New message variation: HCP6867E.
QUERY CACHEFW	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E. **Upwardly compatible:** • **[V2.2]** Response indicates if the cache fast write function is suspended for the subsystem.
QUERY CAPABILITY	**Upwardly compatible:** • **[V6.1]** Additional response information for nominal CPU capability and capacity changes. • **[V6.3 VM65577]** CPU capability numbers in the response can be integers or decimal fractions. If the number is too large or too small to display, “OVERFLOW” or “UNDERFLO” is displayed.
QUERY CFLINKS	**Upwardly compatible:** • **[V6.2]** If executed with the AT command, the NAME *msgprocid* operands are required.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY CHPID	**Upwardly compatible:** • **[V2.3]** New operand: TYPE. • **[V2.3]** New responses if TYPE is specified. • **[V2.4]** New responses for the TYPE operand to indicate OSA-Express®, OSA Direct Express, cluster-bus-sender, and FICON channels. • **[V3.1]** New responses for the TYPE operand to indicate coupling-facility-sender, cluster-bus-receiver, internal-coupling-sender, internal-coupling-receiver, direct-system-device, emulated-I/O, cluster-bus-peer, coupling-facility-peer, and internal-coupling-peer channels. • **[V4.3]** New response for the TYPE operand of FCP to indicate a SCSI device adapter. • **[V4.4]** New operand: PCHID *nnnn*. • **[V4.4]** New message: HCP1006E. • **[V5.1]** New response for the TYPE operand to indicate an Open Systems Adapter 3270-Console (OSC) channel. • **[V5.2]** New response for the TYPE operand to indicate an Open Systems Adapter NCP (OSN) channel. • **[V5.3]** New response for the TYPE operand to indicate a coupling over InfiniBand (CIB) channel. • **[V6.1]** Updated response for IEDN and INMN. • **[V6.2]** New responses for the TYPE operand to indicate a zHPF channel. • **[V6.3 VM65417]** New operand: NETID. • **[V6.3 VM65417]** New message variation: HCP1006E. • **[V6.4]** New response for the TYPE operand to indicate a coupling over PCIe (CS5) channel. • **[V6.4 VM65865]** New response for the TYPE operand to indicate a coupling over RoCE (CL5) channel.
QUERY COLLECT	**Upwardly compatible:** • **[V6.2]** New responses.
QUERY CONFIGMODE	**Upwardly compatible:** • **[V5.3]** DEFINE LPAR, DELETE LPAR, and SET LSYSTEM commands added to configuration mode.
QUERY CONTROLLER	**INCOMPATIBLE:** • **[V5.1]** Format of the response has changed. • **[V5.2]** Output for command was updated to include the GVRP attribute. • **[V5.4]** Format of the response has changed. **Upwardly compatible:** • **[V5.3]** New response indicates level of the TCP/IP stack running. • **[V6.1]** The response displays the ensemble capability for an internal controller. • **[V6.2]** Response displays HiperSockets Bridge Port capability and status. • **[V6.3]** The response includes VEPA information and indicates if the controlled device is an active device for a virtual switch. • **[V6.3]** Message HCP2832E is added. • **[V6.3 VM65583]** Updated example to show SHARED_LINKAGG. • **[V6.4]** Capabilities NO_ENSEMBLE and NO_INMN always displayed in the response.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY CPCMDS	**INCOMPATIBLE:** • **[V5.2]** Address of CMDBK in response is host logical.
QUERY CPLANGLIST	**Upwardly compatible:** • **[V4.1]** Responses might have new meanings.
QUERY CPLEVEL	**Upwardly compatible:** • **[V2.2]** New operands to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** Response includes 4-digit years for FULLDATE and ISODATE operands. • **[V2.2]** If the user's default date format is changed from SHORTDATE to FULLDATE or ISODATE, and a date format option is not specified on the QUERY CPLEVEL command, the output from QUERY CPLEVEL uses the new default date format. This causes the date to be expanded to include the 4-digit year and also causes fields to the right of the date to be shifted. • **[V3.1]** Response contains a new field after the service level to identify the CP image that is running: '[32-bit]' or '[64-bit]'. • **[V5.1]** Image field in the response always contains '[64-bit]'. • **[V6.4]** Response contains changed values for the new product level.
QUERY CPLOAD	**Upwardly compatible:** • **[V3.1]** Responses have changed because you can no longer load CP from a nucleus, only from a module.
QUERY CPOWNED	**Upwardly compatible:** • **[V6.2]** New response for SSI.
QUERY CPUAFFINITY	**Upwardly compatible:** • **[V6.2]** If executed with the AT command, the *userid* operand is required.
QUERY CPUPOOL	**Upwardly compatible:** • **[V6.3 VM65680]** Response displays CPU pool limits in cores instead of CPUs.
QUERY CPXLOAD	**Upwardly compatible:** • **[V5.1]** Displayed address of customer-written CP routine is host logical.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY CRYPTO	**INCOMPATIBLE:** • **[V4.2]** Revised responses. The keyentry information is no longer issued. • **[V6.1]** DOMAINS and CAMQS operands have no function; documentation removed. • **[V6.1]** Many responses removed or changed. **Upwardly compatible:** • **[V2.4]** New operand: CAMQS. • **[V2.4]** Two new responses. • **[V4.2]** Added APQS parameter. • **[V5.1]** Responses revised to indicate installation of Crypto Adjunct Processor. • **[V5.2]** Response revised with new AP type: CEX2A (Crypto Express2 Accelerator). • **[V6.1]** Response revised with new AP types: CEX3A (Crypto Express3 Accelerator) and CEX3C (Crypto Express3 Coprocessor). • **[V6.1]** New USERS parameter and response. • **[V6.2]** Operand changed from APQS to DOMAINS. • **[V6.2]** Responses updated for DOMAINS operand and two-digit domain values. • **[V6.3]** Response updated with new AP types for Crypto Express4S: CEX4A (accelerator), CEX4C (IBM CCA coprocessor), and CEX4S (neither of the other modes; used for IBM EP11 coprocessor). • **[V6.3 VM65577]** Response updated with new AP types for Crypto Express5S: CEX5A, CEX5C, and CEX5S. • **[V6.3 VM65577]** Responses updated for three-digit AP and domain numbers. • **[V6.3 VM65942]** Response updated with new AP types for Crypto Express6S: CEX6A, CEX6C, and CEX6S.
QUERY CU	**Upwardly compatible:** • **[V5.4]** New operands: SELC, *poolid*, DETAILS. New response. New messages: HCP319E, HCP332E. • **[V6.3]** New operand and associated responses: PPRCSN. • **[V6.3 VM65322]** New operand and associated responses and messages: FENCED. • **[V6.4]** New operand and associated response: ALIAS_SHARE.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY DASD	**Upwardly compatible:** • **[V3.1]** New operand: QUIESCED. • **[V3.1]** Response for QUERY DASD DETAILS might contain a new line containing Parallel Access Volumes information. • **[V5.1]** Response for QUERY DASD DETAILS changed for devices attached to a caching storage controller. • **[V5.3]** Added HyperPAV and CU DETAILS to response. • **[V5.3]** Added new response for device expansion. • **[V5.4]** New operand: SPACE-EFFICIENT. New response. • **[V6.3]** Real device number description updated for a 5-digit device value in the alternate subchannel set. • **[V6.3]** Responses updated to include the device subchannel set number. • **[V6.3 VM65322]** Response for QUERY DASD DETAILS will include Soft Fence and Query Host Access information if supported on the device. • **[V6.3 VM65544]** Response for QUERY DASD DETAILS will indicate if the volume is a Multi-Target PPRC primary device. • **[V6.4]** New operand and associated response and messages: CHARACTERISTICS. Responses updated to include the serial number.
QUERY DASDFW	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E. **Upwardly compatible:** • **[V2.2]** Response indicates if the DASD fast write function is suspended for the subsystem.
QUERY DASD RESERVE	**Upwardly compatible:** • **[V5.4]** Command removed. Information about RESERVE has been moved into the QUERY DASD command.
QUERY DUMP	**Upwardly compatible:** • **[V5.2]** Responses revised to include XF operand, removed PRINTER, ALL and NOPRINT operands. • **[V5.2]** Changed message: HCP9259I.
QUERY DUPLEX	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E.
QUERY EDEVICE	**Upwardly compatible:** • **[V5.2]** Responses have changed to provide information for 1750 and 2107 emulated devices. • **[V5.3]** Responses have been added and changed to provide more detail for emulated devices. • **[V5.3]** Added new response for device expansion. • **[V6.2]** Two responses updated for device equivalency ID (EQID). • **[V6.3]** Response for DETAILS updated to include STATUS condition. • **[V6.3]** New operand: STORAGE; new response about the pools of reserved storage. • **[V6.4]** New response showing the FLASH attribute. • **[V6.4]** New operands and associated responses and messages: INQUIRY, PAGE. New usage note. Responses updated to include the serial number.
QUERY EXITS	**Upwardly compatible:** • **[V2.4]** Additional response information provided for a dynamic CP exit.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY FCP	**Upwardly compatible:** • **[V5.1]** New operand: AGENT. • **[V5.1]** New response for AGENT operand. • **[V5.2]** New WWPN value in responses. • **[V6.1]** New WWPN operand and responses.
QUERY FENCES	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E. **Upwardly compatible:** • **[V6.3 VM65322]** Response indicates if the soft fence state exists for the specified device.
QUERY FRAMES	**INCOMPATIBLE:** • **[V5.2]** Response completely changed in structure and content. **Upwardly compatible:** • **[V3.1]** Responses changed to support 64-bit architecture. • **[V5.1]** The V=R field of the response will always be zero. The V=R area is not supported. • **[V5.1]** The PGNUC field of the response indicates additional frames being used by the CP nucleus. However, the CP nucleus is no longer pageable. • **[V6.3 VM65417]** Responses have been added to provide more details for output.
QUERY GRAF	**INCOMPATIBLE:** • **[V3.1]** The fields of an IP address included in the response no longer contain leading zeros.
QUERY HSA	**INCOMPATIBLE:** • **[V6.2]** Command removed. This function is not supported.
QUERY HYPERSWAP	**Upwardly compatible:** • **[V6.1]** Updated response for automatic quiesce. • **[V6.3]** New operand and associated response: CAPABILITIES. • **[V6.3]** Real device number can be 5 digits if the device is in the alternate subchannel set. • **[V6.3 VM65544]** New CAPABILITIES response token MT to indicate support for Multi-Target PPRC DASD. • **[V6.3 VM65674]** New CAPABILITIES response token OFFSEC to indicate support for OFFSEC operand on HYPERSWAP command.
QUERY IMG	**Upwardly compatible:** • **[V2.2]** New operands to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** Response includes 4-digit years for FULLDATE and ISODATE operands. • **[V2.2]** If the user's default date format is changed from SHORTDATE to FULLDATE or ISODATE, and a date format option is not specified on the QUERY IMG command, the output from QUERY IMG uses the new default date format. This causes the date to be expanded to include the 4-digit year and also causes fields to the right of the date to be shifted.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY IOASSIST	**INCOMPATIBLE:** • **[V5.1]** Command removed. I/O assist is not supported. **Upwardly compatible:** • **[V6.3 VM65417]** A new QUERY IOASSIST command has been added whose function is similar to but different from the command that was removed in V5.1.
QUERY ISLINK	**Upwardly compatible:** • **[V6.2]** New operand: NODe *nodeid*. • **[V6.2]** New message variation: HCP2720E.
QUERY IUCV	**Upwardly compatible:** • **[V6.2]** If executed with the AT command, the *userid* or **sysserv* operand is required.
QUERY KEYALIAS	**Upwardly compatible:** • **[V5.3]** Use the QUERY KEYALIAS command to display current key label alias definitions. Key label aliases are used on the ATTACH command to identify the key label(s) to be used for encrypting data written to a tape. • **[V5.3]** New message: HCP9968E.
QUERY LAN	**INCOMPATIBLE:** • **[V4.4]** In the response, the *ipaddr (n)* address can contain an IPv6 address. • **[V5.1]** Format of the response has changed. • **[V6.4]** Removed operands and associated responses: IEDN, INMN. **Upwardly compatible:** • **[V4.3]** New ACCOUNTING ON and ACCOUNTING OFF responses indicate whether accounting records will be generated for this LAN. • **[V4.4]** New operand: VLAN *vlanid*. • **[V5.2]** New PROMISCUOUS operand and updated response. • **[V5.3]** More information added to the DISCARDED parameter. • **[V5.4]** New operand: USErid *userid*. • **[V6.1]** The response displays ensemble MAC prefix and MACPROTECT settings. • **[V6.1]** New messages: HCP003E, HCP2783E. • **[V6.2]** Port number added to the response. • **[V6.4 VM65925]** Updated Purpose; operands, ACCESSLIST and PROMISCUOUS; and response.
QUERY LDEVS	**INCOMPATIBLE:** • **[V3.1]** The fields of an IP address included in the response no longer contain leading zeros. **Upwardly compatible:** • **[V2.3]** The response might include the IP address for TCP/IP.
QUERY LINKS	**Upwardly compatible:** • **[V6.2]** New message: HCP6638I. • **[V6.2]** New response for SSI.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY LPARS	**Upwardly compatible:** • **[V4.4]** New operand: CSS *nn*. • **[V4.4]** New responses. • **[V5.2]** New operand: RESERVED. • **[V5.2]** New responses. • **[V5.4]** Updated response, new response. • **[V6.1]** New operands: CPCname, NETworkname. • **[V6.1]** New responses for CPCname and NETworkname.
QUERY MAXSPOOL	**Upwardly compatible:** • **[V6.2]** If executed with the AT command, the SYSTEM or *userid* operand is required.
QUERY MDCACHE	**INCOMPATIBLE:** • **[V6.4]** XSTORE information removed from response. Expanded storage (XSTORE) is not supported. **Upwardly compatible:** • **[V5.2]** Additional meaning for value in message HCP1921E. • **[V6.2]** Response updated for SSI. • **[V6.2]** New response to indicate if caching was disabled by transport mode I/O.
QUERY MDISK	**Upwardly compatible:** • **[V6.2]** New operand: RESERVED. • **[V6.2]** New responses for SSI.
QUERY MEMASSIST	**Upwardly compatible:** • **[V6.2]** If executed with the AT command, the *userid* operand is required.
QUERY MONITOR	**Upwardly compatible:** • **[V6.2]** Updated example: The command is extended to show the status of monitoring for the ISFC domain. • **[V6.2]** When PROCESSOR is enabled for SAMPLE, the command displays either CPUMFC or NOCPUMFC. • **[V6.2]** Updated example and responses for SSI. • **[V6.3 VM65417]** Updated example and responses for PCIF.
QUERY MULTITHREAD	**Upwardly compatible:** • **[V6.4]** New INITIAL operand and associated response. • **[V6.4]** New message HCP6513I to indicate configuration is in transition.
QUERY NAMES	**INCOMPATIBLE:** • **[V3.1]** The fields of an IP address included in the response no longer contain leading zeros. **Upwardly compatible:** • **[V6.2]** AT option not allowed when issued via the AT command. • **[V6.2]** Example added for SSI. • **[V6.2]** Responses updated for SSI. • **[V6.2]** New message: HCP1142E.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY NLS	**Upwardly compatible:** • **[V2.2]** New operands to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** Response includes 4-digit years for FULLDATE and ISODATE operands. • **[V2.2]** If the user's default date format is changed from SHORTDATE to FULLDATE or ISODATE, and a date format option is not specified on the QUERY NLS command, the output from QUERY NLS uses the new default date format. This causes the date to be expanded to include the 4-digit year and also causes fields to the right of the date to be shifted.
QUERY NSS	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP1375I. **Upwardly compatible:** • **[V2.2]** New operands to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** Response includes 4-digit years for FULLDATE and ISODATE operands. • **[V2.2]** If the user's default date format is changed from SHORTDATE to FULLDATE or ISODATE, and a date format option is not specified on the QUERY NSS command, the output from QUERY NSS uses the new default date format. This causes the date to be expanded to include the 4-digit year and also causes fields to the right of the date to be shifted. • **[V5.4]** New system data file type DCSSG in response to indicate DCSS above 2047 MB. • **[V5.4]** New MAP response format for DCSS above 2047 MB. • **[V6.4]** Response updated to identify MACHMODE Z.
QUERY NVS	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E.
QUERY OBSERVER	**Upwardly compatible:** • **[V6.2]** New message: HCP1149E. • **[V6.2]** If executed with the AT command, the *userid* or ALL operand is required.
QUERY OSA	**INCOMPATIBLE:** • **[V6.4]** Removed operand: ENSEMBLE. **Upwardly compatible:** • **[V5.2]** New operands: TYPE, ALL, HIPERSOCKETS, OSA, OSN. • **[V5.2]** Three new responses. • **[V5.4]** Updated responses. • **[V6.1]** New TYPE operand: ENSEMBLE. • **[V6.1]** Updated responses. • **[V6.2]** Updated responses for ENSEMBLE and HIPERSOCKETS to indicate bridge-capable devices. • **[V6.3 VM65583]** New operand: SYSTEM.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY PATHS	**Upwardly compatible:** • **[V5.2]** New and changed responses. • **[V6.2]** Updated response to indicate if CHPIDs are enabled for transport mode I/O. • **[V6.3]** Real device number description updated for a 5-digit device value in the alternate subchannel set.
QUERY PAV	**Upwardly compatible:** • **[V5.3]** Added HyperPAV to responses.
QUERY PCIFUNCTION	**Upwardly compatible:** • **[V6.4]** Updated response to include UID field for PCI function. • **[V6.4 VM65942]** Updated response to include PORT field for PCI function.
QUERY PINNED	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E.
QUERY PORT	**Upwardly compatible:** • **[V5.3]** Use the QUERY PORT command to display information about link aggregation groups or devices that have been defined for virtual switches on the system. • **[V5.3]** New message: HCP2837E. • **[V5.3]** Changed message: HCP2783E. • **[V5.4]** Updated example to show ifIndex. • **[V6.3]** The response indicates if the controlled device is an active device for a virtual switch. • **[V6.3 VM65583]** New operand: GROUP *groupname.instance*. • **[V6.3 VM65583]** New and updated examples to show EXCLUSIVE. • **[V6.3 VM65583]** Updated response to include new GROUP *groupname[.instance]* format. New response for when DETAILS are requested for a shared port group. • **[V6.3 VM65583]** New message: HCP3185I. • **[V6.4]** New responses displaying COUNTERS CLEAR information.
QUERY PRIVCLASS	**Upwardly compatible:** • **[V6.2]** If issued with the AT command, there is no default and *userid* is required.
QUERY PROCESSORS	**Upwardly compatible:** • **[V5.3]** Response includes the real CPU type. • **[V5.4]** New operand and associated response: EXPANDED. • **[V5.4]** New real CPU type: ICF. • **[V6.3]** New operand: TOPOLOGY. • **[V6.3]** New response information. • **[V6.3 VM65586, VM65696]** If multithreading is enabled, response includes core information. • **[V6.4]** Expanded response updated to identify partition mode General.
QUERY PROMPT	**INCOMPATIBLE:** • **[V5.1]** Operand removed: AFTER_POWEROFF
QUERY QIOASSIST	**Upwardly compatible:** • **[V6.2]** If issued with the AT command, the *userid* operand is required.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY READER / PRINTER / PUNCH	**INCOMPATIBLE:** • **[V5.1]** Response for file being used by SPTAPE has been removed. SPTAPE is not supported. **Upwardly compatible:** • **[V2.2]** New operands to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** Responses include 4-digit years for FULLDATE and ISODATE operands. • **[V2.2]** New operands DIST and NODIST, available only with operands FULLDATE and ISODATE, specify whether the distribution code is to be included in the response. The default is NODIST, so the output record fits within an 80-character buffer. • **[V2.2]** If the user's default date format is changed from SHORTDATE to FULLDATE or ISODATE, and a date format option is not specified on the QUERY RDR|PRT|PUN command, the output from QUERY RDR|PRT|PUN uses the new default date format. This causes the date to be expanded to include the 4-digit year, the NAME and TYPE fields to the right of the date to be shifted, and the distribution code to be omitted (by default).
QUERY (Real Device)	**INCOMPATIBLE:** • **[V5.1]** Response for tape device being used by SPTAPE has been removed. SPTAPE is not supported. • **[V5.1]** Message removed: HCP2601E. **Upwardly compatible:** • **[V4.3]** New ID operand to display the device and control unit identifiers for a specified device address, if they are known. • **[V4.3]** New response to indicate tapes attached MULTIUSER. • **[V6.2]** Response with ID includes the device equivalency ID (EQID). • **[V6.3 VM65583]** Updated response for Open Systems Adapter Devices to include new status.
QUERY REORDER	**INCOMPATIBLE:** • **[V6.3]** The reorder process is no longer supported. The command is retained for compatibility, but the response always indicates that reorder is OFF.
QUERY RESERVED	**INCOMPATIBLE:** • **[V6.3]** Responses have new formats and include new information. • **[V6.3]** Message removed: HCP007E. **Upwardly Compatible:** • **[V6.3]** New operands and associated responses: DCSS, NSS, SPOOLID, SYSMAX, USER.
QUERY RSAW	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E.
QUERY SECUSER	**Upwardly compatible:** • **[V6.2]** If issued with the AT command, the ALL or *userid* operand is required. • **[V6.2]** New message: HCP1149E.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY SET	**Upwardly compatible:** • **[V5.1]** Additional response information provided for TIMEBOMB status. • **[V5.1]** Because NOTRAN, IOASSIST, and CCWTRAN can no longer be set, these fields in the response are retained only for compatibility and will always indicate NOTRAN OFF, IOASSIST OFF, and CCWTRAN ON. • **[V6.3 VM65417]** IOASSIST OFF removed from response. • **[V6.4]** Response updated to identify MACHINE type Z. • **[V6.4]** Response updated to include DIALDROP status. • **[V6.4 VM65942]** An ESA or XA virtual machine may be running in ESA/390-compatibility mode instead of full ESA/390 mode.
QUERY SHARE	**INCOMPATIBLE:** • **[V5.4]** Response includes the virtual CPU types.
QUERY SHUTDOWNTIME	**INCOMPATIBLE:** • **[V6.4]** The minimum abbreviation for SHUTDOWNTIME is changed from SHUTDOWN to SHUTDOWNT. The output of QUERY SHUTDOWNT is unchanged. • **[V6.4]** The help file for QUERY SHUTDOWNTIME is changed from SHUTDOWN HELCPQU to SHUTDOWT HELPCPQU. Help file SHUTDOWN HELPCPQU is now used for the QUERY SHUTDOWN command. **Upwardly Compatible:** • **[V6.4]** Command syntax variation QUERY SHUTDOWN TIME returns the same information.
QUERY SIGNALS	**Upwardly compatible:** • **[V6.4]** Response can specify SYSTEM as the issuer to indicate a signal from the hardware.
QUERY SPACES	**Upwardly compatible:** • **[V3.1]** Response might contain new values to indicate larger address space.
QUERY SRM	**INCOMPATIBLE:** • **[V6.4]** XSTORE operand removed. Expanded storage (XSTORE) is not supported. • **[V6.4]** New message: HCP1401I. **Upwardly Compatible:** • **[V6.1]** New operand and associated response: LIMITHARD. • **[V6.3]** New operands and associated responses to display HiperDispatch settings: CPUPAD, DSPWDMETHOD, EXCESSUSE, POLARIZATION. • **[V6.3]** Response for MAXWSS can include a larger PAGES value. • **[V6.4]** XSTORE always displayed as dashes in ALL response.
QUERY STGEXEMPT	**Upwardly compatible:** • **[V6.2]** If issued with the AT command, the *userid* or ALL operand is required.
QUERY STORAGE	**Upwardly compatible:** • **[V3.1]** Response might contain new values to indicate greater storage. • **[V5.4]** Response contains new fields to indicate the sizes of configured, standby, and reserved real storage and the real storage increment.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY SYSTEM	**Upwardly compatible:** • **[V5.3]** New response and updated responses for HyperPAV alias. • **[V6.2]** In an SSI cluster where real DASD are shared among the members, QUERY SYSTEM will display only the user IDs of users on the local member.
QUERY TAPES	**INCOMPATIBLE:** • **[V5.1]** Response for tape drive being used by SPTAPE has been removed. SPTAPE is not supported. **Upwardly compatible:** • **[V4.3]** New response to indicate tapes attached MULTIUSER. • **[V5.3]** New operand displays detailed information about real tapes: DETAILS. • **[V5.3]** New response for the tape drive and tapes not part of the library. • **[V6.2]** New response when the DETAILS option is chosen, showing the current tape block position and relative position (for the 3592 Model E07 and higher).
QUERY TIME	**Upwardly compatible:** • **[V2.2]** New operands to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** Response includes 4-digit years for FULLDATE and ISODATE operands. • **[V2.2]** If the user's default date format is changed from SHORTDATE to FULLDATE or ISODATE, and a date format option is not specified on the QUERY TIME command, the output from QUERY TIME uses the new default date format. This causes the date to be expanded to include the 4-digit year. • **[V6.3 VM65586, VM65696]** If multithreading is enabled, virtual processor time and total processor time are reported as MT-1 equivalent time.
QUERY TIMEZONES	**Upwardly compatible:** • **[V6.1]** Response changed to list only the time zones provided by STP as well as the default GMT/UTC when in an STP environment that derives timezone information from the STP server.
QUERY TOKEN	**Upwardly compatible:** • **[V3.1]** New message: HCP6815E.
QUERY TRACE	**Upwardly compatible:** • **[V3.1]** Responses might contain new values and expanded address field for 64-bit. • **[V6.4]** Response changed to indicate transactional execution trace status.
QUERY TRFILES	**Upwardly compatible:** • **[V2.2]** New operands to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** Response includes 4-digit years for FULLDATE and ISODATE operands. • **[V2.2]** If the user's default date format is changed from SHORTDATE to FULLDATE or ISODATE, and a date format option is not specified on the QUERY TRFILES command, the output from QUERY TRFILES uses the new default date format. This causes the date to be expanded to include the 4-digit year and also causes fields to the right of the date to be shifted. • **[V6.3 VM65417]** Updated response with a new FILETYPE: PCI.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY TRSOURCE	**Upwardly compatible:** • **[V5.2]** New LAN operand and response. • **[V6.3 VM65417]** New operand: TYPE PCI. • **[V6.3 VM65417]** New responses for traces of type PCI.
QUERY UCR	**Upwardly compatible:** • **[V2.2]** New operands to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** Response includes 4-digit years for FULLDATE and ISODATE operands. • **[V2.2]** If the user's default date format is changed from SHORTDATE to FULLDATE or ISODATE, and a date format option is not specified on the QUERY UCR command, the output from QUERY UCR uses the new default date format. This causes the date to be expanded to include the 4-digit year and also causes fields to the right of the date to be shifted.
QUERY USERID	**Upwardly compatible:** • **[V3.1]** Response might contain additional fields to indicate IP address. • **[V6.2]** New operand: AT *. • **[V6.2]** AT option not allowed when issued via the AT command. • **[V6.2]** New responses for SSI. • **[V6.2]** New messages: HCP045E, HCP361E, HCP1142E, HCP2971I. • **[V6.2]** New message variation: HCP2970E.
QUERY USERS	**INCOMPATIBLE:** • **[V3.1]** The fields of an IP address included in the response no longer contain leading zeros. **Upwardly compatible:** • **[V6.2]** AT operand information updated for SSI. • **[V6.2]** Responses updated for SSI. • **[V6.2]** New messages: HCP2970E, HCP2971I. • **[V6.3]** Response 1 updated to increase the value for the number of users to 6 digits.
QUERY VCONFIG	**Upwardly compatible:** • **[V6.4]** Response updated to identify MODE = GENERAL.
QUERY VECTOR	**INCOMPATIBLE:** • **[V5.2]** Command removed. ESA/390 Vector Facility is not supported.
QUERY VIRTUAL ALL	**INCOMPATIBLE:** • **[V5.2]** Changed message: HCP662I. • **[V6.2]** Messages removed: HCP662I, HCP1705I. **Upwardly compatible:** • **[V2.3]** Supports virtual message processors. • **[V4.3]** New response for the TYPE operand of FCP to indicate a SCSI device adapter. • **[V5.3]** ASCII console included in response. • **[V6.4]** XSTORE included in response for compatibility and always displayed as dashes. Expanded storage (XSTORE) is not supported.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY VIRTUAL CONSOLE	**INCOMPATIBLE:** • **[V2.3]** Response might include a new line containing TCP/IP information.
QUERY VIRTUAL CPUS	**Upwardly compatible:** • **[V5.3]** Response will include the virtual CPU type, whether secondary CPUs are stopped, and whether CPUAFFINITY is in effect, suppressed, or off. • **[V5.4]** New type CPU: ICF.
QUERY VIRTUAL CRYPTO	**INCOMPATIBLE** • **[V4.2]** One response was deleted. Other responses were revised to include Direct Attached Crypto information. • **[V6.2]** Removed message: HCP1705I. **Upwardly compatible:** • **[V2.4]** New response. • **[V5.2]** Response revised with new AP type: CEX2A (Crypto Express2 Accelerator). • **[V6.1]** Response revised: PCICC, PCICA, and PCIXCC removed as possible values for AP type. • **[V6.1]** Message removed: HCP662I. • **[V6.1]** Response revised with new AP types: CEX3A (Crypto Express3 Accelerator) and CEX3C (Crypto Express3 Coprocessor). • **[V6.2]** Response updated for two-digit domain values. • **[V6.3]** Response updated with new AP types for Crypto Express4S: CEX4A (accelerator), CEX4C (IBM CCA coprocessor), and CEX4S (neither of the other modes; used for IBM EP11 coprocessor). • **[V6.3 VM65577]** Response updated with new AP types for Crypto Express5S: CEX5A, CEX5C, and CEX5S. • **[V6.3 VM65577]** Responses updated for three-digit AP and domain numbers. • **[V6.3 VM65942]** Response updated with new AP types for Crypto Express6S: CEX6A, CEX6C, and CEX6S.
QUERY VIRTUAL CTCA	**INCOMPATIBLE:** • **[V3.1]** Response contains new subclass field. • **[V5.1]** Response no longer indicates eligibility for I/O assist. I/O assist is not supported. **Upwardly compatible:** • **[V4.4]** Response contains new subclass field: FCTC.
QUERY VIRTUAL DASD	**INCOMPATIBLE:** • **[V5.1]** Response no longer indicates eligibility for I/O assist. I/O assist is not supported. **Upwardly compatible:** • **[V5.2]** New PAV responses. • **[V5.3]** Added HyperPAV to response. • **[V5.4]** New SEF information in response.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY (Virtual Device)	**INCOMPATIBLE:** • **[V5.1]** Response no longer indicates eligibility for I/O assist. I/O assist is not supported. **Upwardly compatible:** • **[V2.3]** Response indicates REAL-MPLF or SIMULATED-MPLF if enabled for the device. • **[V2.3]** Supports virtual message devices. The device type MSGD appears in the response. • **[V4.2]** Changed response. • **[V4.3]** New response for the TYPE operand of FCP to indicate a SCSI device adapter. • **[V5.2]** New PAV responses. • **[V5.3]** Added HyperPAV to response. • **[V6.2]** New response when the DETAILS option is chosen, showing the current tape block position and relative position (for the 3592 Model E07 and higher).
QUERY VIRTUAL DUPLEX	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E.
QUERY VIRTUAL FCP	**Upwardly compatible:** • **[V5.1]** Response might include new BYTES information for QDIO devices. • **[V5.2]** Response contains new TOKEN, DEVTYPE, CHPID, and QEBSM information. • **[V6.3]** Response contains new DATA ROUTER ELIGIBLE and DATA ROUTER ACTIVE information.
QUERY VIRTUAL FLASHCOPY	**Upwardly compatible:** • **[V5.4]** Updated authorization and purpose. Updated operands: *vdev*, ALL. New operands: CREATED, OWNER *userid*. New usage notes. New responses. New message: HCP2465I.
QUERY VIRTUAL GRAF	**INCOMPATIBLE:** • **[V3.1]** The fields of an IP address included in the response no longer contain leading zeros. • **[V5.1]** Response no longer indicates eligibility for I/O assist. I/O assist is not supported.
QUERY VIRTUAL LINES	**INCOMPATIBLE:** • **[V5.1]** Response no longer indicates eligibility for I/O assist. I/O assist is not supported.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY VIRTUAL NIC	**INCOMPATIBLE:** • **[V5.1]** The format of the response has changed. **Upwardly compatible:** • **[V4.4]** In the response, the *ipaddr (n)* address can contain an IPv6 address. • **[V4.4]** Updated example to include VLAN information. • **[V5.4]** In the responses, added *macstatus*. • **[V6.1]** The response includes new TYPEs. • **[V6.2]** The response includes additional data. • **[V6.3]** New flag value in response: conflict. • **[V6.4]** New responses displaying COUNTERS CLEAR information. • **[V6.4]** Adapter types IEDN and INMN will not be displayed because they are not supported.
QUERY VIRTUAL OSA	**INCOMPATIBLE:** • **[V5.1]** Response no longer indicates eligibility for I/O assist. I/O assist is not supported. **Upwardly compatible:** • **[V2.4]** New lines in the response for OSA devices that use the Queued-Direct-I/O (QDIO) Facility. • **[V4.2]** Supports virtual OSA devices. • **[V5.1]** Response might include new BYTES information for QDIO devices. • **[V5.2]** Response contains new DEVTYPE, CHPID, TOKEN, and QEBSM information.
QUERY VIRTUAL PAV	**Upwardly compatible:** • **[V5.3]** Added HyperPAV to responses.
QUERY VIRTUAL PRINTER	**INCOMPATIBLE:** • **[V5.1]** Response no longer indicates eligibility for I/O assist. I/O assist is not supported.
QUERY VIRTUAL PUNCH	**INCOMPATIBLE:** • **[V5.1]** Response no longer indicates eligibility for I/O assist. I/O assist is not supported.
QUERY VIRTUAL READER	**INCOMPATIBLE:** • **[V5.1]** Response no longer indicates eligibility for I/O assist. I/O assist is not supported.
QUERY VIRTUAL STORAGE	**Upwardly compatible:** • **[V3.1]** New response to support the new CONFIGURATION operand on the DEFINE STORAGE command. • **[V5.4]** Response contains new fields to indicate the sizes of maximum, standby, and reserved virtual machine storage and the virtual machine storage increment. • **[V5.4]** New options: ELEMENT *n*, ELEMENTS. • **[V5.4]** Response contains new lines to indicate the status of the storage elements. • **[V5.4]** New messages: HCP002E, HCP003E, HCP013E, and HCP026E. Message HCP263E is no longer issued.
QUERY VIRTUAL TAPES	**INCOMPATIBLE:** • **[V5.1]** Response no longer indicates eligibility for I/O assist. I/O assist is not supported.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY VIRTUAL VECTOR	**INCOMPATIBLE:** • **[V5.2]** Command removed. ESA/390 Vector Facility is not supported.
QUERY VIRTUAL XSTORE	**INCOMPATIBLE:** • **[V6.4]** Command removed. Expanded storage is not supported.
QUERY VMDUMP	**Upwardly compatible:** • **[V5.2]** Response revised for virtual machines in ESA/390 or z/Architecture mode and their address ranges.
QUERY VMLAN	**INCOMPATIBLE:** • **[V5.1]** Format of the response has changed. • **[V6.4]** Updated example and response to remove VMLAN Unified Resource Manager. **Upwardly compatible:** • **[V4.3]** New System Accounting and User Accounting responses indicate the current default setting for account record generation for VMLANs. • **[V4.4]** Updated purpose to include virtual switch information. • **[V6.1]** Response includes ensemble information. • **[V6.3 VM65583]** Updated example and response to show IVL membership. • **[V6.4 VM65925]** Updated example and response to show Directory Network Authorization (DNA).
QUERY VRFREE	**INCOMPATIBLE:** • **[V5.1]** Command removed. The V=R recovery area is not supported.

Table 23. Changes to CP Commands (continued)

Command	Changes
QUERY VSWITCH	**INCOMPATIBLE:** • **[V5.1]** Format of the response has changed. • **[V6.4]** Removed operands: IEDN, INMN, OSDSIM. • **[V6.4]** Updated response. **Upwardly compatible:** • **[V5.2]** New PROMISCUOUS operand and updated response. • **[V5.2]** New GVRP attribute for a VLAN-AWARE virtual switch. • **[V5.3]** More information added to the DISCARDED parameter. • **[V5.3]** Updated example and response to show ISOLATION status. • **[V5.4]** New operands: USErid *userid*, RDEV *nnnn*. • **[V5.4]** Updated examples and responses to show VLAN_counters. • **[V5.4]** New error type for the RDEV display; ISOLATION status is always displayed. • **[V6.1]** The response includes new VSWITCH types, the UPLINK port, VLAN default and NATIVE indications, MACPROTECT, and OSDSIM. • **[V6.2]** The response includes additional data. • **[V6.2]** Updated operand descriptions, response and examples. • **[V6.3]** The response includes VEPA information and indicates if the controlled device is an active device for a virtual switch. • **[V6.3]** New flag value in response: conflict. • **[V6.3 VM65583]** New TYPE operand: IVL. • **[V6.3 VM65583]** New example for the IVL virtual switch. Updated examples and responses. • **[V6.3 VM65583]** New message: HCP3185I. • **[V6.4]** New responses displaying COUNTERS CLEAR information. • **[V6.4 VM65925]** Updated operands: ACCESSLIST | PORTNUMBER.
QUERY V=R	**INCOMPATIBLE:** • **[V5.1]** Command removed. The V=R area is not supported.
QUERY XSTORE	**INCOMPATIBLE:** • **[V6.4]** Command removed. Expanded storage is not supported.
REDEFINE	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP811I.
RESET	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP811I. **Upwardly compatible:** • **[V6.2]** New operands: MDISK, Userid. • **[V6.2]** RESET *vdev* cannot be issued via the AT command. • **[V6.2]** USERID *userid* option is required if RESET RESERVE MDISK is issued via the AT command. • **[V6.2]** New response for RESERVE. • **[V6.2]** New messages: HCP006E, HCP020E, HCP045E. • **[V6.3 VM65417]** New operand: PCIFUNCTION. • **[V6.3 VM65417]** New response for PCI. • **[V6.3 VM65417]** New message variations: HCP022E, HCP040E.

Table 23. Changes to CP Commands (continued)

Command	Changes
RETAIN XSTORE	**INCOMPATIBLE:** • **[V6.4]** Command removed. Expanded storage is not supported.
SAVESEG	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP1350E.
SAVESYS	**INCOMPATIBLE:** • **[V5.1]** Text of message HCP1368E has been changed. **Upwardly compatible:** • **[V6.3 VM65733]** New usage note: The SAVESYS command does not save vector registers. • **[V6.4]** New message variation: HCP1368E.
SEND	**Upwardly compatible:** • **[V3.1]** Class C user can send input to any virtual machine. • **[V6.2]** New operand: AT • **[V6.2]** New messages: HCP068E, HCP1149E, HCP2970E, HCP2971I.
SET (in general)	**INCOMPATIBLE:** • **[V5.1]** Operands removed: CCWTRAN, IOASSIST, NOTRANS. • **[V6.1]** CRYPTO operand has no function; documentation removed. • Also see specific SET commands listed below. **Upwardly compatible:** • **[V3.1]** New operands: PSWTRANS, UNDERSCORE. • **[V4.2]** New operands: CFLINK, LAN. • **[V4.3]** New operands: IOPRIORITY, VMLAN. • **[V4.4]** New operands: LOADDEV, VSWITCH. • **[V5.1]** New operand: EDEVICE. • **[V5.2]** New operand: NIC. • **[V5.3]** New operand: MEMASSIST. • **[V5.3]** New operand: LSYSTEM. • **[V5.4]** New operand: STORAGE. • **[V5.4]** New operand: DUMPDEV. • **[V6.1]** New operand: REORDER. • **[V6.1]** New operand: SHUTSIGNAL. • **[V6.2]** New operands: SSI, VMRELOCATE. • **[V6.3]** New operand: AGELIST. • **[V6.3 VM65418]** New operand: CPUPOOL. • **[V6.3 VM65417]** New operand: IOASSIST. **Note:** The function of the new IOASSIST operand is similar to but different from the IOASSIST operand that was removed in V5.1. • **[V6.4]** New operand and associated responses: PAGING. • **[V6.4]** New operand: VARIABLE. • Also see specific SET commands listed below.
SET ABEND	**Upwardly compatible:** • **[V5.2]** New operands: OFF, OFF *aaannn*, SOFT *aaannn*
SET AGELIST	**Upwardly compatible:** • **[V6.4]** New operand and associated response: KEEPSLOT. New and revised usage notes.

Table 23. Changes to CP Commands (continued)

Command	Changes
SET CACHE	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E.
SET CACHEFW	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E.
SET CCWTRAN	**INCOMPATIBLE:** • **[V5.1]** Command removed. V=R virtual machines are not supported.
SET CONFIGMODE	**Upwardly compatible:** • **[V5.3]** DEFINE LPAR, DELETE LPAR, and SET LSYSTEM commands added to configuration mode.
SET CPLANGUAGE	**Upwardly compatible:** • **[V4.1]** Added new parameters to allow the user to specify whether the language is to be changed for the user or for the whole system.
SET CPTRACE	**INCOMPATIBLE:** • **[V6.3]** Trace codes removed: 4005, 4006, 4007, 4008. • **[V5.2]** Trace codes removed: 2820, 2860, 2C20, 4000. • **[V5.3]** Trace codes removed: 4003, 4004. • **[V6.3]** Trace code for SIGP changed from AE00 to AE01. • **[V6.3]** Trace code for Time Stamp changed from FFFE to FFFD. • **[V6.4]** Trace code 4107 removed. Expanded storage (XSTORE) is not supported. • **[V6.4]** Trace codes removed: 2890, 28A0, 2C90, 2CA0 and 2CB0. • **[V6.4]** Unused trace codes 8600 and 8700 were removed. **Upwardly compatible:** • **[V2.4]** New trace category and trace codes for QDIO instructions. • **[V3.1]** New trace codes. • **[V5.2]** New trace codes. • **[V5.3]** New trace codes. • **[V6.2]** New trace codes. • **[V6.2]** New VNET trace codes: 7402, 7403, 7404. • **[V6.3]** New SVC trace codes: 3610, 3611, 3612, 3613. • **[V6.3 VM65417]** New PCI trace codes: 1200-120C, 1211-1213, 1230-1237 and 123A-123D. • **[V6.3 VM65586, VM65696]** New SVC trace code: 3615. • **[V6.3 VM65583]** New VNET trace codes: 7405, 7406, 7407. • **[V6.4]** New CALLRET trace codes: 2891, 28A1, 28A2, 2C91, 2CA1, 2CB1. • **[V6.4]** New RUNU trace codes: 8B00, 8B01, 8B02, 9A01. • **[V6.4]** New CONTAINR trace codes: 6026, 6027, 6028, 6029, 6030. • **[V6.4]** New SVC trace code: 3614.
SET CPUAFFINITY	**Upwardly compatible:** • **[V6.2]** If issued with the AT command, the *userid* operand is required. • **[V6.3 VM65418]** New message: HCP1473E.
SET CPUPOOL	**Upwardly compatible:** • **[V6.3 VM65680]** When multithreading is enabled, prorated core time is used in the limiting calculation for the CPU pool. When multithreading is not enabled, raw CPU time is used instead.

Table 23. Changes to CP Commands (continued)

Command	Changes
SET CRYPTO	**INCOMPATIBLE:** • **[V6.1]** Command has no function; documentation removed.
SET CU	**Upwardly compatible:** • **[V6.3 VM65322]** New operand and associated response and messages: UNFENCE. • **[V6.4]** New operands, usage notes, and example: ALIAS, MDISK_SHARE, PAGING_SHARE.
SET DASDFW	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E.
SET DUMP	**INCOMPATIBLE:** • **[V5.1]** V=R operand and corresponding value in response removed. The V=R area is not supported. • **[V6.2]** This command no longer supports 3424 or 9348 tape drives. **Upwardly compatible:** • **[V3.1]** New operand: XF. • **[V3.1]** New message: HCP1851E. • **[V4.3]** New message: HCP1917E. • **[V5.2]** The operands ALL and NOPRINT were removed. • **[V5.2]** Changed message: HCP9259I. • **[V6.2]** New message: HCP144E. • **[V6.3]** Maximum number of DASD real device numbers increased to 32.
SET EDEVICE	**Upwardly compatible:** • **[V5.2]** New operands: 1750, PREFERRED, NOTPREFERRED. • **[V5.2]** New response. • **[V5.2]** New message: HCP1014E. • **[V5.2]** Message variation: HCP8701I. • **[V5.3]** New message: HCP2777I. • **[V6.1]** New operand: XIV. • **[V6.2]** New operands: EQID, NOEQID. • **[V6.2]** New responses for device equivalency ID (EQID). • **[V6.2]** New message: HCP048E. • **[V6.2]** Descriptions updated for 1750, 2105, 2107, and 2145 devices. • **[V6.4]** New operand: FLASH. • **[V6.4]** Usage note added for SCSI devices.
SET IOASSIST	**INCOMPATIBLE:** • **[V5.1]** Command removed. I/O assist is not supported. **Upwardly compatible:** • **[V6.3 VM65417]** A new SET IOASSIST command has been added whose function is similar to but different from the command that was removed in V5.1.

Table 23. Changes to CP Commands (continued)

Command	Changes
SET IPLPARMS	**INCOMPATIBLE:** • **[V5.1]** IPL parameter ARCH390 is ignored when IPLing a z/VM V5.1 or later CP module because there is no 32-bit version. **Upwardly compatible:** • **[V4.4]** New operand: NOHCD. • **[V6.2]** New conditional IPL parameters: *selector*:*value*(*IPL_parameter_substring*) • **[V6.2]** New message variations: HCP6770E. • **[V6.4]** New IPL parameter: PAGING63. • **[V6.4]** New IPL parameter: IPLVAR=*value* • **[V6.4]** New message variation: HCP6770E.
SET KEYALIAS	**Upwardly compatible:** • **[V5.3]** Use the SET KEYALIAS command to associate a shorter or more meaningful alias name with a key label. Key label aliases are used on the ATTACH command to identify the key label(s) to be used for encrypting data written to a tape. • **[V5.3]** New message: HCP9968E.
SET LAN	**Upwardly compatible:** • **[V4.3]** New ACCOUNTING ON and ACCOUNTING OFF operands allow Class B users to control whether accounting records are to be created for the LAN being defined. • **[V5.2]** New operands: PROMISCUOUS, NOPROMISCUOUS. • **[V6.1]** New MACPROTECT operand. • **[V6.2]** If executed with the AT command, the OWNERID *ownerid* or OWNERID SYSTEM operands are required.
SET LOADDEV	**Upwardly compatible:** • **[V5.1]** New operand allows Class G users to specify the SCPDATA to be passed to the program to be loaded during a guest IPL.
SET MACHINE	**Upwardly compatible:** • **[V4.1]** 370 operand no longer supported. • **[V6.4]** New Z operand and associated response. • **[V6.4]** New message variation: HCP1016E. • **[V6.4 VM65942]** An ESA or XA virtual machine may be put into ESA/390-compatibility mode instead of full ESA/390 mode.
SET MDCACHE	**INCOMPATIBLE:** • **[V6.4]** XSTORE operand and associated response removed. Expanded storage (XSTORE) is not supported. **Upwardly compatible:** • **[V3.1]** Supports defining a minidisk cache larger than 2 GB for z/Architecture virtual machines. • **[V5.2]** Additional meaning for value in message HCP1921E. • **[V5.3]** Additional meaning for value in message HCP1921E. • **[V6.2]** New message: HCP6630E. • **[V6.2]** Up to 2 MB of real storage can be used for minidisk caching.

Table 23. Changes to CP Commands (continued)

Command	Changes
SET MITIME	**Upwardly compatible:** • **[V6.4]** Command has no effect on EDEVICEs. • **[V6.4]** New message: HCP287I.
SET MEMASSIST	**Upwardly compatible:** • **[V6.2]** If issued with the AT command, the ALL or *userid* operand is required.
SET NIC	**Upwardly compatible:** • **[V5.4]** Privilege Class: B format has been added. • **[V6.1]** New MACPROTECT operands. Updated message: HCP2850E • **[V6.2]** New variations for message HCP2850E. • **[V6.2]** MACPROTECT description updated. • **[V6.4]** Removed usage note. • **[V6.4]** Removed message: HCP2850E.
SET NOTRANS	**INCOMPATIBLE:** • **[V5.1]** Command removed. V=R virtual machines are not supported.
SET NVS	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E. I/O assist is not supported.
SET OBSERVER	**Upwardly compatible:** • **[V6.2]** New messages: HCP6768I, HCP1149E.
SET PORT GROUP	**Upwardly compatible:** • **[V5.3]** Use the SET PORT GROUP command to define or change the OSA-Express devices that make up a link aggregation group and to set the attributes of a link aggregation group. • **[V5.3]** New message: HCP2857E. • **[V5.3]** Changed messages: HCP2783E, HCP2799E, HCP6706E. • **[V5.4]** New format of the RDEV operand. • **[V5.4]** Updated operand description: RDEV. • **[V6.2]** New message variation: HCP2838I. • **[V6.3]** New message variation: HCP2838I. • **[V6.3 VM65583]** New operands: EXCLUSIVE, SHARED. • **[V6.3 VM65583]** New messages: HCP3162E, HCP3167E, HCP3173E, HCP3178I, HCP3180E, HCP3183E, HCP3184E, HCP3186E, HCP3194I, HCP3197E, HCP3198E, HCP3200E, HCP3220E.
SET PROMPT	**INCOMPATIBLE:** • **[V5.1]** Operand removed: AFTER_POWEROFF.
SET QIOASSIST	**Upwardly compatible:** • **[V6.2]** If issued with the AT command, the ALL or *userid* operand is required.
SET QUICKDSP	**Upwardly compatible:** • **[V6.4]** Function now less meaningful because virtual processor management improvements mean no users stay in the eligible list more than an instant.
SET RDEVICE (in general)	**Upwardly compatible:** • **[V3.1]** New messages: HCP6580E, HCP6862E, HCP6870E. • **[V5.4]** New device type: 3215.

Table 23. Changes to CP Commands (continued)

Command	Changes
SET RDEVICE Advanced Function Printers	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID. • **[V6.2]** New message: HCP048E.
SET RDEVICE Card Punches	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID. • **[V6.2]** New message: HCP048E.
SET RDEVICE Card Readers	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID. • **[V6.2]** New message: HCP048E.
SET RDEVICE Communication Controllers	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID. • **[V6.2]** New message: HCP048E.
SET RDEVICE DASD	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID. • **[V6.2]** New message: HCP048E.
SET RDEVICE Graphic Display Devices	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID. • **[V6.2]** New message: HCP048E.
SET RDEVICE Impact Printers	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID. • **[V6.2]** New message: HCP048E.
SET RDEVICE Integrated Communication Adapters	**INCOMPATIBLE:** • **[V5.1]** Operands removed: TYPE ICA_BSC and TYPE ICA_SDLC **Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID. • **[V6.2]** New message: HCP048E.
SET RDEVICE Special Devices	**INCOMPATIBLE:** • **[V6.1]** LAN_ADAPTER operand has no function; documentation removed. **Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID. • **[V6.2]** New message: HCP048E.
SET RDEVICE Tape Units	**INCOMPATIBLE:** • **[V5.1]** Operands removed: 3420 and the associated MODEL and DUAL_DENSITY **Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID. • **[V6.2]** New message: HCP048E.
SET RDEVICE Terminals	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID. • **[V6.2]** New message: HCP048E.

Table 23. Changes to CP Commands (continued)

Command	Changes
SET RDEVICE Unsupported Devices	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID. • **[V6.2]** New message: HCP048E. • **[V6.2]** Descriptions updated for DPS and RESERVE_RELEASE operands.
SET RDEVICE 3800 Printers	**Upwardly compatible:** • **[V6.2]** New operands: EQID, NOEQID. • **[V6.2]** New message: HCP048E.
SET REORDER	**INCOMPATIBLE:** • **[V6.3]** The reorder process is no longer supported. Reorder is always OFF. The command is retained for compatibility, but issuing SET REORDER ON results in message HCP6005E.
SET RESERVED	**INCOMPATIBLE:** • **[V6.3]** Responses have new formats and include new information. • **[V6.3]** Messages removed: HCP026E, HCP2266E, HCP2267E. **Upwardly Compatible:** • **[V6.3]** New operands and associated responses: DCSS, NSS, SPOOLID, SYSMAX, USER, *storsize*. • **[V6.3]** New guidelines for defining reserved storage.
SET SECUSER	**Upwardly compatible:** • **[V4.2]** New message: HCP017E. • **[V6.2]** New messages: HCP1149E, HCP6768I.
SET SHARE	**INCOMPATIBLE:** • **[V5.4]** Response includes the virtual CPU types. **Upwardly compatible:** • **[V4.2]** New message: HCP017E. • **[V5.4]** New CPU type: ICF. • **[V6.1]** LIMITHARD description updated. • **[V6.3 VM65680]** When multithreading is enabled, prorated core time is used in the consumption limiting calculation for a virtual machine being consumption limited. When multithreading is not enabled, raw CPU time is used instead.

Table 23. Changes to CP Commands (continued)

Command	Changes
SET SRM	**INCOMPATIBLE:** • **[V6.2]** LIMITHARD default changed from DEADLINE to CONSUMPTION. • **[V6.4]** XSTORE operand and associated response removed. Expanded storage (XSTORE) is not supported. • **[V6.4]** New message: HCP1401I. **Upwardly compatible:** • **[V6.1]** New operand and associated response: LIMITHARD. • **[V6.2]** New default STORBUF percentage values. • **[V6.3]** New operands and associated responses: CPUPAD, DSPWDMETHOD, EXCESSUSE, POLARIZATION. • **[V6.3 VM65418]** New message: HCP1054E. • **[V6.3 VM65586, VM65696]** New message variations: HCP1050E. • **[V6.4]** DSPBUF, LDUBUF, and STORBUF functions now less meaningful because virtual processor management improvements mean no users stay in the eligible list more than an instant.
SET SSI	**Upwardly compatible:** • **[V6.3]** New operand: FORCE. • **[V6.3]** Updated message: HCP1664E. • **[V6.3 VM65712]** New operand: PDRVOLUME. • **[V6.3 VM65712]** New messages: HCP1630I, HCP1631I, HCP1640E.
SET STORAGE	**Upwardly compatible:** • **[V5.4]** New message: HCP1134I. • **[V6.2]** New message: HCP1133I.
SET TIMEZONE	**Upwardly compatible:** • **[V6.1]** New message: HCP0987E.
SET VCONFIG	**Upwardly compatible:** • **[V6.4]** New operand and associated response: GENERAL.
SET VMLAN	**INCOMPATIBLE:** • **[V6.4]** Removed operand: ENSEMBLE. • **[V6.4]** Messages removed: HCP3013E, HCP3014E, HCP3015E, HCP3026E, HCP3027E, HCP3028E, HCP3029E, HCP3030E, HCP3035E, HCP3038E. **Upwardly compatible:** • **[V6.1]** New operands: MACPROTECT and ENSEMBLE. New messages: HCP3013E, HCP3014E, HCP3015E. • **[V6.2]** New messages: HCP3027E, HCP3028E, HCP3029E, HCP3030E, HCP3035E, HCP3038E, HCP6645E. • **[V6.2]** MACPROTECT description updated. • **[V6.3 VM65583]** New operand: TRACESIZE. • **[V6.3 VM65583]** New message: HCP3195E. • **[V6.4 VM65925]** New operand: DNA DISABLE | ENABLE.

Table 23. Changes to CP Commands (continued)

Command	Changes
SET VSWITCH	**INCOMPATIBLE:** • **[V5.1]** 'ANY' is not supported as a value for *vlanid*. It was formerly the default. The default VLAN for the user is now the VLAN ID specified on the DEFINE VSWITCH command or configuration statement. • **[V5.4]** New format of the RDEV operand. • **[V5.4]** 'DROP' and 'FORWARD' are not supported as a value for ISOLATION. The valid options are 'ON' or 'OFF'. • **[V6.2]** Message HCP2798I removed. • **[V6.4]** Removed operand: OSDSIM. • **[V6.4]** Removed messages: HCP3020E, HCP3021E, HCP3022E (some variations). **Upwardly compatible:** • **[V5.1]** New operand: PORTTTYP. • **[V5.2]** New operands: GVRP, NOGVR, PPROMISCUOUS, NOPROMISCUOUS. • **[V5.2]** Updated message: HCP2846E. • **[V5.3]** New operands: GROUP | NOGROUP, MACID, RESTORE. • **[V5.3]** New message: HCP2850E. • **[V5.3]** Updated messages: HCP2783E, HCP2846E, HCP2858I. • **[V5.3]** New operand: ISOLATION. New message: HCP3000E. • **[V5.4]** New operand: VLAN_counters. Updated message: HCP2846E. New usage notes on VLAN_counters. • **[V5.4]** Updated operand description: RDEV. • **[V5.4]** New messages: HCP3002E, HCP3003E, HCP3004E. • **[V6.1]** New operands: OSDSIM, UPLINK, NIC, MACPROTECT. • **[V6.1]** New messages: HCP3021E, HCP3022E. Updated messages: HCP2832E, HCP2838I, HCP2850E. • **[V6.2]** New operands: PORTNUMBER and VLANID. • **[V6.2]** Updated messages: HCP2846E and HCP3022E. New messages: HCP3031I, HCP3032I, HCP3033E, HCP3036E, HCP3037E, HCP3039I, and HCP3040I. • **[V6.2]** New operands: BRIDGEPORT, PATHMTUDISCOVERY. • **[V6.2]** New messages: HCP2838I, HCP3022E, HCP3045E, HCP3049E, HCP6706E. • **[V6.3]** New operands: SWITCHOVER, VEPA. • **[V6.3]** Updated message: HCP2838I. • **[V6.3]** New message variations: HCP3022E. • **[V6.3 VM65583]** New operand: TRACESIZE. • **[V6.3 VM65583]** New option: IVLPORT with operands: VLAN, PING, HEARTBEAT, RESET. • **[V6.3 VM65583]** New messages: HCP3160E, HCP3161E, HCP3163E, HCP3166E, HCP3167E, HCP3170E, HCP3173E, HCP3174I, HCP3186E, HCP3195E, HCP3196E, HCP3199E, HCP3221E. Updated messages: HCP2846E, HCP3022E. • **[V6.4]** New option COUNTERS CLEAR with operands UPLINK, BRIDGEport, ALL, DATA, DISCARDed, ERRORs

Table 23. Changes to CP Commands (continued)

Command	Changes
SET VSWITCH (continued)	**Upwardly compatible:** • **[V6.4 VM65925]** Updated operands: GRANT, REVOKE, PORTNUMBER, VLANID and Uplink NIC. • **[V6.4 VM65925]** Updated Usage Notes and Examples. • **[V6.4 VM65925]** New message: HCP3046E. • **[V6.4 VM65925]** Updated message: HCP3022E. • **[V6.4 VM65925]** Deleted messages: HCP3036E and HCP3037E.
SET WRKALLEG	**Upwardly compatible:** • **[V6.2]** New message: HCP6645E.
SET 370ACCOM	**Upwardly compatible:** • **[V6.4]** If MACHINE Z guest issues SET 370ACCOM ON, the command is rejected with error message HCP2632E.
SHUTDOWN	**INCOMPATIBLE:** • **[V5.1]** ORIGIN operand is accepted when loading z/VM V5.1 or later CP, but CP will relocate itself to location X'2000'. • **[V5.1]** Messages removed: HCP6880E, HCP9403E, HCP9405E, HCP9415I, HCP9417I, HCP9418E. **Upwardly compatible:** • **[V3.1]** New operand: WAIT. • **[V3.1]** If REIPL is specified, an automatic warmstart is done from a CP module, not a nucleus. • **[V4.3]** New operands to cancel a scheduled shutdown or send a shutdown signal to users. • **[V5.4]** New SYSTEM operand to help prevent inadvertently shutting down the wrong system. • **[V5.4]** New messages: HCP6021E, HCP6022E. • **[V6.2]** New message: HCP6028E. • **[V6.2]** New message: HCP959I. • **[V6.2]** Updated messages: HCP961W and HCP9277I. • **[V6.3]** New operand: NOCKPT. • **[V6.4]** New usage notes. New messages: HCP2116I, HCP2117I.
SIGNAL	**Upwardly compatible:** • **[V6.4]** New usage note. New messages: HCP2118I.
SNAPDUMP	**Upwardly compatible:** • **[V5.2]** Changed message: HCP9267E. • **[V6.4 VM65989]** New operand: PGMBKS. Updated usage note. New messages: HCP002E, HCP026E.
SPTAPE	**INCOMPATIBLE:** • **[V5.1]** Command removed. SPTAPE is not supported. Tapes created with SPTAPE on an earlier VM system cannot be restored to a z/VM V5.1 or later system.
SPXTAPE (in general)	**Upwardly compatible:** • **[V4.3]** New message: HCP1917E.
START	**Upwardly compatible:** • **[V6.4 VM66029]** Renamed help files for subcommands. See "[V6.4 APAR] z/VM HELP Quality Improvement" on page 238.

Table 23. Changes to CP Commands (continued)

Command	Changes
STORE (in general)	**INCOMPATIBLE:** • **[V5.2]** Message changed: HCP645E. • **[V5.2]** Messages removed: HCP749E, HCP785E, HCP1701E. • **[V5.2]** Operand removed: VECTOR. • **[V6.1]** CDX operand has no function; documentation removed. • **[V6.2]** Message removed: HCP645E. • Also see specific STORE commands listed below. **Upwardly compatible:** • **[V3.1]** Message HCP6150E might specify z/Architecture mode. • **[V3.1]** New operands: PSWA, PSWG. • **[V6.1]** Message removed: HCP1702E. • Also see specific STORE commands listed below.
STORE CDX	**INCOMPATIBLE:** • **[V6.1]** Command has no function; documentation removed.
STORE ESA/XC Storage	See STORE (Guest Storage - ESA/XC).
STORE (Guest Storage - ESA/390)	**INCOMPATIBLE:** • **[V5.2]** Messages removed: HCP645E, HCP749E, HCP785E, HCP1701E. **Upwardly compatible:** • **[V3.1]** New operands for indirect addressing and storing data in ASCII format. • **[V3.1]** Supports 64-bit addressing. • **[V3.1]** New z/Architecture version of this command. • **[V5.2]** Default is zero for operand *hexloc*.
STORE (Guest Storage - ESA/XC)	**Upwardly compatible:** • **[V3.1]** New operands for indirect addressing and storing data in ASCII format. • **[V3.1]** Supports 64-bit addressing. • **[V5.2]** Default is zero for operand *hexloc*.
STORE (Guest Storage - z/Architecture)	**Upwardly compatible:** • **[V5.2]** Default is zero for operand *hexloc*.
STORE (Host Storage)	**INCOMPATIBLE:** • **[V3.1]** New operands for indirect addressing and storing data in ASCII format. • **[V3.1]** Supports 64-bit addressing. • **[V5.2]** New required operand, L or R, to specify host logical storage or host real storage. • **[V5.2]** New messages: HCP160E, HCP164E, HCP6065E. **Upwardly compatible:** • **[V3.1]** New operands for indirect addressing and storing data in ASCII format. • **[V3.1]** Supports 64-bit addressing. • **[V5.2]** Default is zero for operand *hexloc*.

Table 23. Changes to CP Commands (continued)

Command	Changes
STORE PSW	**INCOMPATIBLE:** • **[V4.1]** Operands removed: CAW, CSW. 370 virtual machines are not supported. **Upwardly compatible:** • **[V3.1]** For a z/Architecture guest, translates an ESA/390 format PSW into a z/Architecture format PSW (if the SET PSWTRANS ALL or SET PSWTRANS STORE command was previously issued).
STORE (Registers)	**INCOMPATIBLE:** • **[V5.1]** Text of message HCP6150E has changed. **Upwardly compatible:** • **[V2.4]** New operands: FPC *hexword*. • **[V2.4]** Yreg operands accept register numbers 0-15 if the IEEE Floating Point hardware feature is installed on the processor. • **[V2.4]** New messages: HCP6153E, HCP6154E. • **[V3.1]** New operands: GG*reg1*, GH*reg1*, XG*reg1*, XH*reg1*. • **[V3.1]** Message HCP6150E might specify z/Architecture mode. • **[V6.3 VM65733]** Updated operand: Y*reg1*. The Vector Facility for z/Architecture support enables access to the additional floating-point (AFP) registers without requiring prior guest program use. • **[V6.3 VM65733]** Messages removed: HCP6153E, HCP6154E.
STORE STATUS	**Upwardly compatible:** • **[V2.4]** Stores the address of the extended save area at address 212 (X'D4'). This save area contains floating-point registers 0-15 and the floating-point control register. • **[V3.1]** Stores virtual machine data for a z/Architecture virtual machine. • **[V6.3 VM65733]** Updated note. The extended save area address is used if it is provided.
STORE VECTOR	**INCOMPATIBLE:** • **[V5.2]** Command removed. ESA/390 Vector Facility is not supported.
SYNCMDRS	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP2601E. I/O assist is not supported.
SYSTEM	**Upwardly compatible:** • **[V3.1]** Stores virtual machine data for a z/Architecture virtual machine. • **[V5.3]** SYSTEM RESTART is modified to prevent starting secondary CPUs that are not operational. • **[V5.3]** New message variation: HCP1455E. • **[V5.4]** New option: DETACH.
TERMINAL (in general)	**Upwardly compatible:** • **[V4.4]** New operands: SYS3270, SYSGRAF

Table 23. Changes to CP Commands (continued)

Command	Changes
TRACE (in general)	**INCOMPATIBLE:** • **[V2.2]** New message: HCP1038E. • **[V5.1]** Messages removed: HCP1038E, HCP1039E. • **[V5.2]** Response for Machine Check Codes has changed: codes VF and VS are unused. ESA/390 Vector Facility is not supported. • **[V5.4]** All *mnem* fields in TRACE command responses are now 7 characters wide, in which the mnemonic value is left-justified. **Upwardly compatible:** • **[V2.3]** New command responses. • **[V2.4]** New response for floating-point register information. • **[V3.1]** Message HCP6150E might specify z/Architecture mode. • **[V3.1]** New operand: GG. • **[V3.1]** New option: ASCE. • **[V3.1]** New response formats for 64-bit. • **[V4.3]** New operand: AIF. • **[V5.2]** New message: HCP1040E. • **[V6.4]** New TXSUSPEND and NOTXSUSPEND operands. • **[V6.4]** New messages: HCP2106I, HCP2107E. • Also see specific TRACE commands listed below.
TRACE *mnemonic1*	**Upwardly compatible:** • **[V2.3]** New mnemonic: BSA. • **[V3.1]** New mnemonics: EPSW, LCTLG, LPSWE, STCTG, STFL, STSI. • **[V5.2]** New mnemonic: STFLE. • **[V6.4]** New mnemonic: PFMF.
TRACE IO	**Upwardly compatible:** • **[V5.3]** New command responses for MIDAWs. • **[V6.2]** The default IODATA for each CCW is 64. • **[V6.2]** New and changed responses.
TRACE TABLE	**Upwardly compatible:** • **[V3.1]** New response formats for 64-bit. • **[V5.2]** The contents of the breaking-event-address register (BEAR) are displayed after the traceback table if the PER-3 facility is available.
TRSAVE	**Upwardly compatible:** • **[V4.3]** New message: HCP1917E. • **[V6.2]** If issued with the AT command, the TO *userid* operands are required.
TRSOURCE	**Upwardly compatible:** • **[V5.3]** New message variation: HCP371I. • **[V6.3 VM65417]** New message variation: HCP371I.

Table 23. Changes to CP Commands (continued)

Command	Changes
TRSOURCE ID	**INCOMPATIBLE:** • **[V6.2]** Default value for IODATA changed from 0 bytes to 64 bytes. **Upwardly compatible:** • **[V3.1]** New values can be specified on the DL operand for 64-bit registers or indirect addressing. • **[V5.2]** New LAN operand and new options. • **[V6.3 VM65417]** New trace type: PCI.
UNCOUPLE	**Upwardly compatible:** • **[V4.4]** Updated purpose to include virtual switch information and added a new response.
UNDEDICATE	**INCOMPATIBLE:** • **[V5.1]** Default of V=R user removed. USER operand must be specified. V=R virtual machines are not supported. • **[V5.1]** Message removed: HCP893E. • **[V5.3]** Text for message HCP898I has changed.
UNLOCK	**INCOMPATIBLE:** • **[V3.1]** RIO370 operand no longer supported. • **[V5.1]** SYSTEM operand removed. The CP nucleus is no longer pageable. Modules in the nucleus cannot be locked or unlocked. • **[V5.1]** SYMBOL and *symbol* operands have no effect. All symbols are now considered resident (locked) and cannot be unlocked. • **[V5.1]** VIRT=REAL and V=R operands removed. The V=R area is not supported. • **[V5.1]** Text of message HCP295E has changed. • **[V5.1]** Messages removed: HCP200E, HCP204E, HCP205I, HCP1556I. • **[V5.2]** New required operand, LOGICAL or REAL, to specify whether guest storage is to be unlocked from host logical storage or host real storage. • **[V5.2]** Message removed: HCP295E.
VARY (in general)	**INCOMPATIBLE:** • See specific VARY commands listed below. **Upwardly compatible:** • **[V6.3 VM65417]** New operand: PCIFUNCTION. • **[V6.3 VM65586, VM65696]** New CORE operand for adding or removing a core when multithreading is enabled. • **[V6.4 VM66029]** Renamed help files for subcommands. See "[V6.4 APAR] z/VM HELP Quality Improvement" on page 238. • Also see specific VARY commands listed below.
VARY PATH	**Upwardly compatible:** • **[V6.2]** New response to indicate incompatible paths (due to incompatible zHPF feature) cannot be varied online logically without specifying the FORCE option. • **[V6.3]** Real device number description updated for a 5-digit device value in the alternate subchannel set.
VARY PCIFUNCTION	**Upwardly compatible:** • **[V6.3 VM65577]** New message: HCP1167E

Table 23. Changes to CP Commands (continued)

Command	Changes
VARY PROCESSOR	**INCOMPATIBLE:** • **[V5.2]** Responses for Vector Facility availability have been removed. ESA/390 Vector Facility is not supported. • **[V5.3]** Text for messages HCP892E and HCP896E has changed. • **[V6.3 VM65586, VM65696]** This command is not valid when multithreading is enabled; use VARY CORE instead. • **[V6.3 VM65586, VM65696]** New message: HCP1321E. **Upwardly compatible:** • **[V5.3]** This command is changed to prevent the last active primary processor from being varied offline. • **[V5.3]** New message: HCP884E.
VARY (Real Device)	**Upwardly compatible:** • **[V3.1]** New message: HCP6861I. • **[V4.3]** New message: HCP1917E. • **[V5.3]** New message: HCP8701I. • **[V6.1]** New message: HCP6294I. • **[V6.2]** New message: HCP6627I. • **[V6.3]** Real device number description updated for a 5-digit device value in the alternate subchannel set. • **[V6.3]** New message: HCP6295I. • **[V6.4]** Removed message: HCP6294I. • **[V6.4]** New message variation: HCP8701I.
VARY CORE	**Upwardly compatible:** • **[V6.4]** Updated Purpose to include information on multithreading.
VARY SUBCHANNEL	**Upwardly compatible:** • **[V6.3]** Real device number description updated for a 5-digit device value in the alternate subchannel set.
VARY VECTOR	**INCOMPATIBLE:** • **[V5.2]** Command removed. ESA/390 Vector Facility is not supported.
VMDUMP	**INCOMPATIBLE:** • **[V5.2]** VMDUMP dumps of z/Architecture guests are not compatible with previous releases. • **[V5.2]** Messages removed: HCP745I, HCP746E. **Upwardly compatible:** • **[V3.1]** Can dump larger storage areas and discontiguous storage areas. • **[V5.2]** Maximum address for ESA/390 mode virtual machines and z/Architecture mode virtual machines has changed. • **[V5.2]** Supports the SET UNDERSCORE command setting to control the separation of 16-digit output fields. • **[V5.2]** New message: HCP162I. • **[V5.2]** Changed messages: HCP009E, HCP161E, HCP6186I and HCP6191I. • **[V5.2]** Removed message: HCP162E. • **[V6.2]** Execution can be halted with the CPHX or FORCE command. • **[V6.2]** New message: HCP6790I. • **[V6.3 VM65733]** New usage note: The VMDUMP command does not dump vector registers.

Table 23. Changes to CP Commands (continued)

Command	Changes
VMRELOCATE	**Upwardly compatible:** • **[V6.3]** New message HCP3150I. New message variations for HCP1982I. • **[V6.3 VM65417]** New message: HCP1161I. • **[V6.3 VM65418]** New message: HCP1961I. • **[V6.3 VM65733]** New message: HCP1987I. • **[V6.4]** Deleted messages: HCP1962I, HCP3150I. • **[V6.4 VM65925]** New message variation: HCP1982I.
WARNING	**Upwardly compatible:** • **[V6.2]** AT option not allowed when issued via the AT command. • **[V6.2]** New message: HCP1142E.
XAUTOLOG	**INCOMPATIBLE:** • **[V4.1]** 370 operand no longer supported. • **[V5.1]** Messages removed: HCP811I, HCP1508I, HCP1511I, HCP1556I. **Upwardly compatible:** • **[V3.1]** Can specify larger amount of storage on the STORAGE operand. • **[V6.4]** New MACHINE mode: Z. • **[V6.4]** New message variation: HCP1512E. • **[V6.4 VM65982]** Updated usage note. • **[V6.4 VM65942]** An ESA or XA virtual machine may be put into ESA/390-compatibility mode instead of full ESA/390 mode.
XSPOOL QUERY	**INCOMPATIBLE:** • **[V6.3]** Removed. CSE is no longer supported.
XSPOOL QUERY SHARE	**INCOMPATIBLE:** • **[V6.3]** Removed. CSE is no longer supported.
XSPOOL REMOVE	**INCOMPATIBLE:** • **[V6.3]** Removed. CSE is no longer supported.
XSPOOL SYNC	**INCOMPATIBLE:** • **[V6.3]** Removed. CSE is no longer supported.
XSPOOL UNLOCK	**INCOMPATIBLE:** • **[V6.3]** Removed. CSE is no longer supported.
XSPOOL XLIST	**INCOMPATIBLE:** • **[V6.3]** Removed. CSE is no longer supported.
XSPOOL XLISTADD	**INCOMPATIBLE:** • **[V6.3]** Removed. CSE is no longer supported.
XSPOOL XLISTDEL	**INCOMPATIBLE:** • **[V6.3]** Removed. CSE is no longer supported.

CP Utilities

Table 24 on page 350 lists CP utilities that have changed. For additional information see *z/VM: CP Commands and Utilities Reference*.

Table 24. Changes to CP Utilities

Utility	Changes
CPFMTXA	**INCOMPATIBLE:** • **[V5.1]** Message removed: HCP6203E. • **[V5.1]** Fields in the FBA allocation map have been expanded to accommodate larger disk sizes. • **[V6.2]** CPFMTXA EXEC moved from MAINT 190 to PMAINT 551. **Upwardly compatible:** • **[V5.1]** Allocation specifications for FBA devices, formerly limited to 6 digits, now accept up to 9 digits. • **[V5.1]** TDSK, PERM, and PARM allocations may reside above page 16,777,215. DRCT, PAGE, and SPOL allocations may not reside above page 16,777,215. • **[V6.2]** New messages: HCP6628E, HCP6632E, HCP6633E, HCP6639E. • **[V6.2]** New OWNER operand and options.
CPSYNTAX	**Upwardly compatible:** • **[V5.3]** Purpose expanded to include limitations. • **[V6.2]** New LPAR operand and enhancements added to support the new system configuration file statements SSI, BEGIN, and END. • **[V6.4]** New message: HCP2149E. • **[V6.4]** New message variations: HCP1001E, HCP2768E, HCP6706E.
DDR	**INCOMPATIBLE:** • **[V5.1]** NUCLEUS function statement removed. The CP nucleus cannot be dumped to or restored from tape. • **[V5.1]** Messages removed: HCP722E, HCP723E, HCP724E. • **[V6.2]** Previously, a response or message was issued when the input device for a COPY or RESTORE function was larger than the output device, even if the data being copied or restored was a subset. Now, a response or message is issued when the input device is larger than the output device only if the ALL operand was specified on the COPY or RESTORE function statement. **Upwardly compatible:** • **[V4.2]** New operand to compress DDR data: LZCOMPACT. • **[V6.4]** New message: HCP9037W. • **[V6.4]** New MODE option values: EFMT1, EFMT2, EFMT3 and EFMT4.

Table 24. Changes to CP Utilities (continued)

Utility	Changes
DIRECTXA	**INCOMPATIBLE:** • **[V2.3]** If a USER statement within the directory specifies a logical line edit symbol that is not valid (a letter A-Z, number 0-9, or bytes X'OE' (shift out) or X'0F' (shift in)), DIRECTXA issues message HCP786I, uses the system default line edit symbol, and continues processing. If no error prevents the directory from being written, DIRECTXA returns to CMS with RC=9. • **[V5.1]** A request for a V=R or V=F virtual machine in the directory file being processed will result in the user being logged on as a V=V virtual machine. V=R and V=F virtual machines are no longer supported. • **[V6.2]** DIRECTXA MODULE moved from MAINT 190 to PMAINT 551. • **[V6.3 VM65577]** New version of DIRECTXA requires more directory space on the DASD volume that contains the object directory. An additional 60 bytes per user plus 2 pages for system use is required. The count of users should include each user defined with a USER, IDENTITY, or POOL statement. **Upwardly compatible:** • **[V3.1]** New operands: &SYSRES *volid*. • **[V3.1]** New message: HCP493E. • **[V5.1]** Supports larger FBA disk size in directory MDISK statements. • **[V6.2]** SYSAFNID is not allowed if the SSI option is specified on the DIRECTORY statement in the user directory being processed. • **[V6.2]** New messages: HCP495E, HCP496E, HCP497E, HCP498E, HCP499E, HCP664E, HCP665E, HCP666E, HCP667E, HCP668E, HCP669I, HCP670E, HCP671E, HCP672E, HCP675E, HCP676E, HCP796E, HCP797E, and HCP1799E. • **[V6.2]** Changed messages: HCP481E, HCP751E, HCP752E, HCP753E, HCP754E, HCP755E, HCP758E, HCP772E, HCP773E, HCP776E, HCP777E, HCP783I, HCP786I, HCP789E, HCP790E, HCP791E, HCP792E, HCP795E, HCP798E, HCP1755E, HCP1757E, HCP1758E, HCP1759E, HCP1760E, HCP1762E, HCP1764E, HCP1766E, HCP1770E, HCP1771E, HCP1772E, HCP1776E, HCP1785E, HCP1796E, HCP1798E, HCP1879E, HCP1880E, HCP1882E, HCP2826E, and HCP2829E. • **[V6.2]** Deleted messages: HCP781E, HCP1761E, and HCP1881E. • **[V6.2]** Return code 4 is issued if an attempt is made to update the real user directory with an SSI-enabled or not-SSI-enabled source directory that is not valid in the current system configuration. • **[V6.2]** Updated formula for computing storage for a directory. • **[V6.4]** Response contains changed value for the new release level. • **[V6.4 APAR VM65925]** Updated message: HCP1882E.

Table 24. Changes to CP Utilities (continued)

Utility	Changes
DISKMAP	**INCOMPATIBLE:** • **[V6.2]** DISKMAP EXEC moved from MAINT 190 to PMAINT 551. • **[V6.2]** Output file record length has been changed from 80 characters to 100 characters to accommodate columns for new SUBCONFIG and MEMBER information. **Upwardly compatible:** • **[V6.2]** From an authorization perspective, the USERID value specifies the user ID to which the MDISK belongs. For multiconfiguration virtual machines, the MDISKs can be defined within either the IDENTITY or SUBCONFIG stanzas. In both cases the USERID value specifies the name from the IDENTITY statement. If the SUBCONFIG statement is not referenced by any BUILD statement, the USERID value contains a value of ***NOUSER!**. • **[V6.2]** The SUBCONFIG value in the output is blank unless the MDISK was defined within a SUBCONFIG stanza. • **[V6.2]** The MEMBER value in the output is blank unless the MDISK was defined within a SUBCONFIG stanza. The value displayed is the system ID from the BUILD statement. If there is no corresponding BUILD statement, this value is blank.
DUMPLD2	**Upwardly compatible:** • **[V6.1]** New message: HCP8247E. • **[V6.3]** New operand: DASD.
DUMPLOAD	**Upwardly compatible:** • **[V5.1]** New message: HCP8185I. • **[V5.2]** Can process new dump formats. • **[V6.2]** PRINT/NOPRINT operand description updated.
FORMSSI	**Upwardly compatible:** • **[V6.3 VM65712]** Updated messages: HCP6609E, HCP6619I. • **[V6.3 VM65712]** New message: HCP6620E.
HCPLDR	**INCOMPATIBLE:** • **[V5.1]** Operands removed: 370, ESA, XA, XC. All modules created with HCPLDR will run in ESA, XA, and XC virtual machines. 370 virtual machines are not supported. • **[V5.1]** RLDSAVE operand is accepted when generating the CP module but has no effect. The CP nucleus is no longer relocatable. CP will always relocate itself to location X'2000'. **Upwardly compatible:** • **[V3.1]** New operands: &SYSRES *volid*. • **[V3.1]** New message: HCP493E. • **[V5.1]** Output report "CSECT'S WITH SIZE GREATER THAN CONDITIONAL PAGE BOUNDARY" is generated when the MAP option is specified but the PAGEB option is not specified. For CP, only @MAPSTART should be listed. Any other modules listed are in error. • **[V6.4]** Response contains changed value for the new release level.

Table 24. Changes to CP Utilities (continued)

Utility	Changes
HCPSADMP	**INCOMPATIBLE:** • **[V6.2]** Support to IPL from device types 3350, 3375, 9345, 3424 and 9348 has been removed. Supported IPL device types are 3380 and 3390 DASD and 3480, 3490, and 3590 tape. Also, support to dump to device types 3424 and 9348 has been removed. Supported dump device types are 3480, 3490, and 3590 tape. **Upwardly compatible:** • **[V6.4]** HCPSADMP cannot be run in a Z mode virtual machine. • **[V6.4]** New message: HCP8208E.
INSTALL	**INCOMPATIBLE:** • **[V6.2]** Renamed to INSTTAPE. TAPE operand removed.
INSTDEF	**INCOMPATIBLE:** • **[V5.1]** Utility removed.
IOCP	**Upwardly compatible:** • **[V5.3]** Default DYN or NODYN option set according to processor type.
MONWRITE	**Upwardly compatible:** • **[V6.1]** Added usage note about CLOSE restriction.
MOVE2SFS	**Upwardly compatible:** • **[V6.2]** This utility requires an operand and has no default.
OVERRIDE	**INCOMPATIBLE:** • **[V6.3]** This utility has been removed. UCR spool files and the OVERRIDE utility are no longer supported. Use the MODIFY statement or command to change command privilege classes.

Table 24. Changes to CP Utilities (continued)

Utility	Changes
SALIPL	**INCOMPATIBLE:** • **[V2.2]** Message HCP039E has been deleted and replaced by new message HCP394E with the same text. • **[V2.2]** Date field on the file list panel displays the year with 4 digits. • **[V5.1]** ORIGIN option is accepted when loading CP for z/VM V5.1 or later, and the default address of X'1000' still applies, but CP will relocate itself to location X'2000'. • **[V5.1]** IPLPARMS value ARCH390 is ignored when IPLing a z/VM V5.1 or later CP module because there is no 32-bit version. • **[V5.1]** The number of blocks on FBA DASD (SCSI or not) that SALIPL uses for SAPL has increased. SALIPL now writes to blocks 5-207 on CP-formatted FBA DASD. In previous releases, SALIPL wrote to blocks 5-31. You must ensure that no other functions (such as CP directory, warmstart, checkpoint, paging, spooling, user minidisks, or CP parm disks) are using the area to which SALIPL writes. Use ICKDSF or CPFMTXA to allocate pages 4 through 25 as PERM. • **[V5.1]** On a CMS minidisk on FBA DASD, SALIPL now writes to blocks 5-207 of the RECOMP area. In previous releases, SALIPL used blocks 5-31. To accommodate the larger size of SAPL, you might need to increase the size of the RECOMP area on the CMS minidisk. **Upwardly compatible:** • **[V6.3]** New DEVICE option which can have a 5-digit device address. • **[V6.3]** New message: HCP021E. • **[V6.3 VM65577]** New message: HCP9061W. • **[V6.4]** Default ORIGIN in z/Architecture-only environment is X'2000'. • **[V6.4]** New message: HCP9037W.
SDINST	**INCOMPATIBLE:** • **[V6.4]** The user running the SDINST utility must have at least 256 MB of storage, not 128 MB. **Upwardly compatible:** • **[V6.4]** Renamed image files for the new release: SADU64, SSPK64, SSPI64, SSPP64. • **[V6.4]** New message: HCP8639E.
TRACERED	**Upwardly compatible:** • **[V5.2]** New output for LAN trace. • **[V5.2]** TRACERED selects either CP or guest trace entries by time. If the timestamp of an entry is identical to that of the previous entry, as can occur when the Store-Clock-Fast Facility is installed and enabled during tracing, an '=' sign is displayed immediately following the timestamp, signifying that the order of occurrence of those entries is unpredictable. • **[V5.3]** Output for I/O trace updated to include MIDAWs.

Dynamic I/O Return Codes

Table 25 on page 355 lists dynamic I/O return codes that have changed. For additional information, see *z/VM: I/O Configuration*.

Table 25. Changes to Dynamic I/O Return Codes

Return Code	Changes
0103	**Upwardly compatible:** • **[V5.1]** Issued if you tried to modify or delete an I/O control unit, but the control unit number specified is associated with a coupling-facility control unit. • **[V5.1]** Issued if you tried to modify or delete a coupling-facility control unit but the control unit number specified is associated with an I/O control unit.
0105	**Upwardly compatible:** • **[V4.2]** Issued if the specified channel path for a FICON_CTC control unit is not a FICON channel. • **[V4.4]** Issued if you tried to define or modify the I/O or coupling-facility control unit with a mixture of shared and unshared channel paths (including channel path that might be defined in other channel-subsystem images).
0106	**Upwardly compatible:** • **[V6.3 VM65417]** Issued if you tried to add a PCI function, but you specified one that is already defined in the I/O configuration.
0107	**Upwardly compatible:** • **[V4.2]** Issued if more than one channel path is specified for a FICON_CTC control unit. • **[V6.3 VM65417]** Issued if one of the following occurred: – You tried to define or change a PCI function but one or more of the logical partitions that you specified is not currently defined in the active I/O configuration. – You tried to add a PCI function or change an existing one by adding access to one or more logical partitions, and one or more logical partitions were defined in the initial access list that were not defined in the access list. – You tried to change an existing PCI function by deleting one or more logical partitions from the access list, but one or more logical partitions do not have access to the PCI function. – You tried to add a nonshared PCI function or change an existing nonshared PCI function by adding access to one or more logical partitions, and more than one logical partition was specified as having initial access.
010A	**Upwardly compatible:** • **[V6.3 VM65417]** Issued if you tried to change or delete an existing PCI function, but the PCI function you specified is not defined in the I/O configuration.

Table 25. Changes to Dynamic I/O Return Codes (continued)

Return Code	Changes
010D	**Upwardly compatible:** • **[V4.4]** Issued if you tried to add, delete, or change a control unit, but one of the following occurred: – For a channel path to which the specified I/O control unit is attached, you exceeded the limit on the number of unit addresses supported for that type of channel path. – One or more of the specified channel paths for the specified I/O control unit is a coupling-facility channel path. – One or more of the specified channel paths for the specified coupling-facility control unit is one of several types of coupling-facility-receiver channel paths that is not dynamically changeable on your machine model. – Your machine model does not support the dynamic-I/O configuration change that you requested. – One or more of the specified channel paths for the specified coupling-facility control unit is an I/O channel path. – The combination of coupling-facility channel-path types that would result from the requested change for the specified coupling-facility control unit (including channel-path types defined in other channel-subsystem images for the control unit) is not supported on your machine model. • **[V6.3 VM65417]** Issued if your machine model does not support the dynamic I/O configuration change that you requested.
010F	**Upwardly compatible:** • **[V6.3 VM65417]** Issued if you tried to delete a PCI function or you tried to change a PCI function by deleting one or more logical partition's access to it, and the PCI function is currently configured to one or more of those logical partitions. The system's attempt to deconfigure the PCI function from one or more logical partitions was unsuccessful.
0110	**Upwardly compatible:** • **[V6.3 VM65417]** Issued if one of the following occurred: – The machine could not get access to certain resources required to make the requested configuration change. – The processor was temporarily unable to make the requested configuration change. – You issued a MODIFY PCIFUNCTION command to request conditional deletion of access by one or more logical partitions from the specified PCI function. However, the PCI function is currently configured to one or more of those logical partitions.
0112	**Upwardly compatible:** • **[V6.3 VM65417]** Issued if one of the following occurred: – The specified PCI function, if added to your configuration, would exceed the maximum number of PCI functions that can be configured to your machine. – The value of the specified PCI function ID exceeds the maximum value supported by your machine. – The value of the specified virtual function number (VFN) would exceed the maximum value supported by your machine.

Table 25. Changes to Dynamic I/O Return Codes (continued)

Return Code	Changes
0113	**Upwardly compatible:** • **[V4.4]** Issued if you tried to delete a coupling-facility control unit from your configuration, but the the control unit is defined in more than one channel-subsystem image.
011C	**Upwardly compatible:** • **[V6.3 VM65417]** Issued if the specified physical channel identifier (PCHID) is not a valid PCHID in your configuration • **[V6.3 VM65577]** Issued if the specified physical channel identifier (PCHID) is not a valid PCHID in your configuration or was not specified when required.
011D	**Upwardly compatible:** • **[V6.3 VM65417]** Issued if you tried to add a PCI function to your configuration, but the specified physical channel identifier (PCHID) or PCHID/Virtual Function Number (VFN) combination is already defined in your configuration.
012C	**Upwardly compatible:** • **[V6.3 VM65577]** Issued if the specified virtual channel identifier (VCHID) is not a valid VCHID in your configuration or was not specified when required.
012D	**Upwardly compatible:** • **[V6.3 VM65577]** Issued if the specified virtual channel identifier (VCHID) is a valid VCHID but has already been used in your configuration.
0320	**Upwardly compatible:** • **[V6.3 VM65577]** Issued if the specified network identifiers (NETID) are not the same value as those already defined for another virtual function which is defined on the same PCHID. • **[V6.4 VM65942]** Issued if you tried to add a PCI function and specified a PORT and a NETID, but existing PCI functions of the same type already exist for the corresponding port with a different NETID value.
0321	**Upwardly compatible:** • **[V6.3 VM65716]** Issued if the network identifiers (NETID) value was specified for an ISM PCI function and the value was already used for another ISM PCHID. • **[V6.4]** UID checking is enabled and the given UID value has already been used by another PCI function in the configuration or was not specified. • **[V6.4 VM65942]** Issued if you tried to add a PCI function of TYPE ISM with a NETID, but the NETID is already associated with another ISM PCI function in your configuration.
0322	**Upwardly compatible:** • **[V6.4 VM65942]** Issued if you tried to add a PCI function and specified a UID. UID checking is enabled for the system and the UID value failed the checking rules.
0323	**Upwardly compatible:** • **[V6.4 VM65942]** Issued if you tried to add a PCI function and specified a PORT that was invalid.

DIAGNOSE Codes

Changes to DIAGNOSE codes are identified in the following topics:

• "DIAGNOSE Codes for Customer Use" on page 358

• "Reserved DIAGNOSE Codes" on page 363

DIAGNOSE Codes for Customer Use

Table 26 lists DIAGNOSE codes for customer use that have changed. For additional information, see *z/VM: CP Programming Services*.

Table 26. Changes to DIAGNOSE Codes

Code	Changes
X'00' Storage Extended Identification Code	**Upwardly compatible:** • **[V2.2]** In the program product bit map, bit 13 (X'0004000000000000') indicates whether Year 2000 support is present in CP. • **[V3.1]** In the execution environment field of the extended-identification code bit map, bit 1 (formerly reserved) indicates the CP image that is running: 0=32-bit, 1=64-bit. • **[V6.4]** Licensed program bit map field of the extended-identification code bit map contains a changed value for the new product level. • **[V6.4]** Release information field of the extended-identification code bit map contains the new release level.
X'04' Examine Host Storage	**INCOMPATIBLE:** • **[V5.2]** Ry.2 bit setting indicates the storage address type: 0 = host logical, 1 = host real. If the loaded module is greater than 4 KB in length, the frames backing the host logical storage are not necessarily contiguous in host real storage. **Upwardly compatible:** • **[V3.1]** Supports 64-bit addressing.
X'08' Virtual Console Function	**Upwardly compatible:** • **[V4.3]** Supports 64-bit addressing.
X'10' Release Pages	**Upwardly compatible:** • **[V5.2]** Supports 64-bit addressing. New Protection exception. • **[V5.3]** Removed restriction on address range.
X'14' Input Spool File Manipulation	**Upwardly compatible:** • **[V2.2]** For subcodes X'0004', X'0008', X'0FFE', and X'0FFF', a one-byte century indicator was added to the SFBLOK data area.
X'18' Standard DASD I/O	**INCOMPATIBLE:** • **[V5.3]** This DIAGNOSE code does not support HyperPAV alias.

Table 26. Changes to DIAGNOSE Codes (continued)

Code	Changes
X'3C' Activate z/VM CP Directory	**Upwardly compatible:** • **[V6.2]** Condition code 2 means the DASD address specified in the volume label was invalid. The volume label and allocation map should be restored to the values that existed prior to the update. • **[V6.2]** Condition code 3 means a fatal I/O error occurred when z/VM attempted to read the CP directory. The volume label and allocation map should be restored to the values that existed prior to the update. • Cause of a specification exception: If Rx is unchanged: The update is not for the current online directory volume and the issuer of the DIAGNOSE X'3C' does not have the appropriate privilege class. If Rx=0: The update is for the current online directory volume and one or more of the following is true: – The issuer of the DIAGNOSE X'3C' does not have the appropriate privilege class. – The new directory is not SSI-enabled, which is not valid in the current system configuration. – The new directory is SSI-enabled, which is not valid in the current system configuration. Rx should be tested, and if it is zero, the volume label and allocation map should be restored by the function that issued the DIAGNOSE X'3C' to the values that existed prior to the update.
X'44' Voluntary Time Slice End	**Upwardly compatible:** • **[V4.3]** Supports 64-bit addressing.
X'4C' Generate Accounting Records	**Upwardly compatible:** • **[V4.3]** Supports 64-bit addressing. • **[V4.3]** New accounting record identification code: C'CC' — the virtual machine network data transmission records.
X'60' Determine Virtual Machine Storage Size	**Upwardly compatible:** • **[V5.2]** For 64-bit addressing mode, subcode X'0000000C' of Diagnose X'260' may be used to obtain the virtual machine storage size. Subcode X'00000000' of Diagnose X'260' may be used as an alternative to Diagnose X'60' for a virtual machine in 24-bit or 31-bit addressing mode.
X'64' Named Saved Segment Manipulation	**Upwardly compatible:** • **[V4.1]** For exit with error, return code X'54D' is reserved (not used). • **[V5.4]** New subcodes to support 64-bit output addresses.
X'7C' Logical Device Support Facility	**Upwardly compatible:** • **[V2.3]** For the INITIATE function, bit 3 of the first byte of Rx+1 indicates that Ry+1 contains the IP address associated with the logical device. • **[V4.4]** Bit 4 of Ry+1 points to a 16-byte field containing an IPv6 address. If bit 3 is on, bit 4 is ignored.

Table 26. Changes to DIAGNOSE Codes (continued)

Code	Changes
X'84' Directory Update-in-Place	**INCOMPATIBLE:** • **[V2.3]** For EDITCHAR operation, letters A-Z, numbers 0-9, and bytes X'OE' (shift out) and X'0F' (shift in) cannot be specified as logical line edit symbols (line-end, line-delete, character-delete, and escape). • **[V5.1]** For MACHINE operation, 370 option is not valid. 370 virtual machines are not supported. **Upwardly compatible:** • **[V2.2]** Can now replace the user's default date format setting. • **[V2.2]** New operation: DATEFMT. • **[V2.2]** New return codes due to new function (in hex): 122, 123, 124. • **[V5.1]** For OPTIONS operation, Virt=real option is ignored. V=R virtual machine is not supported. • **[V5.1]** Return code X'108' is reserved (not used). V=R and V=F virtual machines are not supported. • **[V6.2]** Updated for IDENTITY and SUBCONFIG IDs. • **[V6.2]** SHARE operation removed. • **[V6.4]** New virtual machine type for MACHINE operation: Z. • **[V6.4]** XSTORE operation is supported for compatibility only; the operation has no function. Expanded storage (XSTORE) is not supported.
X'88' Validate User Authorization/Link Minidisk	**Upwardly compatible:** • **[V5.3]** Subcodes X'08' and X'FF' were added to support password phrase validation using DMSPASS.
X'90' Read Symbol Table	**INCOMPATIBLE:** • **[V5.2]** The returned symbol address is a host logical storage address. If the CSECT at the symbol is greater than 4 KB in length, the frames backing the host logical storage are not necessarily contiguous in host real storage. **Upwardly compatible:** • **[V5.1]** Condition code 1 no longer returned. All symbols are now resident.
X'94' VMDUMP and Symptom Record Service	**Upwardly compatible:** • **[V5.2]** New parameters: DUMPDW *list-addr* and SRDW *symptom-rec-addr*. • **[V6.2]** Execution can be halted with the CPHX or FORCE command.

Table 26. Changes to DIAGNOSE Codes (continued)

Code	Changes
X'98' Real I/O	**INCOMPATIBLE:** • **[V5.1]** Cannot be used in 24-bit addressing mode. **Upwardly compatible:** • **[V3.1]** New subfunctions: BLOCK DIAGNOSE, MULTIPLE REQUEST BLOCK. • **[V5.1]** For the LOCK and UNLOCK subfunctions, return code 1 is reserved (not used). V=R and V=F virtual machines are not supported. • **[V5.1]** For the LOCK subfunction, return code 3 is reserved (not used). 24-bit addressing mode guests cannot use this DIAGNOSE code, and 370 virtual machines are not supported. • **[V5.1]** For the Block Diagnose X'98' subfunction, global return code 2 is reserved (not used). V=R and V=F virtual machines are not supported. • **[V5.1]** For the Block Diagnose X'98' subfunction, global return codes 3 and 4 are reserved (not used). 24-bit addressing mode guests cannot use this DIAGNOSE code, and 370 virtual machines are not supported. • **[V5.2]** New Release and Lock Subfunction (Block Diagnose X'98' Multiple Request Block FCN Code X'20') allows you to lock guest pages into host real storage above or below 2 GB. The existing DIAGNOSE code X'98' LOCK subfunction continues to support lock requests below 2 GB only.
X'A4' Synchronous I/O (Standard CMS Blocksize)	**INCOMPATIBLE:** • **[V5.3]** This DIAGNOSE code does not support HyperPAV alias.
X'A8' Synchronous I/O (for All Devices)	**Upwardly compatible:** • **[V5.3]** I/O Request Flag - support for format-2 IDAW.
X'BC' Open and Query Spool File Characteristics	**Upwardly compatible:** • **[V2.2]** Depending on the specified buffer length, following the SECLABEL field the user's buffer will include the full (4-digit-year) date and the ISO date.
X'D8' Read Spool File Blocks	**Upwardly compatible:** • **[V2.2]** For subcode X'0000', a one-byte century indicator was added to the SFBLOK data area.
X'DC' Control Application Monitor Record Collection	**Upwardly compatible:** • **[V4.3]** Supports 64-bit addressing.
X'FC' Channel Path Reconfiguration Interface	**Incompatible:** • **[V6.2]** This DIAGNOSE code is no longer supported.

Table 26. Changes to DIAGNOSE Codes (continued)

Code	Changes
X'210' Retrieve Device Information	**Upwardly compatible:** • **[V3.1]** A byte code X'0C' of the VRDCBLOK control block has been added so that application programs can find out the underlying real hardware associated with tape devices operating in emulation mode. • **[V3.1]** New messages: HCP2327I, HCP2328I, HCP2340I, HCP2341I, HCP2342I, HCP2343I, HCP2344I, HCP2345I, HCP2346I, HCP2347I, HCP2348I, HCP2349I, HCP2350I, HCP2351I, HCP2352I, HCP2353I, HCP2354I, HCP2355I, HCP2356I, HCP2357I, HCP2358I, HCP2359I, HCP2360I, HCP2361I, HCP2362I, HCP6365I. • **[V4.4]** New VRDCUNDV values for the 3590 Model H. • **[V5.1]** New VRDCUNDV values for the 3592 Model J. • **[V5.3]** New VRDCBLOK value returned to indicate MIDAW support. • **[V5.4]** New VRDCBLOK value for the 3592 Model E06.
X'250' Block I/O (Standard Blocksize)	**INCOMPATIBLE:** • **[V5.3]** The DIAGNOSE code does not support HyperPAV alias. **Upwardly compatible:** • **[V5.2]** 64-bit support.
X'254' Access Real Subsystem	**INCOMPATIBLE:** • **[V4.3]** This DIAGNOSE code is no longer supported as a Programming Interface for customers.
X'258' Page-Reference Services	**Upwardly compatible:** • **[V4.3]** Supports 64-bit addressing.
X'260' Access Certain Virtual Machine Information	**Upwardly compatible:** • **[V5.2]** 64-bit addressing support. Updated subcodes: X'00000000', X'00000004', X'00000008'. New subcode: X'0000000C'.

Table 26. Changes to DIAGNOSE Codes (continued)

Code	Changes
X'26C' Access Certain System Information	**Upwardly compatible:** • **[V5.3]** Virtual Switch SNMP Agent support. Updated subcodes: X'00000004', X'00000008', X'0000000C', X'00000010', X'00000014', X'00000018', X'0000001C' , X'00000020', and X'00000024' have been added. • **[V5.3]** VSWITCH Port Isolation support. Updated subcode: X'00000024' to reflect isolation status. • **[V5.3]** Updated subcodes: X'00000008', X'00000010', X'00000018', X'00000020', and X'00000024' to change the Rx description. • **[V5.4]** Virtual Switch Networking Management support. Updates to Common Usage Notes. Updated subcodes: X'00000008', X'00000010', X'00000018', X'00000020', and X'00000024' to change the Rx description. • **[V5.4]** OSA QDIO Connection Isolation support. Updated subcode: X'00000020', a new error code X'19' (No OSA Connection Isolation) was added. In Subcode X'24' (N) Virtual Port information, DSECT CSIPNSTR, the drop bit in CSIPEST flag has been updated to indicate ISOLation ON is in effect. The FORWARD bit is no longer used and is marked as reserved. • **[V6.1]** Support for IBM zEnterprise System Networks. The following subcodes were updated with ensemble information: X'00000008', X'00000010', X'00000018', X'00000020', X'00000024', X'00000030'. • **[V6.2]** Single System Image support. Updated subcodes: X'00000008' , X'00000018', X'00000020', and X'00000030'. • **[V6.2]** Virtual Switch HiperSockets Bridge support. Updated subcodes: X'00000010' , X'00000020', and X'00000024'. • **[V6.3 VM65583]** Multi-VSwitch Link Aggregation Support. Updated subcodes: X'00000008' and X'00000020'. • **[V6.4 VM65925]** NICDEF Security Controls support. Updated subcode: X'00000020'.
X'270' Pseudo Timer Extended	**Upwardly compatible:** • **[V2.4]** Output has been expanded to include the version of DIAGNOSE code X'270', the user's default date format, and the system default date format.
X'2FC' Obtain Certain Guest Performance Data	**Upwardly compatible:** • **[V5.4]** Output has been expanded to include the CPU information.

Reserved DIAGNOSE Codes

Table 27 lists reserved DIAGNOSE codes that have changed. For additional information, see *z/VM: CP Programming Services*.

Note: These DIAGNOSE codes are not supported as Programming Interfaces for customers.

Table 27. Changes to Reserved DIAGNOSE Codes

Code	Changes
X'C4' Handle Class Override Data for UCR file	**INCOMPATIBLE:** • **[V6.3]** This DIAGNOSE code is no longer supported.

Table 27. Changes to Reserved DIAGNOSE Codes (continued)

Code	Changes
X'E0' System Trace File Interface	**Upwardly compatible:** • **[V5.2]** New block descriptor code C'L'.
X'F0' CSE Communication	**INCOMPATIBLE:** • **[V6.3]** Disabled. CSE is no longer supported. All subcodes will return CC=1 with caller's RC=24. No data is returned.
X'2AC' HCD Dynamic I/O	**Upwardly compatible:** • **[V6.3]** PCIe Guest Direct Attach Support. Updated condition code and return code description: RC64 - PCI function is not offline. • **[V6.4]** Updated condition code 1 and return code 8 description to include: RC2 - CHPID type is unknown.

STHYI Instruction

Table 28 identifies fields in the STHYI instruction function code X'0000' response buffer that have changed. For additional information, see *z/VM: CP Programming Services*.

Table 28. Changes to the STHYI Instruction Function Code X'0000' Response Buffer

Section	Changes
Partition	**Upwardly compatible:** • **[V6.3 VM65586, VM65696]** Support for multithreading: – INFPFLG1 field will contain X'80' flag to indicate multithreading is enabled. – INFPSCPS, INFPDCPS, INFPSIFL, and INFPDIFL fields contain count of cores when multithreading is enabled. – INFPWBCP, INFPABCP, INFPWBIF, and INFPABIF fields indicate capped capacity in cores. • **[V6.3 VM65716]** Support for LPAR group absolute capacity capping: – INFPVAL1 field will contain X'08' flag to indicate LPAR group absolute capacity capping is enabled. – INFPLGNM field contains the LPAR group name when LPAR group absolute capacity capping is enabled. – INFPLGCP field contains absolute capacity value for CP CPU type when LPAR group absolute capacity capping is enabled. – INFPLGIF field contains absolute capacity value for IFL CPU type when LPAR group absolute capacity capping is enabled.
Hypervisor	**Upwardly compatible:** • **[V6.3 VM65586, VM65696]** Support for multithreading: – New INFYCPT field indicates threads in use per CP core (valid only when multithreading is enabled). – New INFYIFLT field indicates threads in use per IFL core (valid only when multithreading is enabled). – INFYSCPS, INFYDCPS, INFYSIFL, and INFYDIFL fields contain count of cores when multithreading is enabled. • **[V6.3 VM65680]** If INFYFLG1 field X'40' flag is on, LIMITHARD caps use prorated core time for capping. If flag is off, raw CPU time is used.

Table 28. Changes to the STHYI Instruction Function Code X'0000' Response Buffer (continued)

Section	Changes
Guest	**Upwardly compatible:** • **[V6.3 VM65586, VM65696]** Support for multithreading: – INFGFLG1 field can specify new flags: X'08' indicates virtual CPs are thread dispatched; X'04' indicates virtual IFLs are thread dispatched. – INFGCPCC, INFGIFCC, INFGPCCC, and INFGPICC fields indicate capped capacity in cores.

Accounting Records

Table 29 lists accounting records that have changed. For additional information, see *z/VM: CP Planning and Administration*.

Table 29. Changes to Accounting Records

Accounting Record	Changes
Virtual Machine Resource Usage (Type 1)	**Upwardly compatible:** • **[V6.3 VM65586, VM65696]** When multithreading is enabled on the system, the column 33-36 and 37-40 fields contain the MT-1 equivalent time and a type F record is generated. • **[V6.3 VM65586, VM65696]** Column 67, formerly reserved, reports the number of threads on the Real CPU on which the virtual CPU was dispatched. When multithreading is enabled, additional data is reported on a type F record.
CPU Capability (Type D)	**Upwardly compatible:** • **[V6.1]** New fields: Nominal CPU capability, Capacity-Change Reason, Capacity-Adjustment Indication. • **[V6.3 VM65577]** CPU Capability data is continued on the type E record.
CPU Capability continuation data (Type E)	**Upwardly compatible:** • **[V6.3 VM65577]** New record produced at the same time as the type D record to report more detailed CPU capability data.
Virtual Machine Resource Usage 2 (Type F)	**Upwardly compatible:** • **[V6.3 VM65586, VM65696]** New record that is a continuation of the type 1 record. When multithreading is enabled on the system, a single type F record is produced for each virtual CPU for which one or two type 1 records (for primary and secondary CPU types) are produced. • **[V6.3 VM65680]** When multithreading is enabled, prorated core times are calculated and reported for every virtual machine.

CP Macros

Changes to CP macros are identified in the following topics:

- "System Definition Macros" on page 366
- "VM Data Space Macros" on page 366
- "Other CP Macros" on page 367

System Definition Macros

Table 30 lists System Definition macros that have changed.

Table 30. Changes to System Definition Macros

Macro	Changes
CSESYS	**INCOMPATIBLE:** • **[V6.3]** Not available for customer use. CSE is no longer supported.
CSETRACE	**INCOMPATIBLE:** • **[V6.3]** Not available for customer use. CSE is no longer supported.
CSEUSER	**INCOMPATIBLE:** • **[V6.3]** Not available for customer use. CSE is no longer supported.
SYSCPVOL	**INCOMPATIBLE:** • **[V6.3]** OWN and SHARED keywords are ignored. CSE is no longer supported.
SYSRES	**INCOMPATIBLE** • **[V5.1]** SYSVOL, SYSRES, SYSTYPE, and SYSNUC parameters are ignored. CP cannot be IPLed from a nucleus, only from a module on the parms disk.
SYSSTORE	**INCOMPATIBLE** • **[V3.1]** RIO370 parameter removed. • **[V5.1]** Parameters removed: VRSIZE, VRFREE. The V=R area and V=R recovery area are not supported.

VM Data Space Macros

Table 31 lists VM data space macros that have changed. For additional information, see *z/VM: CP Programming Services*.

Table 31. Changes to VM Data Space Macros

Macro	Changes
ADRSPACE	**Upwardly compatible:** • **[V5.1]** Operation exception cannot occur. 370 virtual machines are not supported.
ADRSPACE ISOLATE	**INCOMPATIBLE:** • **[V5.1]** MACHINE=370 parameter has no supported function. 370 virtual machines are not supported.
ADRSPACE PERMIT	**INCOMPATIBLE:** • **[V5.1]** MACHINE=370 parameter has no supported function. 370 virtual machines are not supported.
ADRSPACE QUERY	**INCOMPATIBLE:** • **[V5.1]** MACHINE=370 parameter has no supported function. 370 virtual machines are not supported.
ALSERV	**Upwardly compatible:** • **[V5.1]** Operation exception cannot occur. 370 virtual machines are not supported.
ALSERV ADD	**INCOMPATIBLE:** • **[V5.1]** MACHINE=370 parameter has no supported function. 370 virtual machines are not supported.

Table 31. Changes to VM Data Space Macros (continued)

Macro	Changes
ALSERV REMOVE	**INCOMPATIBLE:** • **[V5.1]** MACHINE=370 parameter has no supported function. 370 virtual machines are not supported.
MAPMDISK	**Upwardly compatible:** • **[V5.1]** Operation exception cannot occur. 370 virtual machines are not supported.
REFPAGE	**Upwardly compatible:** • **[V5.1]** Operation exception cannot occur. 370 virtual machines are not supported.
VMUDQ	**INCOMPATIBLE:** • **[V5.1]** MACHINE=370 parameter has no supported function. 370 virtual machines are not supported. **Upwardly compatible:** • **[V6.2]** MDISK parameter list has an additional field for MDISK definitions that are defined for the specified SSI member. • **[V6.4 VM65877]** A flag can be set in the LSTMDISK parameter list to return SUBCONFIG information rather than the system name.

Other CP Macros

Table 32 lists other CP macros that have changed. For additional information, see *z/VM: CP Exit Customization*.

Table 32. Changes to Other CP Macros

Macro	Changes
HCPCALL	**INCOMPATIBLE:** • **[V5.1]** HCPCALL now ensures that no extra positional parameters are specified. If any are found, the following message is displayed: `MNOTE 8, 'Extra positional parameter: pppppppp'` If *pppppppp* is blank, this message could indicate that incorrect continuation was used on the HCPCALL invocation. • **[V5.2]** ATTR parameter may be used to specify only entry point attributes, not module running attributes.
HCPPROLG	**INCOMPATIBLE:** • **[V5.1]** PAGEABLE attribute is not supported. All CP modules must be resident. If PAGEABLE is specified, CP changes the attribute to RESIDENT. • **[V5.2]** ATTR parameter no longer allows the DIRECTCALL and INDIRECTCALL keywords. Use the new HCPCALL parameter (HCPCALL=DIRECT or HCPCALL=INDIRECT). • **[V5.2]** TRSOURCE parameter is no longer allowed. Use the new RESTRICT parameter (RESTRICT=TRSOURCE).
HCPTKDEF Parser Token Definition	**Upwardly compatible:** • **[V2.4]** New conversion type: INSTRUCT.

Table 32. Changes to Other CP Macros (continued)

Macro	Changes
MDLATENT Exit Entry Definition	**INCOMPATIBLE:** • **[V5.1]** PAGEABLE parameter is not supported and has been removed. All modules included in the CP nucleus are resident. • **[V5.1]** NUCLEUS=ESA390 parameter is not supported and has been removed. The CP nucleus is 64-bit only. • **[V5.2]** Syntax has been extensively changed. Parameters previously retained for compatibility (FULLREG and NUCLEUS) have been removed. Parameters have been added for specifying the addressing mode (Amode) and translation mode (Tmode). The relationship between module attributes and entry point attributes and how they are specified in this macro have significantly changed. **Upwardly compatible:** • **[V5.1]** RESIDENT parameter is now the default. • **[V5.1]** FULLREG parameter is retained for compatibility and is equivalent to LONGREG. • **[V5.1]** NUCLEUS=BOTH parameter is retained for compatibility and is still the default, but it is now equivalent to NUCLEUS=ESAME.

CP System Services

Table 33 lists CP system services that have changed. For additional information see *z/VM: CP Programming Services*.

Table 33. Changes to CP System Services

System Service	Changes
*BLOCKIO DASD Block I/O System Service	**INCOMPATIBLE:** • **[V5.3]** This system service does not support HyperPAV alias.
*LOGREC Error Logging System Service	**Upwardly compatible:** • **[V5.1]** New record added: Special VM Record (SVMR).
*RPI Access Verification System Service	**INCOMPATIBLE:** • **[V5.1]** ACI entry points HCPRPDEP, HCPDA0RL, HCPDA0UL, and HCPDA0MC are no longer pageable. • **[V6.4]** CP denies access to default VLAN ID when ACILVDL contains X'FFFF' for a VLAN-aware virtual switch. **Upwardly compatible:** • **[V5.3]** New code (ACIRSCHK) and a new field (ACIPASSA) are added to ACIDATA. • **[V5.3]** Two new flags added to ACILOPTS: ACICHGPW and ACIUSEEP. • **[V6.2]** The ACIPARMS control block has been updated. In the ACIPARMS parameter lists, the LOGOFF and LOGON commands have been updated, and a new parameter list has been added for VMRELOCATE.

Table 33. Changes to CP System Services (continued)

System Service	Changes
*VMEVENT Notification of Asynchronous VSWITCH Events	**Upwardly compatible:** • **[V6.2]** A new class is added: Class 4, for Asynchronous VSWITCH events. • **[V6.3 VM65583]** New Class 4 Types added for Multi-VSwitch Link Aggregation Support.

CPXLOAD Directives

Table 34 lists CPXLOAD directives that have changed. For additional information see *z/VM: CP Programming Services*.

Table 34. Changes to CPXLOAD Directives

Directive	Changes
OPTIONS	**Upwardly compatible:** • **[V5.1]** LOCK and NOLOCK operands no longer have any effect. The CP nucleus, including dynamically loaded CP routines, is now completely resident.

VM Dump Tool

Table 35 lists VM Dump Tool functions that have changed. For additional information, see *z/VM: VM Dump Tool*.

Table 35. Changes to VM Dump Tool Command, Subcommands, and Macros

Function	Changes
VMDUMPTL command	**INCOMPATIBLE:** • **[V4.4]** Message HCQ031W changed to HCQ031E. **Upwardly compatible:** • **[V4.3]** New messages: HCQ001W, HCQ006W, HCQ011W, HCQ012W, HCQ040E, HCQ049E. • **[V4.4]** New messages: HCQ106E, HCQ107E, HCQ108E, HCQ109E, HCQ112E, HCQ113E, HCQ114E. • **[V5.1]** New message: HCQ007E • **[V5.4]** New operand: *dumpft*.
BLOCK macro	**INCOMPATIBLE:** • **[V5.4]** The names of bits are now included in BLOCK output for releases prior to V5.3. Any macro which depends on specific output from an older release will need to be updated. **Upwardly compatible:** • **[V4.3]** New options for the FIELDS operand. • **[V4.4]** New messages: HCQ090I, HCQ111E. • **[V5.2]** New optional address space prefix character (L, R or U). • **[V5.2]** New options: NOMap/MAP, CHars/HEx and BIts/NOBits.
CALLERS macro	**Upwardly compatible:** • **[V5.1]** Now accepts up to 17 characters of input.

Table 35. Changes to VM Dump Tool Command, Subcommands, and Macros (continued)

Function	Changes
CCW subcommand	**Upwardly compatible:** • **[V5.2]** New optional address space prefix character (L, R or U). • **[V5.2]** New options: BIts and NOBits.
CHAIN/CHAIN8 subcommand	**Upwardly compatible:** • **[V5.2]** New optional address space prefix character (L, R or U). • **[V5.2]** New options: DATa and NOData.
CODE subcommand	**Upwardly compatible:** • **[V5.4]** The DESCRIBE macro can be used to display a summary of the meaning of an abend code.
CPEBK subcommand	**Upwardly compatible:** • **[V5.1]** Now accepts up to 17 characters of input. • **[V5.2]** New optional address space prefix character (L, R or U).
CPEXITS macro	**INCOMPATIBLE:** • **[V5.4]** Due to changes in CP, the CPEXITS macro now supports dumps only at the level of V5.2 or later.
CPUUSE macro	• **[V6.3]** New output for a parked processor is "CPU XXXX Parked." • **[V6.3 VM65586, VM65696]** When multithreading is enabled, output includes CPU type and core ID. • **[V6.3 VM65586, VM65696]** New message: HCQ163E.
DESCRIBE macro	**Upwardly compatible:** • **[V5.4]** Changes to Usage Notes.
DISPLAY subcommand	**INCOMPATIBLE:** • **[V4.4]** Output from the DISPLAY subcommand, when issued from a macro, is now the same as DISPLAY subcommand output when the subcommand is issued from the command line. **Upwardly compatible:** • **[V5.2]** New optional address space prefix character (L, R or U). • **[V5.2]** The defaults for invocation from the command line or a macro are now the same.
DX macro	**INCOMPATIBLE:** • **[V5.2]** This macro is no longer supported.
EXTRACT subcommand	**INCOMPATIBLE:** • **[V5.1]** The DFIR/DFIZ and SYMPTOM operands require a decimal value. **Upwardly compatible:** • **[V5.1]** New options: MAPA, MAPN, and DISPL. • **[V5.1]** New message: HCQ121E. • **[V5.2]** New optional address space prefix character (L, R or U). • **[V5.2]** New functions: ADDRESS and ASIBK.
FRAMES subcommand	**Upwardly compatible:** • **[V5.1]** New message: HCQ127I. • **[V5.2]** New operands: NOSUM/SUMmary and NO2Gig/2Gig. • **[V6.3]** The page count fields on all output lines have been changed from 8 digits to 10 digits.

Table 35. Changes to VM Dump Tool Command, Subcommands, and Macros (continued)

Function	Changes
FRMAVL macro	**Upwardly compatible:** • **[V6.3]** Numeric digits increased to 20.
GREGS subcommand	**Upwardly compatible:** • **[V4.3]** New operands: MAP, NOMAP, LONG, SHORT.
HCQGDSPL function	**Upwardly compatible:** • **[V5.4]** The new HCQGDSPL function is available to return the displacement of a field from the BLOCK data. A number of VM Dump Tool macros were updated to use this new function. • **[V6.3 VM65586, VM65696]** New NOMSG option prevents the display of messages from HCQGDSPL.
INDQ subcommand	**Upwardly compatible:** • **[V5.1]** This subcommand is no longer supported.
INSTR subcommand	**Upwardly compatible:** • **[V5.2]** New optional address space prefix character (L, R or U). • **[V5.4]** A number of new mnemonics and formats have been added to the INSTR subcommand.
LASTTRAN macro	**Upwardly compatible:** • **[V5.2]** New macro: This macro can be used to display how the last logical or user-defined address was translated.
LOCATE subcommand	**Upwardly compatible:** • **[V4.3]** New message: HCQ004W. • **[V5.2]** New optional address space prefix character (L, R or U). • **[V5.2]** New operand: 2Gig
LOCDISP macro	**Upwardly compatible:** • **[V4.3]** New message: HCQ004W. • **[V5.2]** New optional address space prefix character (L, R or U). • **[V5.2]** New operand: 2Gig
MAP subcommand	**Upwardly compatible:** • **[V4.3]** New messages: HCQ086E, HCQ087E, HCP088E, HCQ100I, HCQ101E, HCP102E, HCP103E, HCP104I, HCP105E. • **[V5.1]** New messages: HCQ116W, HCQ117E, HCQ118E, HCQ123E, HCQ124E, HCQ126I.
PFXSAVE macro	**INCOMPATIBLE:** • **[V5.2]** The CPU keyword is made optional. The minimum abbreviation for Cpu is changed to CPU.
QUERY subcommand	**Upwardly compatible:** • **[V4.4]** New operands: DVFMACRO, IMPDVF. • **[V5.1]** New operand: DEBUG. • **[V5.2]** The new VMDTQRY subcommand is recommended over the QUERY subcommand. All operands of QUERY that have been supported in the past will continue to be supported without change. New operands will be added to VMDTQRY only.
RADIX macro	**Upwardly compatible:** • **[V5.2]** New optional address space prefix character (L, R or U).

Table 35. Changes to VM Dump Tool Command, Subcommands, and Macros (continued)

Function	Changes
RDEVBK subcommand	**Upwardly compatible:** • **[V5.1]** The output now displays the addresses of significant related control blocks rather than the storage of the control blocks.
RSCH subcommand	**Upwardly compatible:** • **[V5.1]** The output now displays the addresses of significant related control blocks rather than the storage of the control blocks.
SET subcommand	**Upwardly compatible:** • **[V4.4]** New operands: DVFMACRO, IMPDVF, XEDITPRE. • **[V4.4]** New message: HCQ110E. • **[V5.1]** New operand: DEBUG. • **[V5.2]** The new VMDTSET subcommand is recommended over the SET subcommand. All operands of SET that have been supported in the past will continue to be supported without change. New operands will be added to VMDTSET only.
SETVAR subcommand	**Upwardly compatible:** • **[V4.3]** New operand: WORD. • **[V5.2]** New optional address space prefix character (L, R or U).
SSASAVE macro	**INCOMPATIBLE:** • **[V5.2]** The CPU keyword is made optional. The minimum abbreviation for Cpu is changed to CPU.
SXSTE macro	**Upwardly compatible:** • **[V5.4]** The SXSTE macro is now supported.
TODCLOCK subcommand	**Upwardly compatible:** • **[V5.2]** New optional address space prefix character (L, R or U).
TRACE subcommand	**INCOMPATIBLE:** • **[V5.3]** Abbreviations are no longer allowed for TYPE or TERM names. • **[V5.4]** The old form of the TRACE subcommand has been removed. VMDTSET TRACE OLD and NEW are both still accepted, but both are handled as NEW. Any macros which depend on specific output from the TRACE OLD option must be updated. All VM Dump Tool modules unique to the old TRACE formatter have been removed. **Upwardly compatible:** • **[V4.3]** New messages: HCP047W, HCQ074E. • **[V5.1]** The FROM and TO operands no longer check the frame type. These operands can now be used when the frame table is not available. • **[V5.3]** There is a new format to the TRACE subcommand. The output of the TRACE subcommand for many trace entry codes is different. The information has been reformatted with more consistent keywords and order. At this time, the old format is still available. • **[V6.2]** New operands: FRToken *token* and FRName *name*.
TWOGFRM macro	**Upwardly compatible:** • **[V6.3]** Numeric digits increased to 20.
VDEVBK subcommand	**Upwardly compatible:** • **[V5.1]** The output now displays the addresses of significant related control blocks rather than the storage of the control blocks.
VIRTUAL macro	**Upwardly compatible:** • **[V4.4]** New operand: DETAILS.

Table 35. Changes to VM Dump Tool Command, Subcommands, and Macros (continued)

Function	Changes
VMDBK subcommand	**INCOMPATIBLE:** • **[V4.4]** Message HCQ018E changed to HCQ018I. **Upwardly compatible:** • **[V4.3]** New message: HCQ018E.
VMDSCAN macro	**Upwardly compatible:** • **[V4.4]** New operand: *userid*.
VMDTQRY subcommand	**INCOMPATIBLE:** • **[V5.4]** The output of the VMDTQRY TRACE LIBRARY subcommand has changed to reflect the addition of the new DEFAULT option. Any macro which depends on specific output from an older release will need to be updated. **Upwardly compatible:** • **[V5.2]** The new VMDTQRY subcommand is the preferred way to get information. New operands will be added only to VMDTQRY, which is preferred over QUERY. • **[V5.4]** New operands: ABLib, BLKlibs, MNEmlib, SEParate. • **[V5.4]** New operands: FILes, RECords. • **[V6.2]** New operand: FRTrace. Updated examples.
VMDTSET subcommand	**Upwardly compatible:** • **[V5.2]** The new VMDTSET subcommand is the preferred way to get information. New operands will be added only to VMDTSET, which is preferred over SET. • **[V5.4]** New operands: SEParate, ABLib, BLKlibs, MNEmlib. Changes to TRACE operands.
VSCH subcommand	**Upwardly compatible:** • **[V5.1]** The output now displays the addresses of significant related control blocks rather than the storage of the control blocks.

CMS Changes

Changes to CMS interfaces are identified in the following topics:

For information about changes to CMS monitor records, see IBM: z/VM data areas, control blocks, and monitor records (www.ibm.com/vm/pubs/ctlblk.html).

General CMS Commands

Table 36 on page 374 lists general CMS commands that have changed. For additional information, see *z/VM: CMS Commands and Utilities Reference*.

Table 36. Changes to General CMS Commands

Command	Changes
ACCESS	**Upwardly compatible:** • **[V5.2]** New message: DMS114S.
BROWSE	**INCOMPATIBLE:** • **[V5.1]** Screen layout has been adjusted to accommodate larger file and disk sizes.
CMSDESK	**INCOMPATIBLE:** • **[V2.2]** Message DMS2302E replaced by new format of DMS622E, same return code. **Upwardly compatible:** • **[V2.2]** Supports three date formats: short date, full date, ISO date. • **[V2.2]** Display of the contents of the File Manager application has been changed. • **[V2.2]** Displays a toolbar on the applications. • **[V2.2]** New and changed menu items for some applications. • **[V3.1]** See "[V3.1] Graphical User Interface Facility Changes" on page 168.
CMSSERV	**INCOMPATIBLE:** • **[V5.2]** Not supported. Use TCP/IP for z/VM.
CONV2WD	**INCOMPATIBLE:** • **[V5.1]** Command removed.
COPYFILE	**INCOMPATIBLE:** • **[V2.2]** New message: DMS516E.
CREATE (in general)	**Upwardly compatible:** • **[V6.4 VM66028]** Renamed help files for subcommands. See "[V6.4 APAR] z/VM HELP Quality Improvement" on page 238.
CREATE DIRECTORY	**Upwardly compatible:** • **[V2.2]** New message for authorization failure from ESM: DMS1331E.
CSLLIST	**Upwardly compatible:** • **[V2.3]** When an = or ? is typed as the first character in the "Cmd" area of a line in the CSLLIST display screen, all characters following the = or ? are ignored.
CSLMAP	**Upwardly compatible:** • **[V2.3]** When an = or ? is typed as the first character in the "Cmd" area of a line in the CSLMAP display screen, all characters following the = or ? are ignored.
DEFAULTS	**Upwardly compatible:** • **[V2.2]** New options supported as parameters for FILELIST and RDRLIST: VMDATE, SHORTDATE, FULLDATE, ISODATE. • **[V2.3]** New options supported as parameters for NETDATA: VMDATE, SHORTDATE, FULLDATE, ISODATE. • **[V6.4]** FOR option default number of records for the PEEK command changed from 200 to 25000.
DELETE	**Upwardly compatible:** • **[V6.4 VM66028, VM66029]** Renamed help files for subcommands. See "[V6.4 APAR] z/VM HELP Quality Improvement" on page 238.

Table 36. Changes to General CMS Commands (continued)

Command	Changes
DIRLIST	**Upwardly compatible:** • **[V2.3]** When an = or ? is typed as the first character in the "Cmd" area of a line in the DIRLIST display screen, all characters following the = or ? are ignored.
ERASE	**Upwardly compatible:** • **[V2.2]** New message for authorization failure from ESM: DMS1332E.
FILEDEF	**Upwardly compatible:** • **[V2.2]** New option: LIBSRV. • **[V2.4]** Allows LRECL definitions up to 65535 bytes for OS variable spanned records (under XLRI processing) and non-OS CMS files. • **[V2.4]** Allows BLKSIZE definitions up to 65535 bytes for non-OS CMS files.
FILELIST	**INCOMPATIBLE:** • **[V2.2]** When FILELIST is specified with the SHARE option, if a pre-VM/ESA V2.2 profile (PROFFSHR XEDIT) resides on a disk accessed ahead of the S-disk, sorts by date or size will not work. IBM recommends that you recreate all non-system FILELIST profiles. See *z/VM: CMS Commands and Utilities Reference*. • **[V2.2]** If you file the file created by FILELIST, that file might contain new and changed fields (on the far right). **Upwardly compatible:** • **[V2.2]** New options to specify date format: VMDATE, SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** BEFORE *date* and AFTER *date* options support 4-digit years. • **[V2.2]** Screens and responses support 4-digit years. • **[V2.2]** If a date format option is not specified on the FILELIST command, the CMS DEFAULTS date format setting for FILELIST will be used. • **[V2.3]** When an = or ? is typed as the first character in the "Cmd" area of a line in the FILELIST display screen, all characters following the = or ? are ignored.
FLIST	**INCOMPATIBLE:** • **[V5.1]** Screen layout has been adjusted to accommodate larger file and disk sizes.
FORMAT	**Upwardly compatible:** • **[V5.2]** Message variation: DMS114S.
GENMOD	**INCOMPATIBLE:** • **[V5.1]** 370 option removed. 370 virtual machines are not supported.
GLOBAL	**Upwardly compatible:** • **[V2.2]** New message for duplicate library name in input list: DMS045W. The duplicates are ignored.

Table 36. Changes to General CMS Commands (continued)

Command	Changes
HELP	**INCOMPATIBLE:** • **[V4.4]** HELP components removed: CMSUTIL, CPOTHER, CPUTIL, SOCKETS. CPOTHER component (HELPCPOT) files renamed to HELPCP. • **[V5.1]** HELP components removed: SPTAPE. • **[V5.2]** HELP components removed: SRPI, VMADMIN. • **[V5.2]** CP LOCATE commands changed from HELP component CP to new HELP component LOCATE. • **[V5.3]** HELP components removed: BOOTPD, RTE. • **[V5.4]** HELP components removed: TFTD, X25IPI. • **[V6.1]** HELP components removed: HMF, HMFQUERY, HMFSET. • **[V6.3]** HELP components removed: DHCPD, LPD, XSPOOL. **Upwardly compatible:** • **[V4.4]** HELP components added: BOOTPD, DHCPD, DNS, FTP, GDDMXD, HELP, IMAPADM, LE, LPD, MPROUTE, MROUTINE, NFS, RTE, SMTP, SNMP, SSLADMIN, TCPIP, TELNET, TFTD, TFTP, UFTD, X25IPI. • **[V5.1]** HELP component added: FCX. • **[V5.2]** HELP component added: IPFORMAT. • **[V5.3]** HELP components added: LDAP, MSOCKETS, RSCSAUTH. • **[V5.4]** HELP components added: DFSMS, VMDT. • **[V6.1]** HELP components added: NSINTER, NSLOOKUP. • **[V6.4]** HELP components added. See "[V6.4 APAR] z/VM HELP Quality Improvement" on page 238.
IDENTIFY	**Upwardly compatible:** • **[V2.2]** New options to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** Responses support 4-digit years. • **[V2.2]** If the user's default date format is changed from SHORTDATE to FULLDATE or ISODATE, and a date format option is not specified on the IDENTIFY command, the output from IDENTIFY uses the new default date format. This causes the date to be expanded to include the 4-digit year and also causes fields to the right of the date to be shifted. • **[V2.3]** New option for TCP/IP mail integration: TCPIP.
LISTDIR	**Upwardly compatible:** • **[V5.2]** New message: DMS1240E.
LISTDS	**Upwardly compatible:** • **[V2.3]** New options to specify date format: VMDATE, SHORTDATE, FULLDATE, ISODATE. • **[V2.3]** Responses support 4-digit years. • **[V2.3]** If the user's default date format is changed to FULLDATE or ISODATE, and the SHORTDATE option is not specified on the LISTDS command, any output from LISTDS with dates uses the new default date format. This causes the dates to be expanded to include the 4-digit year and also causes fields to be shifted.

Table 36. Changes to General CMS Commands (continued)

Command	Changes
LISTFILE	**INCOMPATIBLE:** • **[V2.4]** Message DMS550E changed to DMS765E. **Upwardly compatible:** • **[V2.2]** New options to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** BEFORE *date* and AFTER *date* options support 4-digit years. • **[V2.2]** Responses support 4-digit years. • **[V2.2]** If the user's default date format is changed from SHORTDATE to FULLDATE or ISODATE, and a date format option is not specified on the LISTFILE command, any output from LISTFILE with dates uses the new default date format. This causes the dates to be expanded to include the 4-digit year and also causes fields to be shifted.
LKED	**Upwardly compatible:** • **[V2.4]** Defaults for *value1* and *value2* on the SIZE option have been increased to '400K' and '100K', respectively.
LOAD	**Upwardly compatible:** • **[V2.2]** New message for insufficient storage above 16 MB: DMS891W.
LOADMOD	**INCOMPATIBLE:** • **[V2.2]** Changed message (new text possible): DMS639E.
MACLIST	**Upwardly compatible:** • **[V2.3]** When an = or ? is typed as the first character in the “Cmd” area of a line in the MACLIST display screen, all characters following the = or ? are ignored.
MACLMIG	**INCOMPATIBLE:** • **[V5.1]** Command removed.
MOVEFILE	**Upwardly compatible:** • **[V2.4]** Can process OS variable spanned records (under XLRI processing) and non-OS CMS files with record lengths up to 65535 bytes. • **[V2.4]** Adjusts output file sizes for compatibility between CMS and OS. • **[V2.4]** Allows greater FILEDEF default flexibility for file attributes (RECFM, LRECL, BLKSIZE). • **[V2.4]** If default size values are used, fixes record truncation problems when moving data files from fixed to variable format. • **[V2.4]** New message: DMS1116E. • **[V4.4]** Message DMS2139I has additional text.
NAMES	**Upwardly compatible:** • **[V6.2]** Updated the Node and List options for an IPv6 address.

Table 36. Changes to General CMS Commands (continued)

Command	Changes
NETDATA	**Upwardly compatible:** • **[V2.3]** New options to specify date format: VMDATE, SHORTDATE, FULLDATE, ISODATE. • **[V2.3]** Supports 4-digit-year date formats for entries in the *userid* NETLOG file. • **[V2.3]** Responses support 4-digit years. • **[V2.3]** If the user's default date format is changed to FULLDATE or ISODATE, and the SHORTDATE option is not specified on the NETDATA command, any output from NETDATA with dates uses the new default date format. This causes the dates to be expanded to include the 4-digit year and also causes fields to be shifted.
NOTE	**INCOMPATIBLE:** • **[V4.4]** In the Date field of the note header, the name of the time zone has been replaced with the offset from Coordinated Universal Time (UTC), *zzzzz*. The first character is a plus (+) or minus (-), indicating whether the local time is ahead of (east of) or behind (west of) UTC. The next two characters (digits) indicate the number of hours difference from UTC. The last two characters (digits) indicate the number of additional minutes difference from UTC. **Upwardly compatible:** • **[V2.2]** In the Date field of the note header, the year is now displayed with four digits. • **[V2.3]** TCP/IP domain names accepted as user IDs or as the resolution of nicknames. • **[V2.3]** Supports 4-digit-year date formats for entries in the *userid* NETLOG file. • **[V6.2]** Updated the *name* operand
NUCXLOAD	**INCOMPATIBLE:** • **[V2.2]** Changed message (new text possible): DMS639E.
PEEK	**Upwardly compatible:** • **[V2.3]** TCP/IP origin domain name address used when available and shown on PEEK message line for origin within current space and formatting limitations. • **[V6.4]** Default number of records displayed for PEEK has been changed from 200 to 25000.
PIPE	**Upwardly compatible:** • **[V6.4]** CMS Pipelines has been upgraded and offers new function and capabilities. In general, applications created in earlier z/VM releases that use CMS Pipelines as specified will continue to run without any change to the application. For additional migration information, see *z/VM: CMS Pipelines User's Guide and Reference*. This new CMS Pipelines publication is based on *CMS/TSO Pipelines: Author's Edition*, SL26-0018, and replaces both of the previous z/VM CMS Pipelines publications.
QUERY (in general)	**Upwardly compatible:** • **[V3.1]** New operand: TAPENEVR. • **[V4.3]** New operand: TAPECSL. • **[V5.1]** New operands: EDEVICE, HYPERSWAP. • Also see specific QUERY commands below.

Table 36. Changes to General CMS Commands (continued)

Command	Changes
QUERY CMSLEVEL	**Upwardly compatible:** • **[V6.2]** Response includes added information and indicates when z/CMS is running in the virtual machine. • **[V6.4]** Response contains changed value for the new CMS level.
QUERY CMSREL	**Upwardly compatible:** • **[V6.4]** Response contains changed value for the new release level.
QUERY FILEDEF	**Upwardly compatible:** • **[V2.4]** New operand, ATTRIBUT, and its response.
RDRLIST	**INCOMPATIBLE:** • **[V2.2]** If you file the file created by RDRLIST, that file might contain new and changed fields (on the far right). **Upwardly compatible:** • **[V2.2]** New options to specify date format: VMDATE, SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** Screen supports 4-digit years. • **[V2.2]** If a date format option is not specified on the RDRLIST command, the CMS DEFAULTS date format setting for RDRLIST will be used. • **[V2.3]** When an = or ? is typed as the first character in the "Cmd" area of a line in the RDRLIST display screen, all characters following the = or ? are ignored. • **[V2.3]** TCP/IP origin domain name address used when available and shown on the RDRLIST panel origin area within current space and formatting limitations.
RECEIVE	**Upwardly compatible:** • **[V2.3]** Supports 4-digit-year date formats for entries in the *userid* NETLOG file.
SADT	**INCOMPATIBLE:** • **[V5.1]** Screen layout has been adjusted to accommodate larger file and disk sizes.
SEGMENT	**Upwardly compatible:** • **[V6.4 VM66028]** Renamed help files for subcommands. See "[V6.4 APAR] z/VM HELP Quality Improvement" on page 238.
SENDFILE	**Upwardly compatible:** • **[V2.3]** TCP/IP domain names accepted as user IDs or as the resolution of nicknames. • **[V2.3]** New options to specify the transmission method: SMTP, MIME, UFTSYNC, UFTASYNC. • **[V2.3]** Supports 4-digit-year date formats for entries in the *userid* NETLOG file. • **[V5.2]** Four new options to send files: BINARY-ATTACH, BINARY-INLINE, ASCII-ATTACH, and ASCII-INLINE. • **[V5.2]** Descriptions updated for the TRANSLATE, NODBCS, HANGEUL, KANJI and TCHINESE options. • **[V6.2]** Description updated for *node*.
SET (in general)	**Upwardly compatible:** • **[V3.1]** New operand: TAPENEVR. • **[V4.3]** New operand: TAPECSL.

Table 36. Changes to General CMS Commands (continued)

Command	Changes
SET CMS370AC	**Upwardly compatible:** • **[V6.2]** New message: DMS2632E.
TAPE	**Upwardly compatible:** • **[V2.2]** If the tape is under the control of a Tape Library Dataserver machine, and the DFSMS/VM Removable Media Services (RMS) FSMPPSI CSLLIB is available to CMS, the RUN (rewind and unload) function of TAPE calls the RMS FSMRMDMT (Demount) CSL routine to have the Dataserver unmount the tape. • **[V2.2]** Added message for Tape Library Dataserver support: DMS2147W. • **[V4.4]** Message DMS2139I has additional text.
TELL	**Upwardly compatible:** • **[V2.3]** Accepts a TCP/IP domain name as part of the destination information.
VMFPLC2	**Upwardly compatible:** • **[V2.2]** If the tape is under the control of a Tape Library Dataserver machine, and the DFSMS/VM Removable Media Services (RMS) FSMPPSI CSLLIB is available to CMS, the RUN (rewind and unload) function of VMFPLC2 calls the RMS FSMRMDMT (Demount) CSL routine to have the Dataserver unmount the tape. • **[V2.2]** Added message for Tape Library Dataserver support: DMS2147W. • **[V4.4]** Message DMS2139I has additional text.
VMLINK	**INCOMPATIBLE:** • **[V3.1]** Previously, parameters passed to exits were automatically uppercased before the exit was called. Now they remain in the case entered by the user. • **[V4.4]** A file mode extension specified in the *MODES record in the control file is no longer ignored. It is used unless overridden by the command line or by the :product tag in the NAMES file. • **[V6.2]** The PUSH option is ignored when DETACH, RELEASE, or NOKEEP is specified. **Upwardly compatible:** • **[V2.3]** When an = or ? is typed as the first character in the "Cmd" area of a line in the VMLINK display screen, all characters following the = or ? are ignored. • **[V3.1]** New variables: .EX, .PX. • **[V3.1]** Error message returned for each failing INVOKE command. • **[V3.1]** Message DMS2062I includes the name of the nickname being used when the QUERY option is processed. • **[V3.1]** Additional changes to operation and documentation. See "[V3.1] VMLINK Improvements" on page 168. • **[V4.4]** An asterisk (*) or equals sign (=) can be specified in the *MODES record in the control file to indicate that the default search order, Z-A, should be used. • **[V6.2]** New variables: .AS, .LS.
XEDIT	**INCOMPATIBLE:** • **[V6.2]** For files with file types other than the special file types recognized by the editor, the default case has been changed from uppercase to mixed case.

CMS Utilities

Table 37 lists CMS utilities that have changed. For additional information, see *z/VM: CMS Commands and Utilities Reference*.

Table 37. Changes to CMS Utilities

Utility	Changes
ACCOUNT	**Upwardly compatible:** • **[V5.3]** New options: RCPU and VCPU. • **[V5.4]** ICF support has been added. • **[V6.2]** New operand: CPUCAP. • **[V6.2]** New messages: DMS2265E, DMS2266I. • **[V6.2]** Message variation: DMS2263E.
AUDITOR	**Upwardly compatible:** • **[V6.2]** When using AUDITOR for the first time, uncomment the call to the AUDITOR EXEC.
DIRMAP	**INCOMPATIBLE:** • **[V5.1]** Screen layout has been adjusted to accommodate larger file and disk sizes. • **[V6.2]** DIRMAP MODULE moved from MAINT 190 to PMAINT 551. • **[V6.2]** Output file record length has been changed from 80 characters to 100 characters to accommodate columns for new SUBCONFIG and MEMBER information. **Upwardly compatible:** • **[V6.2]** From an authorization perspective, the Ownerid value specifies the user ID to which the MDISK belongs. For multi-configuration virtual machines, the MDISKs can be defined within either the IDENTITY or SUBCONFIG stanzas. In both cases the Ownerid value specifies the name from the IDENTITY statement. If the SUBCONFIG statement is not referenced by any BUILD statement, the Ownerid value does not contain a valid value. • **[V6.2]** The SUBCONFIG value in the output is blank unless the MDISK was defined within a SUBCONFIG definition. • **[V6.2]** The MEMBER value in the output is * unless the MDISK was defined within a SUBCONFIG definition. The value displayed is the system ID from the BUILD statement. If there is no corresponding BUILD statement the value is ***NOUSER!**. • **[V6.2]** If a LINK statement specifies an IDENTITY as a user ID and a Vaddr that is defined in multiple SUBCONFIG stanzas referenced by the IDENTITY's BUILD statements, then the LINK is shown for every possible match. • **[V6.2]** If the LINK statement is defined within a SUBCONFIG stanza, then the link is not associated with MDISKs or LINKs that are defined following SUBCONFIG statements defined for different systems. • **[V6.2]** If END is specified as the length of a minidisk, the ending cylinder for that disk is determined by the other minidisks defined for that device. If END is the only length used for that device, DIRMAP will use the lowest cylinder length for that device type.
QSYSOWN	**INCOMPATIBLE:** • **[V5.1]** Screen layout has been adjusted to accommodate larger file and disk sizes.

Table 37. Changes to CMS Utilities (continued)

Utility	Changes
SYSWATCH	**INCOMPATIBLE:** • **[V6.1]** Privilege class E required when monitoring PAGE and SPOOL allocation data from the CP INDICATE LOAD command.

CMS File Pool Administration and Operator Commands

Table 38 lists CMS file pool administration and operator commands that have changed. For additional information, see *z/VM: CMS File Pool Planning, Administration, and Operation*.

Table 38. Changes to CMS File Pool Administration and Operator Commands

Command	Changes
AUDIT	**Upwardly compatible:** • **[V2.2]** New operands: *fn ft*, REPLACE. • **[V2.2]** Added messages: DMS024E, DMS1258E, DMS3253I, DMS3254E, DMS3255E. • **[V2.2]** Changed message: DMS3470W (new text possible).
DELETE LOCK	**Upwardly compatible:** • **[V5.2]** New operand: FRom *userid*.
DELETE USER	**Upwardly compatible:** • **[V2.2]** New options: DELAUTH | KEEPAUTH. • **[V2.2]** Added message: DMS2023E.
FILEPOOL RELOAD	**Upwardly compatible:** • **[V2.2]** New message: DMS3455I.
FILEPOOL UNLOAD	**Upwardly compatible:** • **[V2.2]** New message: DMS3455I.

OPENVM Commands

Table 39 lists OPENVM commands that have changed. For additional information, see *z/VM: OpenExtensions Commands Reference*.

Table 39. Changes to OPENVM Commands

Command	Changes
OPENVM DEBUG	**Upwardly compatible:** • **[V3.1]** New operands: ALL, NOALL, DUMP, NODUMP, FILEIO, NOFILEIO, FLOW, NOFLOW, MOUNT, NOMOUNT, NFSREQUEST, NONFSREQUEST, OTHER, NOOTHER, RPCBUFFERS, NORPCBUFFERS, RPCLIBRARY, NORPCLIBRARY, WRAPSIZE 500, WRAPSIZE *n*, FORMAT. • **[V3.1]** Supports tracing NFS and BFS Client events.
OPENVM MOUNT	**Upwardly compatible:** • **[V3.1]** Supports new local NFS options.
OPENVM OWNER	**Upwardly compatible:** • **[V3.1]** New operands: *gid*, *uid*.
OPENVM QUERY MOUNT	**Upwardly compatible:** • **[V3.1]** New options: NODETAILS, DETAILS. • **[V3.1]** Supports mounted NFS file systems.

Table 39. Changes to OPENVM Commands (continued)

Command	Changes
OPENVM RUN	**INCOMPATIBLE:** • **[V2.2]** Changed message (new text possible): DMS639E.

XEDIT Subcommands

Table 40 lists XEDIT subcommands that have changed. For additional information, see *z/VM: XEDIT Commands and Macros Reference*.

Table 40. Changes to XEDIT Subcommands

Subcommand	Changes
FILE	**Upwardly compatible:** • **[V3.1]** Updated return codes: 32, 55. • **[V3.1]** New messages: 1019, 1020. • **[V3.1]** Supports mounted NFS file systems.
GET	**Upwardly compatible:** • **[V3.1]** Added message 1019 and 1020; updated return code 32 and 55. • **[V3.1]** Supports mounted NFS file systems.
LOAD	**INCOMPATIBLE:** • **[V6.2]** For files with file types other than the special file types recognized by the editor, the default case has been changed from uppercase to mixed case. **Upwardly compatible:** • **[V3.1]** Added message 1019 and 1020; updated return code 32 and 55. • **[V3.1]** Supports mounted NFS file systems.
PUT	**Upwardly compatible:** • **[V3.1]** Added message 1019 and 1020; updated return code 32 and 55. • **[V3.1]** Supports mounted NFS file systems.
PUTD	**Upwardly compatible:** • **[V3.1]** Added message 1019 and 1020; updated return code 32 and 55. • **[V3.1]** Supports mounted NFS file systems.
SAVE	**Upwardly compatible:** • **[V3.1]** Added message 1019 and 1020; updated return code 32 and 55. • **[V3.1]** Supports mounted NFS file systems.
XEDIT	**INCOMPATIBLE:** • **[V6.2]** For files with file types other than the special file types recognized by the editor, the default case has been changed from uppercase to mixed case. **Upwardly compatible:** • **[V3.1]** Added message 1019 and 1020; updated return code 32 and 55. • **[V3.1]** Supports mounted NFS file systems.

CMS Routines

Changes to CMS routines are identified in the following topics:

- "General CMS Callable Services" on page 384
- "CMS Multitasking Routines" on page 391

General CMS Callable Services

Table 41 lists general CMS callable services that have changed. For additional information, see *z/VM: CMS Callable Services Reference*, unless otherwise indicated.

Table 41. Changes to General CMS Callable Services

Routine	Changes
DMSCLBLK	**Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** *date* parameter supports 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason codes: 90310, 90320, 90330, 90492, 90495.
DMSCLDBK	**Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** *date* parameter supports 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason codes: 90320, 90330, 90495.
DMSCLOSE	**Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** *date* parameter supports 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason codes: 90320, 90330, 90495.
DMSCRDIR	**Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** *date* parameter supports 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason code: 90495.
DMSCRFIL	**Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** *date* and *create_date* parameters support 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason code: 90495.
DMSCROB	**Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** *date* and *create_date* parameters support 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason code: 90495.
DMSDEUSR	**Upwardly compatible:** • **[V2.2]** New parameters: DELAUTH, KEEPAUTH, and *length4*. • **[V2.2]** New reason code: 98700.

Table 41. Changes to General CMS Callable Services (continued)

Routine	Changes
DMSENUSR	**Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE, *length7*. • **[V2.2]** *date* parameter supports 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason codes: 90310, 90330, 90495.
DMSERP	**Upwardly compatible:** • **[V2.2]** New information names for Year 2000 support (FILE_DATE_CENTURY, ACT_FILE_DATE_CENTRY (note that U is omitted), and YEAR2000_SUPPORT).
DMSEXIDI	**INCOMPATIBLE:** • **[V2.2]** When this routine is called from a REXX program, the date field is returned as 10 characters. Previously, it was returned as 8 characters. **Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** *last_change_date* and *create_date* parameters support 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason codes: 90320, 90330.
DMSEXIFI	**INCOMPATIBLE:** • **[V2.2]** When this routine is called from a REXX program, the date field is returned as 10 characters. Previously, it was returned as 8 characters. **Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** *date*, *dateref*, *create_date*, and *last_change_date* parameters support 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason codes: 90320, 90330.

Table 41. Changes to General CMS Callable Services (continued)

Routine	Changes
DMSEXIST	**Upwardly compatible:** • **[V2.2]** Offsets have changed in the FILE data record (and the record length has increased to 436 bytes): **OFFSET** **Field Name — Change Description** **339 (X'153')** *dec_date_ext* — Previously reserved **343 (X'157')** *date_ext* — Added **353 (X'161')** *iso_date_ext* — Added **363 (X'16B')** *dec_dateref_ext* — Added **367 (X'16F')** *dateref_ext* — Added **377 (X'179')** *iso_dateref_ext* — Added **387 (X'183')** *dec_cr_date_ext* — Added **391 (X'187')** *cr_date_ext* — Added **401 (X'191')** *iso_cr_date_ext* — Added **411 (X'19B')** *dec_last_change_date_ext* — Added **415 (X'19F')** *last_change_date_ext* — Added **425 (X'1A9')** *iso_last_change_date_ext* — Added **435 (X'1B3')** Reserved — Added

Table 41. Changes to General CMS Callable Services (continued)

Routine	Changes
DMSEXIST (continued)	**Upwardly compatible:** • **[V2.2]** Offsets have changed in the DIRECTORY data record (and the record length has increased to 308 bytes): **OFFSET** **Field Name — Change Description** **254 (X'FE')** *dec_last_change_date_ext* — Previously reserved **258 (X'102')** *last_change_date_ext* — Added **268 (X'10C')** *iso_last_change_date_ext* — Added **278 (X'116')** *dec_cr_date_ext* — Added **282 (X'11A')** *cr_date_ext* — Added **292 (X'124')** *iso_cr_date_ext* — Added **302 (X'12E')** Reserved — Added
DMSGETDA	**INCOMPATIBLE:** • **[V2.2]** When this routine is called from a REXX program, the date field is returned as 10 characters. Previously, it was returned as 8 characters. **Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE and *length2*. • **[V2.2]** *date* parameter supports 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason codes: 90310, 90320, 90330.
DMSGETDF	**INCOMPATIBLE:** • **[V2.2]** When this routine is called from a REXX program, the date field is returned as 10 characters. Previously, it was returned as 8 characters. **Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE and *length2*. • **[V2.2]** *date* parameter supports 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason codes: 90310, 90320, 90330.

Table 41. Changes to General CMS Callable Services (continued)

Routine	Changes
DMSGETDI	**Upwardly compatible:** • **[V2.2]** Offsets have changed in the FILE data record (and the record length has increased to 112 bytes): **OFFSET** **Field Name — Change Description** **88 (X'58')** *dec_date_ext* — Added **92 (X'5C')** *date_ext* — Added **102 (X'66')** *iso_date_ext* — Added • **[V2.2]** Offsets have changed in the FILEEXT data record (and the record length has increased to 284 bytes): **OFFSET** **Field Name — Change Description** **185 (X'B9')** *dec_date_ext* — Previously reserved **189 (X'BD')** *date_ext* — Added **199 (X'C7')** *iso_date_ext* — Added **209 (X'D1')** *dec_dateref_ext* — Added **213 (X'D5')** *dateref_ext* — Added **223 (X'DF')** *iso_dateref_ext* — Added **233 (X'E9')** *dec_cr_date_ext* — Added **237 (X'ED')** *cr_date_ext* — Added **247 (X'F7')** *iso_cr_date_ext* — Added **257 (X'101')** *dec_last_change_date_ext* — Added **261 (X'105')** *last_change_date_ext* — Added **271 (X'10F')** *iso_last_change_date_ext* — Added **281 (X'119')** Reserved — Added

Table 41. Changes to General CMS Callable Services (continued)

Routine	Changes
DMSGETDI (continued)	**Upwardly compatible:** • **[V2.2]** Offsets have changed in the SEARCHALL and SEARCHAUTH data records (and the record length has increased to 252 bytes): **OFFSET** **Field Name/Change Description** **226 (X'E2')** Reserved — Added **228 (X'E4')** *dec_date_ext* — Added **232 (X'E8')** *date_ext* — Added **242 (X'F2')** *iso_date_ext* — Added
DMSGETDS	**INCOMPATIBLE:** • **[V2.2]** When this routine is called from a REXX program, the date field is returned as 10 characters. Previously, it was returned as 8 characters. **Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE and *length2*. • **[V2.2]** *date* parameter supports 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason codes: 90310, 90320, 90330.
DMSGETDX	**INCOMPATIBLE:** • **[V2.2]** When this routine is called from a REXX program, the date field is returned as 10 characters. Previously, it was returned as 8 characters. **Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE and *length2*. • **[V2.2]** *date*, *dateref*, *create_date*, and *update_date* parameters support 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason codes: 90310, 90320, 90330.
DMSOPBLK	**INCOMPATIBLE:** • **[V2.2]** When this routine is called from a REXX program, the date field is returned as 10 characters. Previously, it was returned as 8 characters. **Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** *date*, *create_date*, and *dateref* parameters support 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason code: 90495.

Table 41. Changes to General CMS Callable Services (continued)

Routine	Changes
DMSOPDBK	**Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** *create_date* parameter supports 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason code: 90495.
DMSOPEN	**Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** *create_date* parameter supports 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason code: 90495.
DMSQEFL	**Upwardly compatible:** • **[V6.4]** Changed values returned for the *cp_level* and *cms_level* parameters.
DMSQSFSL	**Upwardly compatible:** • **[V6.4]** Changed value returned for the *server_level* parameter.
DMSRDCAT	**Upwardly compatible:** • **[V2.2]** In the OBJECTCAT record for SFS: – New bit settings in FILEFLAGS field to support 4-digit years (century setting for DATE and DATEREF fields). – Reserved CHAR(1) field preceding the LAST_CHANGE_DATE field has changed to: **Field Name** **Field Type/Description** **CHGDATE_CENTURY** CHAR(1) Century byte for LAST_CHANGE_DATE – Reserved CHAR(1) field preceding the CREATIONDATE field has changed to: **Field Name** **Field Type/Description** **CREATIONDATE_CENTURY** CHAR(1) Century byte for CREATIONDATE
DMSSECIT	**INCOMPATIBLE:** • **[V5.3]** Only INIT or TERM is allowed as the value for the command keyword parameter (Parm 8). PROGCHK, TERM NORMAL, and TERM ABNORMAL are no longer allowed. For information about this routine, see *z/VM: CMS File Pool Planning, Administration, and Operation*.
DMSTRUNC	**Upwardly compatible:** • **[V2.2]** New parameters to specify date format: SHORTDATE, FULLDATE, ISODATE. • **[V2.2]** *date* parameter supports 4-digit years (10-character dates) when used with the FULLDATE and ISODATE parameters. • **[V2.2]** New reason code: 90495.

CMS Multitasking Routines

Table 42 lists CMS multitasking routines that have changed. For additional information, see *z/VM: CMS Application Multitasking*.

Table 42. Changes to CMS Multitasking Routines

Routine	Changes
DateTimeSubtract	**Upwardly compatible:** • **[V2.3]** Supports new date and time formats.

OpenExtensions Callable Services

Table 43 lists OpenExtensions callable services that have changed. For additional information, see *z/VM: OpenExtensions Callable Services Reference*.

Table 43. Changes to OpenExtensions Callable Services

Routine	Changes
close (BPX1CLO)	**Upwardly compatible:** • **[V3.1]** Supports sockets.
fcntl (BPX1FCT)	**Upwardly compatible:** • **[V3.1]** Supports sockets.
mount (BPX1MNT)	**Upwardly compatible:** • **[V3.1]** Mounts a local or remote file system.
openvmf (BPX1VM5)	**Upwardly compatible:** • **[V2.4]** New function code: VM5_RESOLVE_PATH.
read (BPX1RED)	**Upwardly compatible:** • **[V3.1]** Supports sockets.
wrote (BPX1WRT)	**Upwardly compatible:** • **[V3.1]** Supports sockets.

Systems Management Routines

In z/VM V6.1 and later, only the sockets-based systems management APIs are supported. The RPC-based systems management APIs are not supported.

CMS Compatibility-Interface Routines

Table 44 lists CMS compatibility-interface routines that have changed. For additional information, see *z/VM: CMS Application Development Guide for Assembler*.

Table 44. Changes to CMS Compatibility-Interface Routines

Routine	Changes
DMSTVS	**Upwardly compatible:** • **[V2.2]** New LIBSRV plist parameter. • **[V2.4]** New message DMS2139I is issued if SENSE data from the tape mount indicates that the mounted tape cartridge might be incorrect for the tape device in use.

CMS Macros

Changes to CMS macros are identified in the following topics:

- "General CMS Macros and Subcommands" on page 392
- "OpenExtensions Macros" on page 393
- "CMS Compatibility-Interface Macros" on page 394

• "CMS OS Simulation Macros and Supervisor Calls" on page 394

General CMS Macros and Subcommands

Table 45 lists general CMS macros and subcommands that have changed. For additional information, see *z/VM: CMS Macros and Functions Reference*.

Table 45. Changes to General CMS Macros and Subcommands

Macro or Subcommand	Changes
ADDENTRY macro	**INCOMPATIBLE:** • **[V5.2]** Not supported.
ADDENTRY subcommand	**INCOMPATIBLE:** • **[V5.2]** Not supported.
CMSLEVEL macro	**INCOMPATIBLE:** • **[V2.2]** Returned CMS level has been frozen at X'0F' (Level 12) for VM/ESA V2.1 and later. Use the new DMSQEFL macro or the DMSQEFL CSL routine instead.
CMSSTOR macro	**Upwardly compatible:** • **[V6.2]** New USERG subpool can be specified by a program running in a z/CMS virtual machine to obtain or release storage above 2 GB. • **[V6.2]** New return code 12 on CMSSTOR OBTAIN.
CPRB macro	**INCOMPATIBLE:** • **[V5.2]** Not supported.
CSMRETCD macro	**INCOMPATIBLE:** • **[V5.2]** Not supported.
DELENTRY macro	**INCOMPATIBLE:** • **[V5.2]** Not supported.
DELENTRY subcommand	**INCOMPATIBLE:** • **[V5.2]** Not supported.
DIRBUFF macro	**Upwardly compatible:** • **[V2.2]** FILE record contains the following new fields: DIRFDAXD, DIRFDAXC, DIRFDAXI, DIRFLV13. • **[V2.2]** FILEEXT record contains the following new fields: DIREDAXD, DIREDAXC, DIREDAXI, DIREDRXD, DIREDRXC, DIREDRXI, DIRECDXD, DIRECDXC, DIRECDXI, DIREDCXD, DIREDCXC, DIREDCXI, DIRELV13. • **[V2.2]** SEARCHALL and SEARCHAUTH records contain the following new fields: DIRSDAXD, DIRSDAXC, DIRSDAXI, DIRSCEND, DIRSLV13.
DMSQEFL macro	**Upwardly compatible:** • **[V6.4]** Changed value returned for the new CMS level.
EXSBUFF macro	**Upwardly compatible:** • **[V2.2]** FILE record contains the following new fields: EXSFDAXD, EXSFDAXC, EXSFDAXI, EXSFDRXD, EXSFDRXC, EXSFDRXI, EXSFCDXD, EXSFCDXC, EXSFCDXI, EXSFDCXD, EXSFDCXC, EXSFDCXI, EXSF2000, EXSFLV13. • **[V2.2]** In the FILE record, the following field has changed: EXSFRES. • **[V2.2]** DIR record contains the following new fields: EXSDDCXD, EXSDDCXC, EXSDDCXI, EXSDCDXD, EXSDCDXC, EXSDCDXI, EXSD2000, EXSDLV13. • In the DIR record, the following field has changed: EXSDRES.

Table 45. Changes to General CMS Macros and Subcommands (continued)

Macro or Subcommand	Changes
FSSTATE macro	**Upwardly compatible:** • **[V2.2]** In the FST flag byte, bit 4 indicates the century (first two digits of the year) the file was last written or updated (0=19*nn*, 1=20*nn*, where *nn* is the 2-digit year). Previously, this bit was not used.
FSTD macro	**Upwardly compatible:** • **[V2.2]** FSTFLAGS section contains the new FSTCNTRY field, which is a bit that indicates the century (first two digits of the year) the file was last written or updated (0=19*nn*, 1=20*nn*, where *nn* is the 2-digit year).
GETREQ subcommand	**INCOMPATIBLE:** • **[V5.2]** Not supported.
LT macro	**INCOMPATIBLE:** • **[V5.2]** A new DMSLT macro has been created (identical to LT) and used in CMS, GCS, and REXX/VM to avoid conflict with the LT opcode of the IBM High Level Assembler, V1.5. If you currently use the LT macro, IBM recommends that you change to using DMSLT. The LT macro is still supplied in DMSGPI MACLIB, but it could be withdrawn in a future release.
SENDREQ macro	**INCOMPATIBLE:** • **[V5.2]** Not supported.
SENDREQ subcommand	**INCOMPATIBLE:** • **[V5.2]** Not supported.
SETREPLY subcommand	**INCOMPATIBLE:** • **[V5.2]** Not supported.
TAPECTL macro	**Upwardly compatible:** • **[V2.2]** If the tape is under the control of a Tape Library Dataserver machine, and the DFSMS/VM Removable Media Services (RMS) FSMPPSI CSLLIB is available to CMS, the RUN (rewind and unload) function of TAPECTL calls the RMS FSMRMDMT (Demount) CSL routine to have the Dataserver unmount the tape.

OpenExtensions Macros

Table 46 lists OpenExtensions macros that have changed. For additional information, see *z/VM: OpenExtensions Callable Services Reference*.

Table 46. Changes to OpenExtensions Macros

Macro	Changes
BPXYERNO	**Upwardly compatible:** • **[V3.1]** Contains equates for new return codes and reason codes.
BPXYOPNF	**Upwardly compatible:** • **[V3.1]** New equates for new flag values for fcntl (BPX1FCT).
BPXYVM5	**Upwardly compatible:** • **[V2.4]** New equate for new VM5_RESOLVE_PATH function code for openvmf (BPX1VM5).

CMS Compatibility-Interface Macros

Table 47 lists CMS compatibility-interface macros that have changed.

Table 47. Changes to CMS Compatibility-Interface Macros

Macro	Changes

DEFNUC	**INCOMPATIBLE:**
	• **[V2.3]** In the default CMS IPL heading (used when the VERSION= parameter is specified without a value in DEFNUC) the date is presented in ISO format (*yyyy-mm-dd*).
	
	Upwardly compatible:
	• **[V6.2]** DEFNUC macro for the z/CMS nucleus is contained in the DMSZNGP file. IBM-supplied DEFNUC responses for z/CMS nucleus generation are identical to responses for ESA/390 CMS nucleus generation except:
	– SYSNAME=ZCMS
	– IPLADDR=990
	– Default version identification constructed for VERSION= is:
	`z/CMS Vv.r.m yyyy-mm-dd hh:mm`
	– Default output file heading constructed for INSTID= is:
	`z/Architecture CMS`
	• **[V6.4]** Default IPL heading entry contains the new release level.
	
	For additional information, see *z/VM: CMS Planning and Administration*.

CMS OS Simulation Macros and Supervisor Calls

Table 48 lists CMS OS Simulation macros that have changed. Table 49 on page 395 lists the CMS OS simulation supervisor calls that have changed. For additional information, see *z/VM: CMS Application Development Guide for Assembler*.

Table 48. Changes to CMS OS Simulation Macros

Macro	Changes
CLOSE	**Upwardly compatible:** • **[V2.4]** Can process OS formatted variable spanned QSAM records (under XLRI) or non-OS CMS files up to 65535 bytes in length.
DCB	**Upwardly compatible:** • **[V2.4]** Can describe and pass both LRI and XLRI conventions for QSAM variable spanned long records (up to 65535 bytes) for subsequent OPEN, CLOSE, GET, or PUT processing. Can also describe non-OS CMS files up to 65535 bytes in length.
GET	**Upwardly compatible:** • **[V2.4]** Can process QSAM I/O to variable spanned QSAM records or non-OS CMS files up to 65535 bytes in length.
OPEN	**Upwardly compatible:** • **[V2.4]** Can process OS formatted variable spanned QSAM records (under XLRI) or non-OS CMS files up to 65535 bytes in length.
PUT	**Upwardly compatible:** • **[V2.4]** Can process QSAM I/O to variable spanned QSAM records or non-OS CMS files up to 65535 bytes in length.

Table 48. Changes to CMS OS Simulation Macros (continued)

Macro	Changes
TIME	**Upwardly compatible:** • **[V2.2]** Second half-byte of the date format is a century indicator, where 0 indicates the 1900's, 1 indicates the 2000's, and 2 indicates the 2100's. This corresponds to the MVS implementation of the TIME macro.

Table 49. Changes to CMS OS Simulation Supervisor Calls

SVC	Changes
SVC 19 (OPEN)	**Upwardly compatible:** • **[V2.4]** Can process OS formatted variable spanned QSAM records (under XLRI) or non-OS CMS files up to 65535 bytes in length.
SVC 20 (CLOSE)	**Upwardly compatible:** • **[V2.4]** Can process OS formatted variable spanned QSAM records (under XLRI) or non-OS CMS files up to 65535 bytes in length.

HELP Facility

Table 50 identifies HELP components (HELP*xxxx* file types) that have been added or deleted in each release.

Table 51 on page 397 lists HELP files that have been deleted in each release. If a specific file name is not identified, all files with the specified file type have been deleted.

Note: Changes prior to z/VM V4.4 are not listed.

Both groups of changes are considered to be **INCOMPATIBLE**:

- Files for a new HELP component might conflict with local help files with identical file IDs.
- Files identified as deleted are removed from the new z/VM release but are not automatically removed from your system if you overlay your current help files with the new set. You need to remove them manually.

New files added to existing HELP components are not listed here because they are all upwardly compatible.

Table 50. Added or Deleted HELP Components

HELP Component	File Type	Change
CPOTHER	HELPCPOT	**[V4.4] Deleted** and all data merged into HELPCP component.
SOCKETS	HELPSOCK	**[V4.4] Deleted** (preferred functions now included in runtime for IBM C/C++ for z/VM).
TCP	HELPTCP	**[V4.4] Deleted** and replaced by TCPIP component (file type HELPTCPI).
DNS	HELPDNS	**[V4.4] Added** for TCP/IP DNS server administrative commands.
FCX	HELPFCX	**[V4.4] Added** for Performance Toolkit subcommands and field descriptions.
GDDMXD	HELPGDDM	**[V4.4] Added** for TCP/IP GDDMXD/VM interface subcommands.
IMAPADM	HELPIMAP	**[V4.4] Added** for TCP/IP IMAP server administrative commands.

Table 50. Added or Deleted HELP Components (continued)

HELP Component	File Type	Change
MPROUTE	HELPMPRO	**[V4.4] Added** for TCP/IP MPROUTE server administrative commands.
NFS	HELPNFS	**[V4.4] Added** for TCP/IP NFS server administrative commands.
SMTP	HELPSMTP	**[V4.4] Added** for TCP/IP SMTP server administrative commands.
SNMP	HELPSNMP	**[V4.4] Added** for TCP/IP SNMP administrative commands.
SSLADMIN	HELPSSLA	**[V4.4] Added** for TCP/IP SSL server administrative commands.
TCPIP	HELPTCPI	**[V4.4] Added** for TCP/IP commands and related functions.
TELNET	HELPTELN	**[V4.4] Added** for TCP/IP Telnet protocol client subcommands.
UFTD	HELPUFTD	**[V4.4] Added** for TCP/IP UFT server administrative (UFTD command) subcommands.
SPTAPE	HELPSPTA	**[V5.1] Deleted** (no longer supported).
SRPI	HELPSRPI	**[V5.2] Deleted** (no longer supported).
VMADMIN	HELPVMAD	**[V5.2] Deleted** (no longer supported).
IPFORMAT	HELPIPFO	**[V5.2] Added** for TCP/IP IPFORMAT subcommands.
LOCATE	HELPLOCA	**[V5.2] Added** for CP LOCATE command operands.
BOOTPD	HELPBOOT	**[V5.3] Deleted** (TCP/IP BOOTP server no longer supported).
RTE	HELPRTE	**[V5.3] Deleted** (TCP/IP RouteD server no longer supported).
LDAP	HELPLDAP	**[V5.3] Added** for TCP/IP LDAP server subcommands.
MSOCKETS	HELPMSOC	**[V5.3] Added** for CMS virtual systems management sockets.
RSCSAUTH	HELPRSCS	**[V5.3] Added** for RSCS Networking dynamic authorization commands.
DFSMS	HELPDFSM	**[V5.4] Added** for DFSMS/VM and DFSMS/VM RMS commands.
TFTD	HELPTFTD	**[V5.4] Deleted** (support withdrawn for TCP/IP TFTP server).
VMDT	HELPVMDT	**[V5.4] Added** for VM Dump Tool toolkit macros and stages (unsupported samples/examples).
X25IPI	HELPX25I	**[V5.4] Deleted** (support withdrawn for TCP/IP X.25 interface).
NSINTER	HELPNSIN	**[V6.1] Added** for TCP/IP NSLOOKUP command interactive session subcommands.
NSLOOKUP	HELPNSLO	**[V6.1] Added** for TCP/IP NSLOOKUP command subcommands.
DHCPD	HELPDHCP	**[V6.3] Deleted** (support withdrawn for TCP/IP DHCP server).
LPD	HELPLPD	**[V6.3] Deleted** (support withdrawn for TCP/IP LPD server).
MROUTINE	HELPMROU	**[V6.3] Deleted** (not supported).
XSPOOL	HELPXSPO	**[V6.3] Deleted** (support withdrawn for CSE).
ASSOCIATE	HELPASSO	**[V6.4 VM66029] Added** for CP ASSOCIATE command subcommands.
CREATE	HELPCREA	**[V6.4 VM66028] Added** for CMS CREATE command subcommands.
DEACTIVE	HELPDEAC	**[V6.4 VM66029] Added** for CP DEACTIVE command subcommands.
DELETE	HELPDELE	**[V6.4 VM66029] Expanded** to include both CP DELETE and CMS DELETE command subcommands.
DRAIN	HELPDRAI	**[V6.4 VM66029] Added** for CP DRAIN command subcommands.
FLASHCOPY	HELPFLAS	**[V6.4 VM66029] Added** for CP FLASHCOPY command subcommands.

Table 50. Added or Deleted HELP Components (continued)

HELP Component	File Type	Change
FREE	HELPFREE	**[V6.4 VM66029] Added** for CP FREE command subcommands.
GIVE	HELPGIVE	**[V6.4 VM66029] Added** for CP GIVE command subcommands.
HOLD	HELPHOLD	**[V6.4 VM66029] Added** for CP HOLD command subcommands.
SEGMENT	HELPSEGM	**[V6.4 VM66028] Added** for CMS SEGMENT command subcommands.
START	HELPSTAR	**[V6.4 VM66029] Added** for CP START command subcommands.
VARY	HELPVARY	**[V6.4 VM66029] Added** for CP VARY command subcommands.

Table 51. Deleted HELP Files

File Name	File Type	Release/Reason
CPOTHER	HELPABBR	**[V4.4]** Data merged into CP HELPABBR.
MACRO	HELPABBR	**[V4.4]** Replaced by MACROS HELPABBR.
ROUTINE	HELPABBR	**[V4.4]** Replaced by ROUTINES HELPABBR.
FTP	HELPCMS	**[V4.4]** Replaced by FTP HELPTCPI.
LPQ	HELPCMS	**[V4.4]** Replaced by LPQ HELPTCPI.
LPR	HELPCMS	**[V4.4]** Replaced by LPR HELPTCPI.
LPRM	HELPCMS	**[V4.4]** Replaced by LPRM HELPTCPI.
LPRSET	HELPCMS	**[V4.4]** Replaced by LPRSET HELPTCPI.
NETSTAT	HELPCMS	**[V4.4]** Replaced by NETSTAT HELPTCPI.
PING	HELPCMS	**[V4.4]** Replaced by PING HELPTCPI.
REXEC	HELPCMS	**[V4.4]** Replaced by REXEC HELPTCPI.
TELNET	HELPCMS	**[V4.4]** Replaced by TELNET HELPTCPI.
TFTP	HELPCMS	**[V4.4]** Replaced by TFTP HELPTCPI.
USER	HELPCP	**[V4.4]** Replaced by MESSAGEU HELPCP.
fn	HELPCPOT	**[V4.4]** Replaced by *fn* HELPCP files.
CFLINK	HELPCPQU	**[V4.4]** Replaced by CFLINKS HELPCPQU.
AMMOTHER	HELPMENU	**[V4.4]** Function no longer supported.
AMMR	HELPMENU	**[V4.4]** Function no longer supported.
COMMAND	HELPMENU	**[V4.4]** Replaced by COMMANDS HELPMENU.
CPOTHER	HELPMENU	**[V4.4]** Data merged into CP HELPMENU.
LANRES	HELPMENU	**[V4.4]** Function no longer supported.
MACRO	HELPMENU	**[V4.4]** Replaced by MACROS HELPMENU.
MENU	HELPMENU	**[V4.4]** Replaced by MENUS HELPMENU.
ROUTINE	HELPMENU	**[V4.4]** Replaced by ROUTINES HELPMENU.
SOCKETS	HELPMENU	**[V4.4]** Preferred functions now in runtime for C/C++ for z/VM.
HCQ031W	HELPMSG	**[V4.4]** Replaced by HCQ031E HELPMSG.
KEYS	HELPPF	**[V4.4]** Replaced by PFKEYS HELPHELP.
fn	HELPSOCK	**[V4.4]** Preferred functions now in runtime for C/C++ for z/VM.
AMMR	HELPTASK	**[V4.4]** Function no longer supported.
AVS	HELPTASK	**[V4.4]**

Table 51. Deleted HELP Files (continued)

File Name	File Type	Release/Reason
CMS	HELPTASK	**[V4.4]**
CP	HELPTASK	**[V4.4]**
DVF	HELPTASK	**[V4.4]**
GCS	HELPTASK	**[V4.4]**
LANRES	HELPTASK	**[V4.4]** Function no longer supported.
MACRO	HELPTASK	**[V4.4]** Replaced by MACROS HELPTASK.
OPTIONS	HELPTASK	**[V4.4]** Replaced by QUERYSET HELPTASK.
PIPELINE	HELPTASK	**[V4.4]**
ROUTINE	HELPTASK	**[V4.4]** Replaced by ROUTINES HELPTASK.
SOCKETS	HELPTASK	**[V4.4]** Preferred functions now in runtime for C/C++ for z/VM.
TSAF	HELPTASK	**[V4.4]**
VMDT	HELPTASK	**[V4.4]**
VMRM	HELPTASK	**[V4.4]**
VMSESE	HELPTASK	**[V4.4]**
VMSESEI	HELPTASK	**[V4.4]**
XSPOOL	HELPTASK	**[V4.4]**
IFCONFIG	HELPTCP	**[V4.4]** Replaced by IFCONFIG HELPTCPI.
OBEYFILE	HELPTCP	**[V4.4]** Replaced by OBEYFILE HELPTCPI.
CACHE	HELPTFTP	**[V4.4]** Replaced by CACHE HELPTFTD.
CLIENTS	HELPTFTP	**[V4.4]** Replaced by CLIENTS HELPTFTD.
CMS	HELPTFTP	**[V4.4]** Replaced by CMS HELPTFTD.
CREATION	HELPTFTP	**[V4.4]** Replaced by CREATION HELPTFTD.
DROPFILE	HELPTFTP	**[V4.4]** Replaced by DROPFILE HELPTFTD.
EXIT	HELPTFTP	**[V4.4]** Replaced by EXIT HELPTFTD.
LIGHT	HELPTFTP	**[V4.4]** Replaced by LIGHT HELPTFTD.
LOADPERM	HELPTFTP	**[V4.4]** Replaced by LOADPERM HELPTFTD.
LOADUSER	HELPTFTP	**[V4.4]** Replaced by LOADUSER HELPTFTD.
STAYUP	HELPTFTP	**[V4.4]** Replaced by STAYUP HELPTFTD.
STOP	HELPTFTP	**[V4.4]** Replaced by STOP HELPTFTD.
XFERMODE	HELPTFTP	**[V4.4]** Replaced by XFERMODE HELPTFTD.
SPTAPE	HELPABBR	**[V5.1]** Function no longer supported.
CONV2WD	HELPCMS	**[V5.1]** Function no longer supported.
MACLMIG	HELPCMS	**[V5.1]** Function no longer supported.
INSTDEF	HELPCP	**[V5.1]** Function no longer supported.
SPTAPE	HELPCP	**[V5.1]** Function no longer supported.
IOASSIST	HELPCPQU	**[V5.1]** Function no longer supported.
VR	HELPCPQU	**[V5.1]** Function no longer supported.
VRFREE	HELPCPQU	**[V5.1]** Function no longer supported.
CCWTRAN	HELPCPSE	**[V5.1]** Function no longer supported.
IOASSIST	HELPCPSE	**[V5.1]** Function no longer supported.

Table 51. Deleted HELP Files (continued)

File Name	File Type	Release/Reason
NOTRANS	HELPCPSE	**[V5.1]** Function no longer supported.
SVMLAN	HELPCPSE	**[V5.1]** Replaced by VMLAN HELPCPSE.
GUESTSTG	HELPDISP	**[V5.1]** Replaced by GUESTZ HELPDISP.
GUESTSTO	HELPDISP	**[V5.1]** Replaced by GUEST390 HELPDISP.
XC	HELPDISP	**[V5.1]** Replaced by GUESTXC HELPDISP.
GUESTSTG	HELPDUMP	**[V5.1]** Replaced by GUESTZ HELPDUMP.
GUESTSTO	HELPDUMP	**[V5.1]** Replaced by GUEST390 HELPDUMP.
STACK	HELPDUMP	**[V5.1]** Replaced by LKS HELPDUMP.
XC	HELPDUMP	**[V5.1]** Replaced by GUESTXC HELPDUMP.
SPTAPE	HELPMENU	**[V5.1]** Function no longer supported.
DVMCSSAA	HELPMROU	**[V5.1]** Replaced by DMSCSSAA HELPMROU.
DVMCSSAD	HELPMROU	**[V5.1]** Replaced by DMSCSSAD HELPMROU.
DVMCSSAQ	HELPMROU	**[V5.1]** Replaced by DMSCSSAQ HELPMROU.
DVMCSSAR	HELPMROU	**[V5.1]** Replaced by DMSCSSAR HELPMROU.
DVMCSSSC	HELPMROU	**[V5.1]** Replaced by DMSCSSSC HELPMROU.
DVMCSSSQ	HELPMROU	**[V5.1]** Replaced by DMSCSSSQ HELPMROU.
DVMCSSSR	HELPMROU	**[V5.1]** Replaced by DMSCSSSR HELPMROU.
DMS716E	HELPMSG	**[V5.1]**
DMS718E	HELPMSG	**[V5.1]**
DMS719E	HELPMSG	**[V5.1]**
DMS720E	HELPMSG	**[V5.1]**
DMS812E	HELPMSG	**[V5.1]**
DMS912I	HELPMSG	**[V5.1]**
DMS1126S	HELPMSG	**[V5.1]**
DMS1250E	HELPMSG	**[V5.1]**
DMS1255T	HELPMSG	**[V5.1]**
DMS2048W	HELPMSG	**[V5.1]**
FPL399E	HELPMSG	**[V5.1]**
HCP200E	HELPMSG	**[V5.1]**
HCP202I	HELPMSG	**[V5.1]**
HCP203E	HELPMSG	**[V5.1]**
HCP204E	HELPMSG	**[V5.1]**
HCP352W	HELPMSG	**[V5.1]** Replaced by HCP352E HELPMSG.
HCP407I	HELPMSG	**[V5.1]**
HCP408E	HELPMSG	**[V5.1]**
HCP412E	HELPMSG	**[V5.1]**
HCP416I	HELPMSG	**[V5.1]**
HCP648E	HELPMSG	**[V5.1]**
HCP722E	HELPMSG	**[V5.1]**
HCP723E	HELPMSG	**[V5.1]**

Table 51. Deleted HELP Files (continued)

File Name	File Type	Release/Reason
HCP724E	HELPMSG	**[V5.1]**
HCP811I	HELPMSG	**[V5.1]**
HCP813I	HELPMSG	**[V5.1]**
HCP820E	HELPMSG	**[V5.1]**
HCP825E	HELPMSG	**[V5.1]**
HCP828W	HELPMSG	**[V5.1]**
HCP893E	HELPMSG	**[V5.1]**
HCP1038E	HELPMSG	**[V5.1]**
HCP1039E	HELPMSG	**[V5.1]**
HCP1106I	HELPMSG	**[V5.1]**
HCP1357E	HELPMSG	**[V5.1]**
HCP1375I	HELPMSG	**[V5.1]**
HCP1450E	HELPMSG	**[V5.1]**
HCP1508I	HELPMSG	**[V5.1]**
HCP1511I	HELPMSG	**[V5.1]**
HCP1555E	HELPMSG	**[V5.1]**
HCP1556I	HELPMSG	**[V5.1]**
HCP1557E	HELPMSG	**[V5.1]**
HCP2161I	HELPMSG	**[V5.1]**
HCP2578E	HELPMSG	**[V5.1]**
HCP2601E	HELPMSG	**[V5.1]**
HCP2818I	HELPMSG	**[V5.1]** Replaced by HCP2818E HELPMSG.
HCP6055E	HELPMSG	**[V5.1]**
HCP6203E	HELPMSG	**[V5.1]**
HCP6232E	HELPMSG	**[V5.1]**
HCP6405E	HELPMSG	**[V5.1]**
HCP6541E	HELPMSG	**[V5.1]**
HCP6748E	HELPMSG	**[V5.1]**
HCP6749E	HELPMSG	**[V5.1]**
HCP6752I	HELPMSG	**[V5.1]**
HCP6766E	HELPMSG	**[V5.1]**
HCP6767E	HELPMSG	**[V5.1]**
HCP6788E	HELPMSG	**[V5.1]**
HCP6880E	HELPMSG	**[V5.1]**
HCP8501E	HELPMSG	**[V5.1]**
HCP9026W	HELPMSG	**[V5.1]**
HCP9027W	HELPMSG	**[V5.1]**
HCP9028W	HELPMSG	**[V5.1]**
HCP9029W	HELPMSG	**[V5.1]**
HCP9400I	HELPMSG	**[V5.1]**

Table 51. Deleted HELP Files (continued)

File Name	File Type	Release/Reason
HCP9401I	HELPMSG	**[V5.1]**
HCP9402E	HELPMSG	**[V5.1]**
HCP9403E	HELPMSG	**[V5.1]**
HCP9404I	HELPMSG	**[V5.1]**
HCP9405E	HELPMSG	**[V5.1]**
HCP9406I	HELPMSG	**[V5.1]**
HCP9407I	HELPMSG	**[V5.1]**
HCP9408E	HELPMSG	**[V5.1]**
HCP9409E	HELPMSG	**[V5.1]**
HCP9410I	HELPMSG	**[V5.1]**
HCP9411I	HELPMSG	**[V5.1]**
HCP9412E	HELPMSG	**[V5.1]**
HCP9413I	HELPMSG	**[V5.1]**
HCP9414I	HELPMSG	**[V5.1]**
HCP9415I	HELPMSG	**[V5.1]**
HCP9417I	HELPMSG	**[V5.1]**
HCP9418E	HELPMSG	**[V5.1]**
HCP9419E	HELPMSG	**[V5.1]**
HCP9420E	HELPMSG	**[V5.1]**
HCP9421E	HELPMSG	**[V5.1]**
HCP9422E	HELPMSG	**[V5.1]**
HCS144E	HELPMSG	**[V5.1]**
VMF2112I	HELPMSG	**[V5.1]** Replaced by VMF2112W HELPMSG.
fn	HELPSPTA	**[V5.1]** Function no longer supported.
GUESTSTG	HELPSTOR	**[V5.1]** Replaced by GUESTZ HELPSTOR.
GUESTSTO	HELPSTOR	**[V5.1]** Replaced by GUEST390 HELPSTOR.
XC	HELPSTOR	**[V5.1]** Replaced by GUESTXC HELPSTOR.
SRPI	HELPABBR	**[V5.2]** Function no longer supported.
VMADMIN	HELPABBR	**[V5.2]** Function no longer supported.
CMSSERV	HELPCMS	**[V5.2]** Function no longer supported.
LOCCMDBK	HELPCP	**[V5.2]** Replaced by CMDBK HELPLOCA.
LOCDGNBK	HELPCP	**[V5.2]** Replaced by DGNBK HELPLOCA.
LOCFILID	HELPCP	**[V5.2]** Replaced by FILID HELPLOCA.
LOCFRMTA	HELPCP	**[V5.2]** Replaced by FRAMETBL HELPLOCA.
LOCLDEV	HELPCP	**[V5.2]** Replaced by LDEV HELPLOCA.
LOCRDEV	HELPCP	**[V5.2]** Replaced by RDEV HELPLOCA.
LOCSHPBK	HELPCP	**[V5.2]** Replaced by SHPBK HELPLOCA.
LOCSNABK	HELPCP	**[V5.2]** Replaced by SNABK HELPLOCA.
LOCSPFBK	HELPCP	**[V5.2]** Replaced by SPFBK HELPLOCA.
LOCSTOR	HELPCP	**[V5.2]** Replaced by STORAGE HELPLOCA.

Table 51. Deleted HELP Files (continued)

File Name	File Type	Release/Reason
LOCSYM	HELPCP	**[V5.2]** Replaced by SYMBOL HELPLOCA.
LOCVDEV	HELPCP	**[V5.2]** Replaced by VDEV HELPLOCA.
LOCVMDBK	HELPCP	**[V5.2]** Replaced by VMDBK HELPLOCA.
LOCVSMBK	HELPCP	**[V5.2]** Replaced by VSMBK HELPLOCA.
LOCXITBK	HELPCP	**[V5.2]** Replaced by XITBK HELPLOCA.
VARYVECT	HELPCP	**[V5.2]** Function no longer supported.
VECTOR	HELPCPQU	**[V5.2]** Function no longer supported.
VIRTVECT	HELPCPQU	**[V5.2]** Function no longer supported.
VECTOR	HELPDEFI	**[V5.2]** Function no longer supported.
VECTOR	HELPDETA	**[V5.2]** Function no longer supported.
VECTOR	HELPDISP	**[V5.2]** Function no longer supported.
VECTOR	HELPDUMP	**[V5.2]** Function no longer supported.
VECTOR	HELPINDI	**[V5.2]** Function no longer supported.
ADDENTRY	HELPMACR	**[V5.2]** Function no longer supported.
CPRB	HELPMACR	**[V5.2]** Function no longer supported.
CSMRETCD	HELPMACR	**[V5.2]** Function no longer supported.
DELENTRY	HELPMACR	**[V5.2]** Function no longer supported.
SENDREQ	HELPMACR	**[V5.2]** Function no longer supported.
SRPI	HELPMENU	**[V5.2]** Function no longer supported.
VMADMIN	HELPMENU	**[V5.2]** Function no longer supported.
HCP162E	HELPMSG	**[V5.2]** Replaced by HCP162I HELPMSG.
HCP633E	HELPMSG	**[V5.2]** Function no longer supported.
HCP634E	HELPMSG	**[V5.2]** Function no longer supported.
HCP635E	HELPMSG	**[V5.2]** Function no longer supported.
HCP636E	HELPMSG	**[V5.2]** Function no longer supported.
HCP637E	HELPMSG	**[V5.2]** Function no longer supported.
HCP638E	HELPMSG	**[V5.2]** Function no longer supported.
HCP649E	HELPMSG	**[V5.2]** Function no longer supported.
HCP664I	HELPMSG	**[V5.2]** Function no longer supported.
HCP745I	HELPMSG	**[V5.2]** Function no longer supported.
HCP746E	HELPMSG	**[V5.2]** Function no longer supported.
HCP749E	HELPMSG	**[V5.2]** Function no longer supported.
HCP785E	HELPMSG	**[V5.2]** Function no longer supported.
HCP1011E	HELPMSG	**[V5.2]** Replaced by HCP1011I HELPMSG.
HCP1701E	HELPMSG	**[V5.2]** Function no longer supported.
HCP6864E	HELPMSG	**[V5.2]** Replaced by HCP6864I HELPMSG.
HCP9257I	HELPMSG	**[V5.2]** Function no longer supported.
HLE*xxxxx*	HELPMSG	**[V5.2]** Function no longer supported.
fn	HELPSRPI	**[V5.2]** Function no longer supported.
VECTOR	HELPSTOR	**[V5.2]** Function no longer supported.

Table 51. Deleted HELP Files (continued)

File Name	File Type	Release/Reason
SRPI	HELPTASK	**[V5.2]** Function no longer supported.
fn	HELPVMAD	**[V5.2]** Function no longer supported.
INDQ	HELPVMDU	**[V5.2]**
XEDIT	HELPVMDU	**[V5.2]**
BOOTP	HELPABBR	**[V5.3]** Function no longer supported.
BOOTPD	HELPABBR	**[V5.3]** Function no longer supported.
ROUTED	HELPABBR	**[V5.3]** Function no longer supported.
fn	HELPBOOT	**[V5.3]** Function no longer supported.
NIC	HELPCPQU	**[V5.3]** Replaced by VIRTNIC HELPCPQU. Accepted abbreviations are VNIC and NIC.
PAVALIAS	HELPCPQU	**[V5.3]** Replaced by VIRTPAV HELPCPQU. Accepted abbreviations are VPAV and PAVALIAS.
DNS	HELPDNS	**[V5.3]**
DNSHELP	HELPDNS	**[V5.3]**
CBAMONS	HELPFCX	**[V5.3]**
CBASRP	HELPFCX	**[V5.3]**
CGENEQUA	HELPFCX	**[V5.3]**
CGENFCXL	HELPFCX	**[V5.3]**
CGENFRMT	HELPFCX	**[V5.3]**
CPERGRM	HELPFCX	**[V5.3]**
CPERSUTL	HELPFCX	**[V5.3]**
CPRSORT	HELPFCX	**[V5.3]**
LIMIT	HELPFCX	**[V5.3]**
SFSMENU	HELPFCX	**[V5.3]**
BOOTP	HELPMENU	**[V5.3]** Function no longer supported.
BOOTPD	HELPMENU	**[V5.3]** Function no longer supported.
COMMANDS	HELPMENU	**[V5.3]**
ROUTED	HELPMENU	**[V5.3]** Function no longer supported.
DMS1302E	HELPMSG	**[V5.3]**
DMS2550E	HELPMSG	**[V5.3]**
DMS3508E	HELPMSG	**[V5.3]**
DMS3553E	HELPMSG	**[V5.3]**
FCX324A	HELPMSG	**[V5.3]**
FCX325A	HELPMSG	**[V5.3]**
FCX122E	HELPMSG	**[V5.3]**
FCX537I	HELPMSG	**[V5.3]**
FCX589E	HELPMSG	**[V5.3]**
FCX611E	HELPMSG	**[V5.3]**
FCX900X	HELPMSG	**[V5.3]**
FCX901X	HELPMSG	**[V5.3]**

Table 51. Deleted HELP Files (continued)

File Name	File Type	Release/Reason
GCT103I	HELPMSG	**[V5.3]**
HCP6406E	HELPMSG	**[V5.3]**
HCP6724E	HELPMSG	**[V5.3]** Replaced by HCP6724W HELPMSG.
HCP6758E	HELPMSG	**[V5.3]** Replaced by HCP6758W HELPMSG.
HCP6774E	HELPMSG	**[V5.3]** Replaced by HCP6774W HELPMSG.
HCP6864E	HELPMSG	**[V5.3]**
HCP8350R	HELPMSG	**[V5.3]**
HCP8355I	HELPMSG	**[V5.3]**
HCP8411I	HELPMSG	**[V5.3]**
SYMBOLS	HELPROUT	**[V5.3]**
fn	HELPRTE	**[V5.3]** Function no longer supported.
BOOTPD	HELPTASK	**[V5.3]** Function no longer supported.
FEATURES	HELPTASK	**[V5.3]**
ROUTED	HELPTASK	**[V5.3]** Function no longer supported.
BOOTPD	HELPTCPI	**[V5.3]** Function no longer supported.
ROUTED	HELPTCPI	**[V5.3]** Function no longer supported.
TFTD	HELPABBR	**[V5.4]** Support withdrawn for TCP/IP TFTP server.
X25IPI	HELPABBR	**[V5.4]** Support withdrawn for TCP/IP X.25 interface.
TFTPD	HELPMENU	**[V5.4]** Support withdrawn for TCP/IP TFTP server.
X25IPI	HELPMENU	**[V5.4]** Support withdrawn for TCP/IP X.25 interface.
HCP2837I	HELPMSG	**[V5.4]**
HCP9256I	HELPMSG	**[V5.4]**
HCS087E	HELPMSG	**[V5.4]** Replaced by HCS087I HELPMSG.
HCS620I	HELPMSG	**[V5.4]** Function no longer supported.
HCS621I	HELPMSG	**[V5.4]** Function no longer supported.
HCS622E	HELPMSG	**[V5.4]** Function no longer supported.
HCS623I	HELPMSG	**[V5.4]** Function no longer supported.
HCS624I	HELPMSG	**[V5.4]** Function no longer supported.
HCS640I	HELPMSG	**[V5.4]** Function no longer supported.
HCS1103E	HELPMSG	**[V5.4]**
HCS1602I	HELPMSG	**[V5.4]** Replaced by HCS1602E HELPMSG.
HCS1603I	HELPMSG	**[V5.4]** Replaced by HCS1603E HELPMSG.
DELETE	HELPSSLA	**[V5.4]** Function no longer supported.
EXPORT	HELPSSLA	**[V5.4]** Function no longer supported.
LOGCLEAR	HELPSSLA	**[V5.4]** Function no longer supported.
LOGSIZE	HELPSSLA	**[V5.4]** Function no longer supported.
NOHALT	HELPSSLA	**[V5.4]** Function no longer supported.
REGENER	HELPSSLA	**[V5.4]** Function no longer supported.
REMOVE	HELPSSLA	**[V5.4]** Function no longer supported.
REQUEST	HELPSSLA	**[V5.4]** Function no longer supported.

Table 51. Deleted HELP Files (continued)

File Name	File Type	Release/Reason
SELF	HELPSSLA	**[V5.4]** Function no longer supported.
STORE	HELPSSLA	**[V5.4]** Function no longer supported.
NDB	HELPTASK	**[V5.4]** Support withdrawn for TCP/IP NDB interface.
TFTPADM	HELPTASK	**[V5.4]** Support withdrawn for TCP/IP TFTP server.
X25IPI	HELPTASK	**[V5.4]** Support withdrawn for TCP/IP X.25 interface.
NDBCLNT	HELPTCPI	**[V5.4]** Support withdrawn for TCP/IP NDB interface.
PORTCLNT	HELPTCPI	**[V5.4]** Support withdrawn for TCP/IP NDB interface.
PORTSRVS	HELPTCPI	**[V5.4]** Support withdrawn for TCP/IP NDB interface.
TFTPD	HELPTCPI	**[V5.4]** Support withdrawn for TCP/IP TFTP server.
fn	HELPTFTD	**[V5.4]** Support withdrawn for TCP/IP TFTP server.
fn	HELPX25I	**[V5.4]** Support withdrawn for TCP/IP X.25 interface.
CRYPTO	HELPCPSE	**[V6.1]** Function not supported.
CRYPTO	HELPDEFI	**[V6.1]** Function not supported.
CRYPTO	HELPDETA	**[V6.1]** Function not supported.
CRYPTO	HELPDISP	**[V6.1]** Function not supported.
CRYPTO	HELPDUMP	**[V6.1]** Function not supported.
CRYPTO	HELPSTOR	**[V6.1]** Function not supported.
HMF	HELPMENU	**[V6.1]** HMF/VM not supported.
HMFQUERY	HELPMENU	**[V6.1]** HMF/VM not supported.
HMFSET	HELPMENU	**[V6.1]** HMF/VM not supported.
PRODUCTS	HELPMENU	**[V6.1]**
HCP2468E	HELPMSG	**[V6.1]**
HCP9051W	HELPMSG	**[V6.1]**
DNS	HELPABBR	**[V6.2]** Function not supported.
HSA	HELPCPQU	**[V6.2]** Command removed.
CLOSECON	HELPDNS	**[V6.2]** Function not supported.
COMMIT	HELPDNS	**[V6.2]** Function not supported.
DUMP	HELPDNS	**[V6.2]** Function not supported.
FLIPTABL	HELPDNS	**[V6.2]** Function not supported.
HELP	HELPDNS	**[V6.2]** Function not supported.
HINTS	HELPDNS	**[V6.2]** Function not supported.
LEVEL	HELPDNS	**[V6.2]** Function not supported.
LIST	HELPDNS	**[V6.2]** Function not supported.
PURGE	HELPDNS	**[V6.2]** Function not supported.
REFRESH	HELPDNS	**[V6.2]** Function not supported.
STATS	HELPDNS	**[V6.2]** Function not supported.
STORAGE	HELPDNS	**[V6.2]** Function not supported.
TRACE	HELPDNS	**[V6.2]** Function not supported.
VMDUMP	HELPDNS	**[V6.2]** Function not supported.
CPERTFTP	HELPFCX	**[V6.2]** Function not supported.

Table 51. Deleted HELP Files (continued)

File Name	File Type	Release/Reason
FPERTFTP	HELPFCX	**[V6.2]** Function not supported.
DNS	HELPMENU	**[V6.2]** Function not supported.
DVH3297	HELPMSG	**[V6.2]**
DVH3299	HELPMSG	**[V6.2]**
DVH3300	HELPMSG	**[V6.2]**
DVH3367	HELPMSG	**[V6.2]**
DVH3368	HELPMSG	**[V6.2]**
DVH3392	HELPMSG	**[V6.2]**
FCX080E	HELPMSG	**[V6.2]**
FCX352E	HELPMSG	**[V6.2]**
FCX650I	HELPMSG	**[V6.2]**
FCX651I	HELPMSG	**[V6.2]**
FCX652I	HELPMSG	**[V6.2]**
FCX653I	HELPMSG	**[V6.2]**
FCX654I	HELPMSG	**[V6.2]**
FCX655I	HELPMSG	**[V6.2]**
FCX659I	HELPMSG	**[V6.2]**
FCX660E	HELPMSG	**[V6.2]**
HCP645E	HELPMSG	**[V6.2]** Function not supported.
HCP659E	HELPMSG	**[V6.2]** Function not supported.
HCP660E	HELPMSG	**[V6.2]** Function not supported.
HCP661E	HELPMSG	**[V6.2]** Function not supported.
HCP662I	HELPMSG	**[V6.2]** Function not supported.
HCP663E	HELPMSG	**[V6.2]** Function not supported.
HCP781E	HELPMSG	**[V6.2]**
HCP1702E	HELPMSG	**[V6.2]** Function not supported.
HCP1705I	HELPMSG	**[V6.2]** Function not supported.
HCP1706I	HELPMSG	**[V6.2]** Function not supported.
HCP1707I	HELPMSG	**[V6.2]** Function not supported.
HCP1708I	HELPMSG	**[V6.2]** Function not supported.
HCP1709E	HELPMSG	**[V6.2]** Function not supported.
HCP1710E	HELPMSG	**[V6.2]** Function not supported.
HCP1711I	HELPMSG	**[V6.2]** Function not supported.
HCP1712I	HELPMSG	**[V6.2]** Function not supported.
HCP1713I	HELPMSG	**[V6.2]** Function not supported.
HCP1716E	HELPMSG	**[V6.2]** Function not supported.
HCP1761I	HELPMSG	**[V6.2]**
HCP1881E	HELPMSG	**[V6.2]**
HCP2723E	HELPMSG	**[V6.2]**
HCP2744E	HELPMSG	**[V6.2]**

Table 51. Deleted HELP Files (continued)

File Name	File Type	Release/Reason
HCP6423E	HELPMSG	**[V6.2]**
HCQ039E	HELPMSG	**[V6.2]**
HCQ200I	HELPMSG	**[V6.2]**
HCQ401I	HELPMSG	**[V6.2]**
IRM038E	HELPMSG	**[V6.2]**
XSP001	HELPMSG	**[V6.2]** Function not supported.
XSP100	HELPMSG	**[V6.2]** Function not supported.
XSP101	HELPMSG	**[V6.2]** Function not supported.
XSP102	HELPMSG	**[V6.2]** Function not supported.
XSP110	HELPMSG	**[V6.2]** Function not supported.
XSP111	HELPMSG	**[V6.2]** Function not supported.
XSP112	HELPMSG	**[V6.2]** Function not supported.
XSP113	HELPMSG	**[V6.2]** Function not supported.
XSP114	HELPMSG	**[V6.2]** Function not supported.
XSP115	HELPMSG	**[V6.2]** Function not supported.
XSP116	HELPMSG	**[V6.2]** Function not supported.
DNSADM	HELPTASK	**[V6.2]** Function not supported.
KERBADM	HELPTASK	**[V6.2]** Function not supported.
KERBEROS	HELPTASK	**[V6.2]** Function not supported.
EXT_SRVT	HELPTCPI	**[V6.2]** Function not supported.
KADMIN	HELPTCPI	**[V6.2]** Function not supported.
KDB_DEST	HELPTCPI	**[V6.2]** Function not supported.
KDB_EDIT	HELPTCPI	**[V6.2]** Function not supported.
KDB_INIT	HELPTCPI	**[V6.2]** Function not supported.
KDB_UTIL	HELPTCPI	**[V6.2]** Function not supported.
KDESTROY	HELPTCPI	**[V6.2]** Function not supported.
KINIT	HELPTCPI	**[V6.2]** Function not supported.
KLIST	HELPTCPI	**[V6.2]** Function not supported.
KPASSWD	HELPTCPI	**[V6.2]** Function not supported.
KSTASH	HELPTCPI	**[V6.2]** Function not supported.
NSMAIN	HELPTCPI	**[V6.2]** Function not supported.
GATHER	HELPXSPO	**[V6.2]** Function not supported.
DHCPD	HELPABBR	**[V6.3]** Support withdrawn for TCP/IP DHCP server.
LPD	HELPABBR	**[V6.3]** Support withdrawn for TCP/IP LPD server.
MROUTINE	HELPABBR	**[V6.3]** Not supported.
XSPOOL	HELPABBR	**[V6.3]** Support withdrawn for CSE.
CLUSTER	HELPCPQU	**[V6.3]** Not used.
CLUSTER	HELPCPSE	**[V6.3]** Not used.
OVERRIDE	HELPCP	**[V6.3]** Support withdrawn for user class restructure (UCR).
XSPOOL	HELPCP	**[V6.3]** Support withdrawn for CSE.

Table 51. Deleted HELP Files (continued)

File Name	File Type	Release/Reason
fn	HELPDHCP	**[V6.3]** Support withdrawn for TCP/IP DHCP server.
CP	HELPLPD	**[V6.3]** Support withdrawn for TCP/IP LPD server.
DHCPD	HELPMENU	**[V6.3]** Support withdrawn for TCP/IP DHCP server.
LPD	HELPMENU	**[V6.3]** Support withdrawn for TCP/IP LPD server.
MROUTINE	HELPMENU	**[V6.3]** Not supported.
XSPOOL	HELPMENU	**[V6.3]** Support withdrawn for CSE.
fn	HELPMROU	**[V6.3]** Not supported.
DMS159T	HELPMSG	**[V6.3]** Not used.
DMS160T	HELPMSG	**[V6.3]** Not used.
DMS161T	HELPMSG	**[V6.3]** Not used.
DMS260T	HELPMSG	**[V6.3]** Not used.
DMS340E	HELPMSG	**[V6.3]** Not used.
DMS625S	HELPMSG	**[V6.3]** Not used.
DMS717I	HELPMSG	**[V6.3]** Not used.
DMS789E	HELPMSG	**[V6.3]** Not used.
DMS902T	HELPMSG	**[V6.3]** Not used.
DMS1077S	HELPMSG	**[V6.3]** Not used.
DMS1136W	HELPMSG	**[V6.3]** Not used.
DMS1407I	HELPMSG	**[V6.3]** Not used.
DMS2017I	HELPMSG	**[V6.3]** Not used.
DMS2017S	HELPMSG	**[V6.3]** Not used.
DMS2245I	HELPMSG	**[V6.3]** Not used.
DMS2413W	HELPMSG	**[V6.3]** Not used.
DMS2803S	HELPMSG	**[V6.3]** Not used.
HCP1800E	HELPMSG	**[V6.3]** Not used.
HCP8300*s*-8499*s*	HELPMSG	**[V6.3]** Installation messages now have format IUG####*s*.
HCP9154E	HELPMSG	**[V6.3]** Not used.
DHCP	HELPTASK	**[V6.3]** Support withdrawn for TCP/IP DHCP server.
LPDADM	HELPTASK	**[V6.3]** Support withdrawn for TCP/IP LPD server.
DHCPD	HELPTCPI	**[V6.3]** Support withdrawn for TCP/IP DHCP server.
LPD	HELPTCPI	**[V6.3]** Support withdrawn for TCP/IP LPD server.
fn	HELPXSPO	**[V6.3]** Support withdrawn for CSE.
DVH3261I	HELPMSG	**[V6.4]** Expanded storage (XSTORE) support removed.
DVH3324E	HELPMSG	**[V6.4]** Expanded storage (XSTORE) support removed.
HCP1108E	HELPMSG	**[V6.4]** Expanded storage (XSTORE) support removed.
HCP1400I	HELPMSG	**[V6.4]** Expanded storage (XSTORE) support removed.
HCP1402I	HELPMSG	**[V6.4]** Expanded storage (XSTORE) support removed.
HCP1403E	HELPMSG	**[V6.4]** Expanded storage (XSTORE) support removed.
HCP1703I	HELPMSG	**[V6.4]** Not used.
HCP1704I	HELPMSG	**[V6.4]** Not used.

Table 51. Deleted HELP Files (continued)

File Name	File Type	Release/Reason
HCP1962I	HELPMSG	**[V6.4]** Expanded storage (XSTORE) support removed.
HCP2266E	HELPMSG	**[V6.4]** Not used.
HCP2267E	HELPMSG	**[V6.4]** Not used.
HCP3013E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3014E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3015E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3016I	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3017I	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3018E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3019I	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3020E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3021E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3024E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3025E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3026E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3027E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3028E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3029E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3030E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3035E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3038E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP3172I	HELPMSG	**[V6.4]** Not used.
HCP3179E	HELPMSG	**[V6.4]** Not used.
HCP3182I	HELPMSG	**[V6.4]** Not used.
HCP3186I	HELPMSG	**[V6.4]** Not used.
HCP6153E	HELPMSG	**[V6.4]** Not used.
HCP6154E	HELPMSG	**[V6.4]** Not used.
HCP6254I	HELPMSG	**[V6.4]** Not used.
HCP6294I	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
HCP6625I	HELPMSG	**[V6.4]** Not used.
HCP6872I	HELPMSG	**[V6.4]** Expanded storage (XSTORE) support removed.
HCP9111I	HELPMSG	**[V6.4]** Expanded storage (XSTORE) support removed.
HCP9153E	HELPMSG	**[V6.4]** Not used.
VSM1000E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
VSM1001E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
VSM1002E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
VSM1003E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
VSM1004E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
VSM1005W	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
VSM1006E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.

Table 51. Deleted HELP Files (continued)

File Name	File Type	Release/Reason
VSM1007E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
VSM1008W	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
VSM1009E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
VSM1010W	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
VSM1011E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.
VSM1012E	HELPMSG	**[V6.4]** Ensembles/zManager support removed.

Dump Viewing Facility Changes

Table 52 lists Dump Viewing Facility functions that have changed. For additional information, see *z/VM: Dump Viewing Facility*.

Table 52. Changes to Dump Viewing Facility Functions

Function	Changes
DUMPSCAN	**INCOMPATIBLE:** • **[V3.1]** DUMPSCAN cannot be used for CP dumps. Use the VM Dump Tool. • **[V3.1]** The following subcommands are no longer supported: CPEBK, FINDCPE, FINDUSER, FRAMETBL, INSPECT, REAL, RIOBLOK, SELECT, SNAPLIST, TRSAVE, VIOBLOK, VIRT, VMDBK, XTRACE. • **[V3.1]** Header record indicates new product name. **Upwardly compatible:** • **[V6.4]** Header record indicates new release level.

GCS Changes

Table 53 lists GCS commands and macros that have changed. For additional information, see *z/VM: Group Control System*.

Table 53. Changes to GCS Commands and Macros

Interface	Changes
GCSLEVEL macro	**Upwardly compatible:** • **[V6.4]** Contains new equates for new release levels.
GETMAIN macro	**INCOMPATIBLE:** • **[V2.2]** Previously, although LOC=RES was documented as the default, the actual default was LOC=BELOW, and all the requested virtual storage was allocated below 16 MB. The actual default has been changed to LOC=RES. If the requester resides above 16 MB, virtual storage might be allocated anywhere. If you have any programs that invoke GETMAIN with the default, you must make sure they can accommodate addresses above 16 MB, or you must recode them to invoke GETMAIN with LOC=BELOW.
QUERY GCSLEVEL command	**Upwardly compatible:** • **[V6.4]** Response contains changed values for the new release.
QUERY MODDATE command	**INCOMPATIBLE:** • **[V2.2]** Full 4-digit year is now returned in the date field of the response instead of a 2-digit year.

Table 53. Changes to GCS Commands and Macros (continued)

Interface	Changes
TIME macro	**Upwardly compatible:** • **[V2.2]** Second half-byte of the date format is a century indicator, where 0 indicates the 1900's, 1 indicates the 2000's, and 2 indicates the 2100's. This corresponds to the MVS implementation of the TIME macro.

REXX/VM Changes

Table 54 lists REXX/VM instructions, functions, and external functions that have changed. For additional information, see *z/VM: REXX/VM Reference*.

Table 54. Changes to REXX/VM Instructions, Functions, and External Functions

Interface	Changes
CMSFLAG	**Upwardly compatible:** • **[V2.2]** New value for *flag*: YEAR2000.
DATE	**Upwardly compatible:** • **[V2.2]** New parameters allow you to specify a date to be converted to a different format. • **[V2.4]** New parameters: *output_separator_char*, *input_separator_char*.
DIAG / DIAGRC	**Upwardly compatible:** • **[V2.2]** New DIAGNOSE code is supported: X'270'.
STORAGE	**Upwardly compatible:** • **[V6.2]** Accepts 16-digit hexadecimal address, allowing a program running in z/CMS to inspect and modify storage above 2 GB.

VMSES/E Changes

Table 55 lists VMSES/E commands that have changed. For additional information, see *z/VM: VMSES/E Introduction and Reference*.

[V4.4]

The CMSSEGS BLDDATA file has been renamed to ZVMSEGS BLDDATA and Language Environment component segments have been added.

Table 55. Changes to VMSES/E Commands

Command	Changes
GENCPBLS	**INCOMPATIBLE:** • **[V5.1]** Default has changed from DUAL to NODUAL. **Upwardly compatible:** • **[V3.1]** New options: ALTCNTRL, DUAL, NODUAL, LOADLIST, ALTLOADLIST, PREEXIT.
ITNVTSTR	**Upwardly compatible:** • **[V2.3]** New operands: KEY, PROD.

Table 55. Changes to VMSES/E Commands (continued)

Command	Changes
LOCALMOD	**Upwardly compatible:** • **[V5.2]** New operands: *infn*, *inft*, REWORK, MODID, CNTRL, VVTFT, ASM F, ASM H, ASM HL. • **[V5.2]** New messages: VMF1306E, VMF1307E, VMF1308E, VMF1309E, VMF1310E. • **[V5.4]** New operand: LIST. • **[V5.4]** New options: REMOVE, PREFIX. • **[V5.4]** New messages: VMF392E, VMF1325E, VMF1326E.
PRODUTL	**INCOMPATIBLE:** • **[V6.2]** This command, previously called TCP2PROD EXEC, has been renamed PRODUTL EXEC.
PUT2PROD	**Upwardly compatible:** • **[V4.3]** New message: VMF1219E. • **[V4.4]** Console messages placed into VMSES/E-formatted $VMFP2P $MSGLOG file. • **[V5.4]** New message: VMF1224E. • **[V6.2]** New operands: SAVECMS and SEGMENTS.
SERVICE	**Upwardly compatible:** • **[V4.3]** New operands: BUILD, ENABLE, DISABLE. • **[V4.3]** New messages: VMF1220W, VMF2308W, VMF2310I, VMF2771E, VMF2774I. • **[V4.4]** Console messages placed into VMSES/E-formatted $VMFSRV $MSGLOG file. • **[V5.1]** New operands: BITMAP, STATUS. • **[V5.3]** New operands: NORECEIVE, BLDNUC, TESTAPPLY. • **[V5.4]** New operand: BUILD LIST. • **[V5.4]** New message: VMF1224E.
VMFAPPLY	**Upwardly compatible:** • **[V3.1]** New option: PREEXIT.
VMFASM	**Upwardly compatible:** • **[V3.1]** New option: PREEXIT.
VMFBDNUC	**Upwardly compatible:** • **[V6.4]** New messages: VMF2251I, VMF2252I, VMF2253E, VMF2253W, VMF2254W, VMF2255W, VMF2256W
VMFBLD	**Upwardly compatible:** • **[V3.1]** New options: PREEXIT, ULOG, WILD.
VMFEXUPD	**Upwardly compatible:** • **[V3.1]** New options: PREEXIT.
VMFHASM	**Upwardly compatible:** • **[V3.1]** New options: PREEXIT.
VMFHLASM	**INCOMPATIBLE:** • **[V5.1]** Default has changed from DUAL to NODUAL. **Upwardly compatible:** • **[V3.1]** New options: ALTCNTRL, DUAL, NODUAL, LOADLIST, ALTLOADLIST, PREEXIT.

Table 55. Changes to VMSES/E Commands (continued)

Command	Changes
VMFINS (in general)	**Upwardly compatible:** • **[V2.3]** New operands: DISABLE, ENABLE. • Also see specific VMFINS commands listed below.
VMFINS DELETE	**INCOMPATIBLE:** • **[V6.2]** Deleted options: RESOURCE, DFNAME, DFTYPE, DFMODE. **Upwardly compatible:** • **[V2.3]** New operand: DISABLE.
VMFINS INSTALL	**INCOMPATIBLE:** • **[V6.2]** Deleted options: RESOURCE, DFNAME, DFTYPE, DFMODE. **Upwardly compatible:** • **[V2.3]** New operands: DISABLE, ENABLE, NOSETUP, SETUP.
VMFINS MIGRATE	**INCOMPATIBLE:** • **[V6.2]** Deleted options: RESOURCE, DFNAME, DFTYPE, DFMODE. **Upwardly compatible:** • **[V2.3]** New operands: DISABLE, ENABLE.
VMFMRDSK	**Upwardly compatible:** • **[V3.1]** New options: PREEXIT.
VMFNLS	**Upwardly compatible:** • **[V3.1]** New options: PREEXIT.
VMFREC	**Upwardly compatible:** • **[V3.1]** New options: PREEXIT.
VMFREM	**Upwardly compatible:** • **[V3.1]** New option: PREEXIT. • **[V4.4]** New operand: MOD. • **[V4.4]** New option: PREFIX. • **[V4.4]** UNAPPLY and UNRECEIVE options can be used for local modifications.
VMFREPL	**Upwardly compatible:** • **[V3.1]** New options: LOG, LOGLVL, MOD, PREEXIT, PREFIX, PTF.
VMFSETUP	**INCOMPATIBLE:** • **[V6.2]** Default changed from NOLINK to LINK. **Upwardly compatible:** • **[V6.3]** New option: LINKRr • **[V2.3]** New options: NOCONS, NOPROMPT, PROMPT.
VMFSIM (in general)	**Upwardly compatible:** • See specific VMFSIM commands listed below.
VMFSIM GETLVL	**Upwardly compatible:** • **[V3.1]** New options: MOD, PTF.

Table 55. Changes to VMSES/E Commands (continued)

Command	Changes
VMFSUFIN	**Upwardly compatible:** • **[V2.4]** New operand: ALL. • **[V2.4]** New options: CORTAPE, RSUTAPE. • **[V4.3]** New option: BUILD. • **[V4.3]** New message: VMF2308W • **[V5.3]** New operands: NOREC and TEST.
VMFSUFTB	**Upwardly compatible:** • **[V4.3]** New operand: SUFFN *fn*.
VMFUPDAT	**Upwardly compatible:** • **[V4.3]** New messages: VMF2300E, VMF2301W, VMF2302E, VMF2303W, VMF2304E, VMF2306W, VMF2307W, VMF2309W. • **[V4.4]** New operands: SRVBLDS and SYSREST. • **[V4.4]** Changes to VMFUPDAT Function Selection Panel. • **[V4.4]** New panels: SRVBLDS Update Panel and SYSREST Update Panel. • **[V5.1]** New operand: SYSMEMO. • **[V5.1]** New panel: SYSMEMO Update Panel. • **[V5.1]** Changed panel: SYSREST Update Panel — added PF key for Browse. • **[V5.1]** Changed panel: VMFUPDAT Function Selection Panel — added SYSMEMO. • **[V5.3]** Changed table: Update SYSSUF Table Entries — added prodlev column. • **[V6.2]** New option: SYSPINV.
VMFVIEW	**INCOMPATIBLE:** • **[V6.3]** Default PF2 and PK14 key assignments have been changed. **Upwardly compatible:** • **[V2.4]** New operands (all have the same function): REMOVE, VMFREM, $VMFREM. • **[V4.2]** New operands (all have the same function): *xxx*, VMF*xxx*, $VMF*xxx*.

Chapter 4. Migration Tasks

Attention: If you will use or have used the upgrade installation procedure to upgrade to z/VM V6.4 or later, then the following topics will not apply to your system. See *z/VM: Installation Guide* for information on tasks you may need to perform after you have completed your upgrade installation.

If you are upgrading your system to a z/VM V6.4 or later system that you installed using the traditional installation procedure, the following topics provide guidance for various tasks and procedures that might apply to your migration.

Converting from HCPRIO, HCPSYS, and HCPBOX to Configuration Files

If you are still using the HCPSYS, HCPRIO, and HCPBOX ASSEMBLE files to define your current VM system, you should convert to using configuration files **before you migrate to the new z/VM system**. Configuration files:

- Provide more flexibility.

  When system definition information is moved out of the CP nucleus into configuration files, much of that information can be changed dynamically using CP commands without having to rebuild CP.
- Support the latest system functions and capabilities.

  Since configuration files were established as the recommended method for defining a VM system, only the configuration files have been enhanced to support new system functions and capabilities. The ASSEMBLE files have not been updated to support those functions.

The HCPSYS and HCPRIO definitions can be replaced by statements in the system configuration file (usually called SYSTEM CONFIG). The HCPSYS ASSEMBLE file included in the CP module on the new z/VM system contains only a SYSEND macroinstruction. The HCPRIO ASSEMBLE file included in the CP module on the new z/VM system contains only an RIOGEN CONS=DYNAMIC macroinstruction.

The system logo definitions in HCPBOX can be replaced by statements in the logo configuration file (usually called LOGO CONFIG). See "Creating a Logo Configuration File" on page 420.

z/VM: CP Planning and Administration describes how to set up and use the system configuration file and logo configuration file, including how to define the *parm disk* on which the configuration files reside.

There are three ways that you can migrate your HCPSYS and HCPRIO data to system configuration statements:

- Create the SYSTEM CONFIG file and system configuration statements manually.
- Use the HCPTSYS and HCPTRIO sample utility programs to read your current HCPSYS and HCPRIO files and create system configuration statements. See "Using the HCPTRIO and HCPTSYS Utilities" on page 419.
- Use the HCPDCON sample utility program to create system configuration statements from your current running system. You can also use the HCPRDEVS sample utility program to create real I/O device definitions from your running system. See "Using the HCPDCON Utility" on page 420 and "Using the HCPRDEV Utility" on page 420.

Note: The HCPTSYS, HCPTRIO, HCPDCON, and HCPRDEVS sample utility programs were supplied with z/VM V4.4 and earlier releases. **They are not included with z/VM V5.1 or later.**

No matter which method you choose, you will have to make some decisions and adjustments. After you create the SYSTEM CONFIG file and make your changes, re-IPL CP. Because the system configuration statements override the HCPSYS and HCPRIO macros that were used in building the CP nucleus, there is no need to rebuild CP. This allows you to migrate your definitions gradually. At any time you can change back to using the HCPSYS and HCPRIO files by renaming the SYSTEM CONFIG file on the parm disk.

"Configuration Statements that Replace HCPSYS Macroinstructions" identifies system configuration statements that provide functions equivalent to macroinstructions in HCPSYS. "Configuration Statements that Replace HCPRIO Macroinstructions" on page 418 identifies system configuration statements that provide functions equivalent to macroinstructions in HCPRIO. These are not complete lists of the supported system configuration statements. Many additional configuration statements exist that provide capabilities for which there are no corresponding HCPSYS or HCPRIO macroinstructions.

For more information about what you can specify in the system configuration file, rules for coding the system configuration file, and complete descriptions of all the supported system configuration statements, see *z/VM: CP Planning and Administration*.

Configuration Statements that Replace HCPSYS Macroinstructions

Table 56 shows the macroinstructions in the HCPSYS ASSEMBLE file that should be migrated to system configuration statements. In some cases, information from one macroinstruction must be migrated into more than one system configuration statement.

Table 56. System Configuration Statements that Replace HCPSYS Macros

HCPSYS Macro	Corresponding System Configuration Statement
CSELDEV	XLINK_DEVICE_DEFAULTS
CSELVOL	XLINK_VOLUME_INCLUDE XLINK_VOLUME_EXCLUDE

Table 56. System Configuration Statements that Replace HCPSYS Macros (continued)

HCPSYS Macro	Corresponding System Configuration Statement
CSESYS	XLINK_SYSTEM_INCLUDE XLINK_SYSTEM_EXCLUDE **Note:** CSE cross-system spool is not supported in z/VM V6.3 and later, but cross-system link (XLINK) is still supported for non-SSI systems. The XLINK_SYSTEM_INCLUDE statement identifies one system, and includes a SLOT operand for defining the position of that system in the list of systems participating in cross-system link.
CSETRACE	None. CSE cross-system spool is not supported in z/VM V6.3 and later.
CSEUSER	None. CSE cross-system spool is not supported in z/VM V6.3 and later.
SYSACNT	SYSTEM_USERIDS
SYSADDIN	CP_ADDON_INITIALIZE_ROUTINES
SYSCPVOL	CP_OWNED **Notes:** 1. The list of CP-owned volumes generated by CP_OWNED statements completely supersedes the list of CP-owned volumes generated by the SYSCPVOL macroinstruction. 2. There are two differences between using the SYSCPVOL macroinstruction and the CP_OWNED statement: a. In the SYSCPVOL macroinstruction, the order in which the volumes are specified is significant. In the CP_OWNED statement, each volume is defined with a specific slot number, so the order of the statements is not significant. b. With the SYSCPVOL macroinstruction, if you need to add volumes to the CP-owned volume list on a running system, you have two choices: 1) Update the SYSCPVOL macroinstruction and IPL your system again. 2) Plan ahead — code volumes you do not have into the SYSCPVOL macroinstruction, and when you need to add a new volume, make sure the 1-character to 6-character volume identifier (*volid*) on the SYSCPVOL macroinstruction matches the volume identifier of the new volume. When you use the CP_OWNED statement, adding a new volume is easier: 1) Use the RESERVED operand to save one or more empty slots. 2) Use the CP DEFINE CPOWNED command to add the new volume. 3) If this is a permanent addition to the list, update your system configuration file so the new volume is automatically included in the CP-owned list then next time you IPL the system.
SYSDUMP	SYSTEM_USERIDS
SYSEREP	SYSTEM_USERIDS
SYSEXCL	USER_VOLUME_EXCLUDE
SYSFCN	PRIV_CLASSES
SYSFORM	USERFORM FORM_DEFAULT
SYSID	SYSTEM_IDENTIFIER
SYSINCL	USER_VOLUME_INCLUDE
SYSJRL (PSUPRS)	FEATURES
SYSJRL (other)	JOURNALING
SYSMAXU	FEATURES

Table 56. System Configuration Statements that Replace HCPSYS Macros (continued)

HCPSYS Macro	Corresponding System Configuration Statement
SYSOPR	SYSTEM_USERIDS
SYSPCLAS	PRINTER_TITLE
SYSRES (SYSCLR)	FEATURES
SYSRES (other)	SYSTEM_RESIDENCE
SYSSYMP	SYSTEM_USERIDS
SYSTIME	TIMEZONE_DEFINITION **Note:** There is one difference between the SYSTIME macroinstruction and the TIMEZONE_DEFINITION statement. Using the SYSTIME macroinstruction allows you to define only one time zone, while the TIMEZONE_DEFINITION statement allows you to define an unlimited number of time zones. However, only one time zone can be active at any one time.
SYSTORE	STORAGE **Note:** The SYSTORE macroinstruction has the RMSIZE parameter that lets you define how much real storage CP should allocate for itself. However, CP now uses the maximum amount of available storage, unless you specify a real storage amount on the RMSIZE parameter of the SYSTORE macroinstruction or you specify the STORE=*nnnn*M parameter of the Stand-Alone Program Loader (SAPL). For more information about SAPL and passing IPL parameters, see *z/VM: System Operation*. Because CP uses the maximum amount of available storage, there is no need for an operand on the STORAGE statement that is equivalent to the RMSIZE parameter of the SYSTORE macroinstruction.
SYSUVOL	USER_VOLUME_SET

Configuration Statements that Replace HCPRIO Macroinstructions

Table 57 shows the macroinstructions in the HCPRIO ASSEMBLE file that should be migrated to system configuration statements.

Table 57. System Configuration Statements that Replace HCPRIO Macros

HCPRIO Macro	Corresponding System Configuration Statement
RDEVICE (ICA)	RDEVICE (ICA)
RDEVICE (SHARED)	DEVICES
RDEVICE (UIRATE)	HOT_IO_RATE **Note:** There are two differences between the UIRATE parameter of the RDEVICE macroinstruction and the HOT_IO_RATE statement: 1. Both the UIRATE parameter and the HOT_IO_RATE statement let you specify the maximum unsolicited interrupt rate for a specific device, but the HOT_IO_RATE statement also lets you set the rate for: • All devices • An entire class of devices • A range of real device numbers. 2. You can use the CP QUERY HOTIO command to display hot I/O rates defined in the system configuration file, but you cannot use this command to display hot I/O rates that you defined using the UIRATE parameter of the RDEVICE macroinstruction.
RDEVICE (other)	RDEVICE **Note:** Only real I/O devices that are too old to be sensed or that require more information must be defined with RDEVICE statements.

Table 57. System Configuration Statements that Replace HCPRIO Macros (continued)

HCPRIO Macro	Corresponding System Configuration Statement
RIOGEN	OPERATOR_CONSOLES **Note:** You can use the EMERGENCY_MESSAGE_CONSOLES statement to define a list of console addresses that CP notifies when there is a system emergency.
SYSFCN	PRIV_CLASSES

Using the HCPTRIO and HCPTSYS Utilities

The HCPTRIO and HCPTSYS sample utility programs convert HCPRIO and HCPSYS definitions to system configuration statements. HCPTRIO and HCPTSYS are REXX programs and can run on any level of VM. The programs read your HCPRIO and HCPSYS files and create files containing statements that you can use in a SYSTEM CONFIG file.

Notes:

1. HCPTRIO and HCPTSYS do not process DMKRIO and DMKSYS files.
2. By default, HCPTRIO examines the HCPRIO ASSEMBLE file on the first accessed CMS minidisk and generates a list of statements for those real devices that do not answer sense ID requests or that do not return enough information. By default, HCPTRIO creates a file called RDEV CONFIG on the first CMS minidisk accessed in R/W mode.
3. By default, HCPTSYS examines the HCPSYS ASSEMBLE file on the first accessed CMS minidisk and generates a list of statements for your system that define its system characteristics (such as accounting, journaling, and time zone) and CSE status. By default, HCPTSYS creates a file called SYSTEM CONFIG on the first CMS minidisk accessed in R/W mode.
4. The statements created by HCPTRIO and HCPTSYS may not define a complete SYSTEM CONFIG file. Because HCPTRIO and HCPTSYS create statements from your existing HCPRIO and HCPSYS macroinstructions, they may not create some new configuration statements that are necessary for IPL. For example, HCPTRIO and HCPTSYS will not create a LOGO_CONFIG statement. You must create that statement, and perhaps others, manually in the SYSTEM CONFIG file. For detailed information about configuration file statements and the minimum contents of the SYSTEM CONFIG file, see *z/VM: CP Planning and Administration*.

In the following example, system configuration statements are created in the RIO CONFIG A file from the macros contained in the file HCPRIO ASSEMBLE D:

```
hcptrio rio config a from hcprio assemble d
```

In this example, the HCPTSYS command creates SYS CONFIG A from HCPSYS ASSEMBLE E:

```
hcptsys sys config a from hcpsys assemble e
```

For more information, see “HCPTRIO” on page 456 and “HCPTSYS” on page 459.

To use the statements created by HCPTRIO and HCPTSYS, add them to the SYSTEM CONFIG file on the parm disk. Make any other changes you desire or need and re-IPL CP.

Using the HCPDCON Utility

The HCPDCON sample utility program examines your running system and generates a file of configuration statements. By default, HCPDCON creates a file named SYSTEM CONFIG on the first CMS minidisk accessed in R/W mode. For more information, see "HCPDCON" on page 452.

Before using the SYSTEM CONFIG file created by HCPDCON, make sure that the file contain all the necessary statements and definitions. For detailed information about configuration file statements and the minimum contents of the SYSTEM CONFIG file, see *z/VM: CP Planning and Administration*. Then copy the SYSTEM CONFIG file to the parm disk and re-IPL CP.

Using the HCPRDEV Utility

The HCPRDEV sample utility program creates a system configuration-like file that contains statements for all real devices on your system that do not answer a sense ID request or do not return enough information. (These devices must be defined to the system.) When you select this function of HCPRDEV, by default it creates a file named RDEVS CONFIG on the first CMS minidisk accessed in R/W mode. For more information, see "HCPRDEVS" on page 454.

Difference with Multiple Systems Using Shared Data

When using CP configurability support, you may see DASD attached to systems other than your new system. For example, your installation can have a large set of DASD that is shared by several systems. All the systems' IOCPs could be set up for all the DASD, but each system actually uses only some of the DASD.

Because CP configuration senses all the DASD you have defined in your IOCP, you can see DASD not attached to your new system. To have your new system ignore DASD not attached to it, identify each one on a USER_VOLUME_EXCLUDE configuration statement.

Creating a Logo Configuration File

You can use the logo configuration file to override all information specified in HCPBOX ASSEMBLE. You can define an alternate source for logo picture files without changing HCPBOX. You can use statements in the logo configuration file to choose logo pictures for logical devices, SNA terminals, and locally-attached terminals. You can also use the logo configuration file to define the contents of the following fields:

- Command area
- Input area at the bottom of the logo screen
- Online message at the top of each logo screen
- Status area

You can use the DRAWLOGO sample utility program to create logo screens for your system. With this utility, you can edit the text of the logo file using XEDIT and modify the 3270 screen attributes in a logo file. DRAWLOGO creates these logo files on the first CMS minidisk accessed in R/W mode. By default, the file type is LOGO.

For more information about creating a logo configuration file and using the DRAWLOGO utility, see *z/VM: CP Planning and Administration*.

Migrating Spool Files and Saved Segments

This section provides you with some options for migrating your spool files, including saved segments, from your old system to your new system.

Note: No matter which option you choose, you should use the SPXTAPE command to create a backup of all the spool files and saved segments on your old system.

Depending on how you plan to migrate to the new system, you can use one of the methods shown in Table 58 to migrate your old system's spool files and saved segments. These methods are described in the following sections.

Table 58. Methods for Migrating Spool Files and Saved Segments

Method	When you can use this method
Using the same warm start and checkpoint areas	• If you are ready to cut over completely from your old system to your new system • If you want to migrate all your spool files and saved segments at once
Using SPXTAPE	• If you are ready to cut over completely from your old system to your new system • If you want to use your new system as a test system before cutting over completely • If you want to migrate all your spool files and saved segments at once • If you want to migrate selected spool files or saved segments
Building saved segments individually	• Anytime

Using the Same Warm Start and Checkpoint Areas to Migrate All Spool Files at Once

If you are ready to cut over completely from your old system to a new production system, you can set up your new system to use the same warm start and checkpoint areas as your old system. Then, when you IPL your new system with the warm start option, all of the spool files, including your saved segments, will be known to your new system.

The following procedure describes the situation where you are currently running the new z/VM release on a test system.

1. Define the warm start area and spool areas for your new production system to be the same as the areas used on your old production system. Define these areas in the CP_OWNED and SYSTEM_RESIDENCE statements in SYSTEM CONFIG. List all DASD from your old system first, just as it is defined. **Make sure that the DASD slot numbers for the spool file volumes are identical between the old production system and the new production system.** Then list any additional volumes directly after.

   Attention: Do not IPL the new production system with the new SYSTEM CONFIG file while the old production system is still up and running.

2. Use SPXTAPE to dump all of your spool files to tape from your old production system. This is for backup purposes. Make sure that no other users are on the system. They could create new spool files that do not get backed up in this step.
3. On your old production system, purge any saved segments, saved systems, or other spool files that you do not want on your new production system.
4. If you will need any of the spool files or saved segments from the test system on your new production system (for example, the segments for the new levels of CMS and GCS), use SPXTAPE to dump the spool files and system data files from the test system to tape.
5. Shut down your old production system.
6. IPL your new production system using the warm start option. Consider using the NOAUTOLOG option in response to the START prompt to prevent any service machines from coming up with the wrong level of CMS.

   Note: The spool files of users not known to your new system but who had spool files in your old system are now owned by the OPERATOR user ID.
7. Use SPXTAPE to load the spool files and saved segments from your test system (that you dumped to tape in step 4) on to your new production system.
8. To see what saved segments and saved systems your new system has, enter:

   ```
   q nss
   ```

   You may see some saved segments or saved systems with the same name or some that you no longer want or need on your new system.

   Using the spool ID, purge any saved segments or saved systems that are duplicates or that you no longer want or need. Do not use the names of the saved segments or saved systems to purge them. Using the name may cause you to purge a saved segment you meant to keep.
9. Either shutdown and reIPL your new system, or enter the following command to start up all of your system's service machines:

   ```
   xautolog autolog1
   ```

   Attention: Do not IPL your old production system. Because it is defined with the same checkpoint and warm start areas as your new system, it can corrupt your new system's spool files.
10. Consider using the VMFSGMAP command to further map your system's saved segment layout and set up VMSES/E to build saved segments. See "Building Saved Segments Individually" on page 423 for a brief description of VMFSGMAP. For further details on managing saved segments see *z/VM: Saved Segments Planning and Administration*.

Using SPXTAPE to Migrate Your Spool Files

You can use the SPXTAPE command to migrate all of your spool files and saved segments, but this method is generally used for migrating selected files. For example, if you want to test out particular saved segments on your new system before bringing them into production, you can migrate only the spool files that contain those segments.

1. On your old system, use SPXTAPE to dump to tape the spool files you want to migrate to your new system, including saved segments and saved systems.
2. IPL your new system. Make sure you have enough spool space allocated to contain the spool files you intend to load.

3. On your new system, use SPXTAPE to load the spool files you dumped in step 1 on page 422.
4. To see what saved segments and saved systems your new system has, enter:

   ```
   q nss
   ```

   You may see some saved segments or saved systems with the same name or some that you no longer want or need on your new system.

   Using the spool ID, purge any saved segments or saved systems that are duplicates or that you no longer want or need. Do not use the names of the saved segments or saved systems to purge them. Using the name may cause you to purge a saved segment you meant to keep.
5. Consider using the VMFSGMAP command to further map your system's saved segment layout and to set up VMSES/E to build saved segments. See "Building Saved Segments Individually" for a brief description of VMFSGMAP. For further details on managing saved segments see *z/VM: Saved Segments Planning and Administration*.

Building Saved Segments Individually

If you prefer, you can rebuild each individual saved segment you need on your new system. The VMFSGMAP command provides a saved segment mapping and planning interface and allows you to define saved segments to VMSES/E. You can then use the VMFBLD command to build saved segments on your system. *z/VM: Saved Segments Planning and Administration* provides details on saved segments, including how to use VMFSGMAP to map and manipulate saved segment layouts and how to use VMFBLD to build saved segments on your system.

To plan and set up your saved segment layout in the new system:

1. Identify all of the products or applications in your new system that require saved segments.
2. Collect all of the saved segment definitions (that is, default DEFSEG commands) for each of the products or applications that will use saved segments in your new system.

   For products, this information will most likely be in the installation information for each product. You need to gather this information for your system's own applications as well. The type of information you would need is shown in Table 59.

Table 59. Example of Saved Segment Information

Type of information needed	Your segment information
Product	z/VM
Segment Name	CMSBAM
Segment Type	Physical
Space Name	DOSBAM
Default Location	B10-B3F
Size (Pages)	30 hex pages (48 decimal)
Run Above 16 MB?	No
Build Tool	VMFBLD

Depending on how you plan to migrate to your new system, you may gather this information in two ways:

- Gradually, as you install or move each product or application

- During your migration planning prior to installing any products

Although it may require more time to plan for your migration, you should consider mapping your entire saved segment layout at once. Otherwise, as you gradually install additional products, you may have to remap and rebuild saved segments on your system.

As you gather saved segment information for your system's products and applications, you may need to take note of additional information about each saved segment that you plan to use in your new system:

- Whether the saved segment is a CMS logical segment or a CP physical segment
- Whether the saved segment can reside above the 16 MB line
- What execs or commands are used to build and save the segment

Note: There are two situations when you may not need to gather the additional information:

- Depending on how your new system and associated products were packaged, a number of saved segments may already be defined on your system. If this is true, do not collect the information for these segments manually. Later, when you use VMFSGMAP, you can pull saved segment information for these saved segments directly into the VMFSGMAP segment mapping tool using the SEGMERGE macro. You can input any additional segment information later if necessary.
- Some products may be VMSES/E-installed. In this case, these products may already have their default saved segment definitions identified for VMFSGMAP; so you do not need to collect this information manually. When you refer to each individual product's installation manual, the manual should clearly identify whether the product's saved segments are already defined for VMFSGMAP.

The information you gather will be useful to you later when you build your saved segments using VMFBLD. You provide this information to VMSES/E using the VMFSGMAP interface.

3. Use VMFSGMAP to add saved segment definitions or to plan for and out your system's saved segment layout.

   With VMFSGMAP, you provide specific information about each saved segment in your system:

   - Where it should reside in storage (the DEFSEG statement)
   - The name of the saved segment and segment space name if it applies
   - Whether it can reside above the 16 MB line
   - Whether CMS logical segments reside in the saved segment
   - What disks must be accessed for the saved segment to be properly built
   - Instructions for how to build and save the saved segment

   Note: Some packaged systems and any products that are VMSES/E enabled already have this information, as well as any other applicable information, identified for you. In general, for these saved segments you update only information about the placement of a saved segment or of a segment space in which the saved segment resides.

   Using VMFSGMAP, you can map and manipulate your system's saved segment layout without affecting your running system. How to invoke VMFSGMAP, use its panel interface, and map out and manipulate saved segment layouts is fully described in *z/VM: Saved Segments Planning and Administration*.

When manipulating your system's saved segment layout:

- Remember that CMS uses storage locations from 15 MB to 20 MB. Be sure not to place any saved segments in that area.
- Make sure that you adequately plan for saved segments that must be run below the 16 MB line.
- Be careful not to overlay saved segments that require one another.

4. Build and save the saved segments you need on your system.

   Saved segments that have been completely defined to VMSES/E with the VMFSGMAP interface can be built using VMFBLD, regardless of whether the saved segment belongs to a product that is completely serviced with VMSES/E. And, for any products or applications that are serviced with VMSES/E, you will be notified whenever service to that product or application requires that a saved segment must be rebuilt. For products or applications that are not VMSES/E-enabled, you can manually provide the necessary information and use VMFBLD to build the saved segment.

   For example, to build the CMSPIPES saved segment, you would enter:

   ```
   vmfbld ppf segbld esasegs segblist cmspipes ( all
   ```

   For examples of how to use VMFBLD to build saved segments on your system, see *z/VM: Saved Segments Planning and Administration*.

Avoiding the Loss of Spool Files and System Data Files During Migration

Changing how you allocate SPOOL space on your new system may cause you to lose spool files, which can include system data files. For example, you may decide to remove a cylinder or extent of SPOOL space during the migration of your system and make it PERM space for a minidisk. When the minidisk owner formats that space, any spool file that was chained through a page that was changed from SPOOL space to PERM space will then be destroyed. To avoid this potential loss of data:

1. Back up your spool files and system data files using SPXTAPE DUMP before reallocating SPOOL space.
2. Reallocate the SPOOL space.
3. Purge your system data files from the old system.
4. Shutdown the old system. **Do NOT use SHUTDOWN REIPL.**
5. IPL the new system with a cold start.
6. Restore the backed up spool files and system data files using SPXTAPE LOAD.

Sharing Data between Your Old System and Your New System

If you plan to migrate users from your old system to the new system in a staged fashion, your production will be split between several systems. One major concern is how you will share data between these systems.

z/VM allows you to share information on a DASD volume in the following ways:

- Among multiple virtual machines using virtual reserve/release.
- Among one virtual machine and operating systems running on other processors using real reserve/release.
- Among multiple virtual machines and operating systems running on other processors using concurrent virtual and real reserve/release. The virtual machines and operating systems must support reserve/release CCWs.

The following sections include examples of using virtual reserve/release and concurrent virtual and real reserve/release. For a complete discussion, see the information about DASD sharing in *z/VM: CP Planning and Administration*.

Reserve/Release Considerations for VSE

z/VM supports virtual reserve/release for minidisks that are not a full pack. Therefore, the cross-system communication (also called the "lock file") volume does not have to be defined as a full pack.

MDISK statements for all DASD you want to mount to VSE as shared (in other words, you want to use the S operand of the IPL ADD statement) must include the V suffix on the link mode. That is, the link mode must be MWV. If this is not done, VSE issues MSG0I23I for the minidisks that do not have link mode MWV on their MDISK statements.

Specifying MWV does not result in any additional overhead because z/VM does not do a reserve/release to any pack unless the guest asks it to. VSE only does a reserve/release to the cross-system communication file (the "lock file") after IPL.

Note that if the cross-system communication file (the "lock file") is shared by more than one CPU, SHARED must be YES on the RDEVICE statement in the system configuration file. Also, for sharing a volume concurrently between real and virtual machines, the volume must be defined as a full-pack minidisk.

Note: z/VM supports virtual reserve/release for the virtual disks in storage function. Virtual disks in storage are temporary FBA minidisks simulated in system storage rather than mapped to real DASD. Therefore, a virtual disk in storage may be faster than other minidisks because it avoids the overhead of I/O operations. VSE guests may benefit from this function by using a virtual disk in storage instead of a permanent minidisk to store label information areas and the cross-system communication file (the "lock file"). The virtual disk in storage function may be used by a guest running any supported version or release of VSE.

What to Do if Reserve/Release Cannot Be Used

In some instances, you will not benefit from or will not be able to share DASD through reserve/release. For instance, if:

- You need to share data among multiple virtual machines on multiple systems and the operating systems running in the virtual machines do not support reserve/release CCWs. CMS is an example of a virtual machine that does not support this type of sharing.
- Your system is constrained and you cannot afford the performance degradation that results from shared DASD (particularly concurrent virtual and real reserve/release).
- You need to share databases (such as SQL/DS or DB2 Server for VM) between several groups of users and not all of the users can be moved to the new system.
- There are incompatibilities that prevent your old system and new system from accessing the same data.

In these cases, you might:

- Replicate the data.
- Maintain the data on only one system and give two user IDs to users who need to access the data.

- Physically attach the DASDs you want to share to the second system and allow write access from only one system. If you have RACF, you can use it to find out who has access and to limit the write access from one system. You should be aware, however, that this method of sharing is not protected by z/VM; you must set up the controls yourself.
- Use the cross-system link (XLINK) support in z/VM. For more information, see *z/VM: CP Planning and Administration*.

Sharing Data among Multilevel Virtual Machines

If you want to share data among virtual machines running on the new system that support reserve/release CCWs, and you do not need to share this data with operating systems running on other processors, use virtual reserve/release. "Using Virtual Reserve/Release" explains how to do this. A virtual machine in which MVS is running is an example of a virtual machine that supports reserve/release CCWs.

If you want to share data among virtual machines running on the new system, and the virtual machines do not support reserve/release CCWs, you cannot use virtual reserve/release. A virtual machine in which CMS is running is an example of a virtual machine that does *not* support reserve/release CCWs. "Without Using Virtual Reserve/Release" on page 428 shows you how to share data between virtual machines that do not support virtual reserve/release.

Using Virtual Reserve/Release

If you want to share data among virtual machines running on the new system that support reserve/release CCWs, and you do not need to share this data with operating systems running on other processors, use virtual reserve/release.

Virtual reserve/release works the same way on the new system as it does on your old system. Figure 1 depicts virtual machines containing MVS that are sharing DASD through virtual reserve/release.

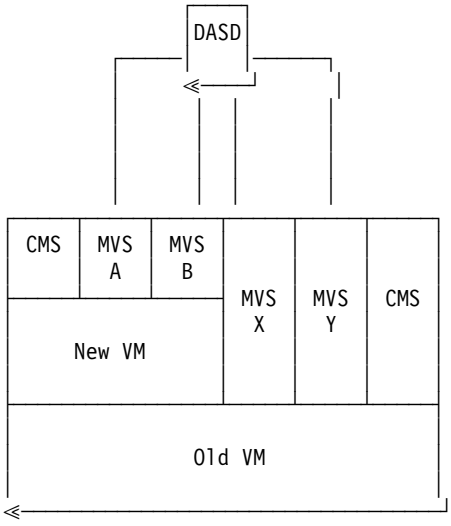

Figure 1. Sharing DASD Using Virtual Reserve/Release

To set up virtual reserve/release, it has to look like concurrent virtual and real reserve/release to the guest, which in the scenario depicted in Figure 1 is your new system.

Do the following on the new (guest) system:

1. Define the DASD as a minidisk for one of the virtual machines, perhaps MVS A, in the new system directory. Specify MWV as the access mode on the MDISK statement. To the second level guest, which is your new z/VM system, this minidisk must be viewed as a full-pack minidisk.
2. Use the SHARED YES operand on the RDEVICE statement in your SYSTEM CONFIG file on the new system guest for this minidisk. For example:

   ```
   Rdevice 0cc0 Type Dasd Shared yes
   ```

   Or use the SET SHARED command on the new system guest for this minidisk. For example:

   ```
   set shared on for 0cc0
   ```
3. Code the LINK statement for the remaining virtual machine, MVS B, in the new system directory. Specify MW as the access mode.

Do the following on the old (first-level) system:

1. Define the DASD where the minidisk resides.
2. Code the MDISK statement in your new system's directory entry in the old system's directory. You must append a V to the primary access mode (read, write, multiple write, and so on) indicating that this minidisk can be shared between virtual machines. For example:

   ```
   MDISK 197 3390 000 400 WORKPK MWV ORANGE
   ```
3. Code the LINK statement in the directory entries for MVS X and MVS Y. For example:

   ```
   LINK NEWESA 197 197 MW
   ```
4. Specify that the DASD will not be shared with another operating system. The default setting of the SHARED option of the RDEVICE statement (SHARED=NO) takes care of this for you.

Now, virtual machines running on your new system (the second-level system) may have write access to the same information as virtual machines running on your old system (the first level system).

Without Using Virtual Reserve/Release

If you want to share data among virtual machines running on the new system, and the virtual machines do not support reserve/release CCWs, such as CMS, you cannot use virtual reserve/release. You define the DASD as a minidisk, giving read/write access to only one of the CMS users running on the new system guest; the others can have only read access. This configuration is shown in Figure 2 on page 429.

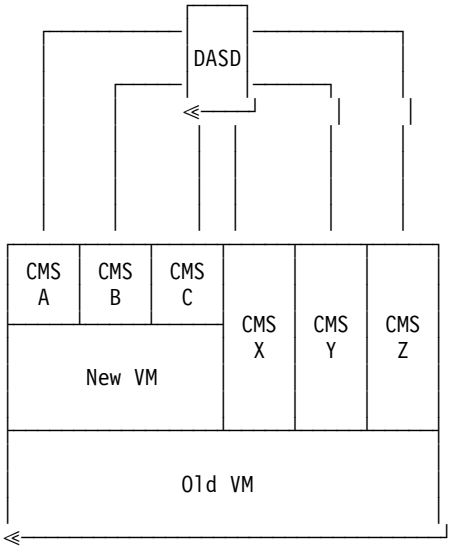

Figure 2. Sharing DASD without Using Virtual Reserve/Release

Do the following on the new (guest) system:

1. Define the DASD as a minidisk for one of the virtual machines, perhaps CMS A, in the new system directory. Specify MW as the access mode on the MDISK statement.
2. Code the LINK statement for the remaining virtual machines, CMS B and CMS C, in the new system directory. Specify RR as the access mode.

Do the following on the old (first-level) system:

1. Define the DASD where the minidisk resides.
2. Code the MDISK statement in your new system's directory entry in the old system's directory. Do not append a V to the primary access mode. For example:

   ```
   MDISK 197 3390 000 400 WORKPK MW ORANGE
   ```

3. Code the LINK statement in the directory entries for CMS X, CMS Y, and CMS Z. For example:

   ```
   LINK NEWESA 197 197 MW
   ```

4. Specify that the DASD will not be shared with another operating system. The default setting of the SHARED option of the RDEVICE statement (SHARED=NO) takes care of this for you.

Now, virtual machines running in the new system may have read access to the same information as virtual machines running on the old system.

Sharing Data among Virtual Machines and Other Systems

To share data among multiple virtual machines running on the new system and on other (older) systems, use concurrent virtual and real reserve/release support. You can do this, however, only if the virtual machines that are sharing the data support reserve/release CCWs. In the example shown in Figure 3 on page 430, the virtual machines containing MVS on the new system can share DASD in this manner with the virtual machine containing MVS on the old system. The virtual machines containing CMS cannot share DASD because CMS does not support reserve/release CCWs. You have to replicate the data for the CMS users or physically attach the DASD and give write access to only one system.

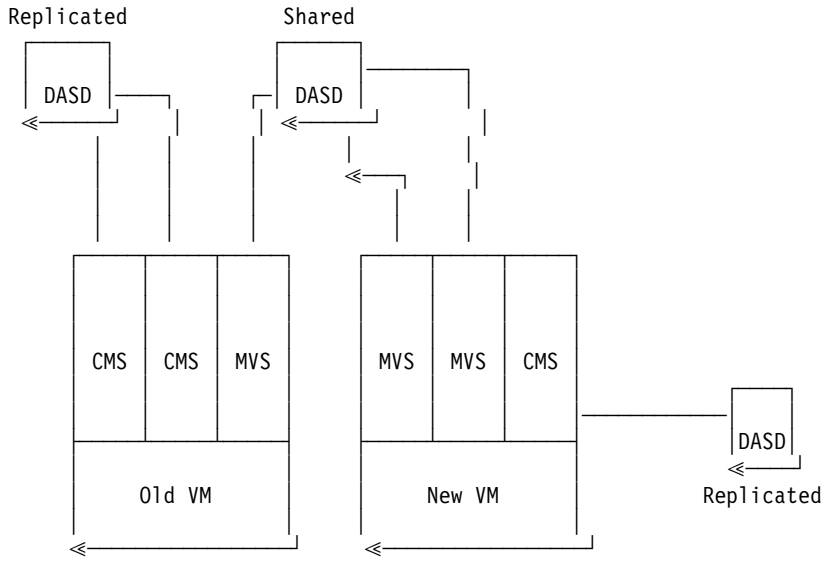

Figure 3. Sharing DASD between Virtual Machines on Multiple Systems

Concurrent virtual and real reserve/release support can be invoked either during system generation or at any time while the system is running.

Do the following to invoke concurrent virtual and real reserve/release while generating the system:

1. Ensure that the operating system running as a guest on the new system and the operating system with which you are sharing the DASD both support reserve/release CCWs.
2. Define the DASD as a shareable full-pack minidisk. To do this, use the MDISK statement in the user directory.

   To define the DASD as a full-pack minidisk, the starting cylinder or block number must be zero and the number of cylinders or blocks must equal or exceed the number of cylinders or blocks on the real device. To define the DASD as virtually shareable, include the V in the mode definition. For example:

   ```
   MDISK 327 3380 000 885 MVS003 MWV
   ```

   or

   ```
   MDISK 328 9336 0000 END MVS003 MWV
   ```
3. Define the DASD as being shareable between multiple real and virtual systems by specifying the SHARED YES option on an RDEVICE statement in SYSTEM CONFIG. For example:

   ```
   Rdevice 0327 Type Dasd Shared yes
   ```

   Note: Unless SHARED YES is specified, CP assumes that the device is not shared.

Migrating Your User Directory

This section helps you migrate your old user directory to the new system. Do not migrate the directory until you finish installing the new system.

Attention: If you are migrating from a pre-V6.2 system, see "[V6.2] SSI Cluster User Identity and Configuration" on page 201.

Object Directory Compatibility

Object directory compatibility is supported when the object directory created with the DIRECTXA MODULE on the new system is used by **other currently supported levels of CP**. This support lets you use a single source directory in a mixed Cross System Extensions (CSE) environment, or when migrating to the new release of CP.

A copy of the new DIRECTXA MODULE should exist on each system that will share a single source directory and depend on object directory compatibility. A new copy of the DIRECTXA MODULE should be distributed to each of the systems whenever service is applied to the HCPDIR ASSEMBLE file on the new system, which is used to create the DIRECTXA MODULE.

Getting Around Directory Differences

If you want to use the same source directory for your old and new systems:

- **Migrating from an Unsupported Release:** Avoid using any directory statements or options that were introduced after your current release. MIXED directory APARs do not exist for any subsequent releases. To include statements or options that were introduced after your release, you must maintain two source directories. The statements and options that are unsupported by your old system should be included only in the source directory for the new system.
- **Migrating from a Supported Release:** You can do either of the following:
  - Use the DIRECTXA MODULE on the new system to create the object directory for both systems, as described above.
  - Avoid using any directory statements or options that were introduced after your release. MIXED directory APARs do not exist for any subsequent releases.

User Directory Control Statement Changes

For a list of control statement changes, refer to the control statements compatibility table (if any) for each of the releases since your old release.

Default Volume Label Changes

The default volume labels have changed from your old release to the new release. See the installation worksheet in *z/VM: Installation Guide*.

Preparing to Migrate Your User Directory

If you have a problem IPLing the new system, it may be because of changes you have made to the directory. If you suspect this, you can use the NODIRECT option when you bring up the new system. This brings up the system without accessing the user directory.

Before you can use NODIRECT, you need to know information about the source directory's minidisk, the CMS system disk, and the system residence volume. Table 60 on page 432 shows the information you need. Record your information now, because it might be too late to get this information if you wait until you have directory problems. See *z/VM: CP Planning and Administration* for more information on loading z/VM V6.4. See *z/VM: Installation Guide* for information on the IBM default location and sizes of minidisks.

Table 60. Examples of Information Needed To Use the NODIRECT Option

	Minidisk or Volume Address	Volume Identifier	Starting Cylinder or Block	Number of Cylinders or Blocks Allocated
Minidisk containing source directory	2CC	VMCOM1	296	5
CMS system disk	190	M01RES	189	107
System residence volume	123	M01RES	0	END

Using Two Source Directories

During your migration, you may choose to maintain either two source directories or one. If you use two directories, one on the old system and one on the new, you need to be very careful that any changes are reflected in both directories, if needed. Using two directories allows you to exploit the new functions of z/VM V6.4 without worrying about backing out.

Considerations for Using a Directory Maintenance Program

If you have already installed a directory maintenance program on your z/VM V6.4 system, or if you are using only one directory, you may have to issue additional commands while migrating the source directory.

Example Using the IBM Directory Maintenance Facility (DirMaint)

For example, several additional steps are needed if DirMaint is used:

1. Before changing the source directory, enter DIRM DISABLE. This ensures that DirMaint does not try to update the directory while you are converting it.
2. Enter DIRM USER BACKUP. This creates a CMS flat file of the current source directory named USER BACKUP and found on DirMaint's 1DB minidisk (G-disk).
3. Update USER BACKUP with the new directory statements.
4. After updating, copy USER BACKUP to DirMaint's 1DF minidisk (E-disk) and rename it to USER INPUT.
5. Erase any USER DIRECT E file and start DirMaint by autologging it or by running DVHBEGIN or DIRM RLDDATA.
6. Enter DIRM ENABLE to allow DirMaint to resume making updates.
7. Issue DIRM DIRECT command to put the directory online.

For complete information on how to make manual changes to the source directory when DirMaint is installed and running, see:

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

Steps for Migrating Your Source User Directory

Whether you use two directories or one during migration, you can migrate your source directory by merging information from the old directory into the new directory. The old directory is left unchanged and can be either discarded or used on the old system. To do this:

1. Make a copy of the source user directory on the old system with a different file name. Move the copy to the new system.

2. Edit the z/VM V6.4 directory and merge pieces from the old directory into it as the old directory entries are migrated. To migrate the old directory entries:
   a. Migrate each system user ID.

      IBM supplies a sample directory with the z/VM product. In this directory, there are several system user IDs defined. For each system user ID that you used in your old system, verify that your directory entries are applicable in z/VM V6.4:
      - If applicable, make the needed changes in the new directory's entry.
      - If not applicable, leave the new directory's entry as is.

   b. For all other user IDs:

      Note: This step may be staged or done in groups as you move the users to the new system.

      1) Increase storage sizes as needed. The minimum storage size for a shared copy of CMS is 256 KB. At least 20 MB is needed to IPL a nonshared copy of CMS, for example to IPL the 190 minidisk.
      2) Verify that AUTOONLY, NOLOG, NOPASS, or LBYONLY is not specified in the password field of the USER statement unless the user is to take advantage of the function the operand provides.
      3) Specify the type of virtual machine you want your user IDs to have. You can specify an ESA (the default), XA (equivalent to ESA), or XC virtual machine.
      4) Look for usage of incompatible directory control statements. See the user directory control statement compatibility tables in this document for the list of incompatible statements. Note any defaults that changed since your old release.
      5) Resolve any incompatibilities by:
         - Removing directory control statements that are no longer supported.
         - Changing directory control statements with changed parameters.
         - Optionally, exploiting new directory control statements and parameters.
      6) Move all migrated user IDs to the new directory.
3. Check the syntax of your new directory using the DIRECTXA command with the EDIT option. For example:

   ```
   directxa user direct (edit
   ```
4. Fix any problems.
5. Put the new directory on-line using the DIRECTXA command, for example,

   ```
   directxa user direct
   ```

If You Cannot IPL Because of a Problem with the User Directory

If you have a problem IPLing z/VM V6.4, you can try using the NODIRECT option. Because z/VM V6.4 does not support a stand-alone directory function, you must use the following steps to recover when you cannot IPL because of a problem with the user directory. The information recorded in Table 60 on page 432 is used in this procedure.

1. IPL and bring up the system using the NODIRECT option. This logs on the primary system operator.

2. Use the DEFINE MDISK command to obtain access to the minidisks containing the directory source file, the CMS system disk, and the CP system residence volume. You recorded the minidisks addresses, starting values, and sizes in Table 60 on page 432.

   For example, using the sample information in Table 60 on page 432:

   ```
   define mdisk as 2cc 296 5 620RES
   define mdisk as 190 189 107 620RES
   define mdisk as 123 0 END 620RES
   ```

3. IPL 190
4. Enter the ACCESS command to access the minidisk that contains the source directory.
5. Fix the problem with the directory.
6. Put the corrected directory on-line using the DIRECTXA command. For example:

   ```
   directxa user direct
   ```

7. Shutdown and re-IPL the system without using the NODIRECT option.

See *z/VM: System Operation* for a complete step-by-step procedure for recovering the user directory after a problem during IPL.

Migrating Your SFS File Pool Servers

The data migration methods described in this section are applicable to installation-defined SFS file pool servers, as well as for data created by your user community in the IBM-supplied VMSYSU file pool.

You will probably want to have *z/VM: CMS File Pool Planning, Administration, and Operation* available as you do this task.

Using Two System Images

For this type of migration, you have both the old system and the new system running at the same time. The new system can be installed on a separate physical processor, in a logical partition, or as a second level system on your old system.

You can do the migration in two ways:

- Move the entire user population for your SFS file pool servers all at once.
- Stage the move by grouping your users, for example, by department, by project, or whatever is appropriate.

Staging the move may be time consuming. Also, if you are not careful, aliases and authorizations may be lost in the move.

You can automate the move somewhat by using the SFSTRANS EXEC, which is documented in *z/VM: CMS File Pool Planning, Administration, and Operation*, as a model. Note that SFSTRANS assumes that you have access to both the old and the new servers.

Procedure

1. Use FILEPOOL BACKUP for each storage group to back up all the data on the old release.

   For assistance, see the section on backing up the user data in *z/VM: CMS File Pool Planning, Administration, and Operation*. This backup can be used on your new server. It can also be used if you need to back out to your old release.

Note: [V2.1] Control data backup files created on your new system are not supported on the old system.

2. Shut down the SFS file pool servers on your old system.

   Enter the STOP operator command. Do not use STOP IMMEDIATE. For example, from the server machine console, enter:

   ```
   stop
   ```

   Or, from a secondary user console, such as MAINT, enter:

   ```
   #cp send vmserv3 stop
   #cp send vmserv4 stop
   ```

3. Install the new level of z/VM.
4. Set up a file pool server machine on the new system.

   If you are moving everyone at once:

   If you can, move the entire DASD pack that contains the SFS minidisks to the new system. Use the same minidisk locations and addresses defined in the directory entry for the server on the old system. If needed, update the directory entry as described in "User Directory Considerations and Changes" on page 437.

   If you cannot move the entire DASD pack to the new processor, configure your directory entry for the new server so that it has identical minidisk addresses and sizes as the server on the old system. If you are moving the minidisks to a DASD of a different type, it may be impossible for the sizes to be exactly the same. In that case, make them slightly larger. Then, use the DFSMS COPY command or the DDR command to move the contents of the old server's minidisks onto the new server's minidisks. See *z/VM: CMS File Pool Planning, Administration, and Operation* for more information about how to prepare for and use DFSMS COPY.

   If you are staging the move by groups:

   Make sure the new file pool server machine has enough physical DASD space to hold the group of users that you want to move. See the section on generating a file pool and server in *z/VM: CMS File Pool Planning, Administration, and Operation* for details on how to do this.

5. **If you are staging the move by groups:**

   Move the users in the group that you want to move to the new file pool. Consider modifying the SFSTRANS EXEC, which is shown in *z/VM: CMS File Pool Planning, Administration, and Operation*, to automate this procedure for you. To move users to the new file pool, do the following:

   a. Enroll the users in the new file pool. Make sure the users have enough file blocks to contain their data.
   b. Re-create each user's directory structure in the new file pool.
   c. Copy or move the users' files from the old file pool to the new file pool.
   d. Remove the users' space from the old file pool.

      If you can, use the DELETE USER command to delete a user from the old file pool. However, if a user needs to write to others' files in the old file pool, you can:

      - Delete the user from the old file pool, then immediately re-enroll the user, or
      - Erase the copied or moved files, and reduce the user's space to 0 using the MODIFY USER command.

      Notes:

1) If you specify the KEEPAUTH option on the DELETE USER command, authorizations that were granted to that user ID are not deleted.
2) Aliases that others have for the moved files are lost, and aliases that the moved user has in the old file pool are lost. If the users still need to share files, they have to access each others' directories. Aliases cannot refer to base files in other file pools.

6. Log on to the file pool server machine. Make sure it uses the new CMS code, which is typically:
   - CMS code: In the CMS segment or on the 190 minidisk, and
   - SFS code: On the 193 minidisk or in the CMSFILES segment.
7. IPL the new CMS.
8. Check the USERS startup parameter in the DMSPARMS file for your server and the MAXCONN value in the server's directory entry against the suggested values found in *z/VM: CMS File Pool Planning, Administration, and Operation*. This can help you avoid potential virtual storage problems.
9. If BACKUP is specified in the DMSPARMS file, issue FILESERV BACKUP to back up control data. You must do this before issuing FILESERV START or you will get message DMS3440E, and FILESERV START will not continue.
   ```
   fileserv backup
   ```
10. Start the SFS file pool servers.

    If FILESERV START is not included in the PROFILE EXEC of the file pool server virtual machine, enter it from the operator's console:
    ```
    fileserv start
    ```
11. Immediately back up your server data. Use FILEPOOL BACKUP for each storage group to back up all the data on the old release.

    Note: Control data created on your old release cannot be used on the new release Storage group (user data) backup files created on your old release can be used on the new release.
12. If you encounter problems and need to back out to your old system, see "Converting an SFS File Pool Server Back to Your Old Release" on page 438.

Using the "Cut and Go" Strategy

The "cut and go" strategy is to move immediately off the old system and right on to the new one. This approach is recommended only if your system is a guest system or is small or very straightforward.

Procedure

1. Use FILEPOOL BACKUP for each storage group to back up all the data on the old release.

   For assistance, see the section on backing up the user data in *z/VM: CMS File Pool Planning, Administration, and Operation*. This backup can be used on your new server. It can also be used if you need to back out to your old release.

   Note: Control data created on your old release cannot be used on the new release. Storage group (user data) backup files created on your old release can be used on the new release. Neither control data nor storage group backup files created on your new system are supported on the old system.
2. Review and make note of the start-up parameters, which are in the DMSPARMS file, for each server on your old system.
3. Shut down your old-system SFS file pool servers.

Enter the STOP operator command. Do not use STOP IMMEDIATE. For example, from the server machine console, enter:

```
stop
```

Or, from a secondary user console, such as MAINT, enter:

```
#cp send vmserv3 stop
#cp send vmserv4 stop
```

4. Install the new level of z/VM.
5. Make sure the server virtual machine has access to the disks where the new CMS code resides, which is typically:
   - CMS code: In the CMS segment or on the 190 minidisk, and
   - SFS code: On the 193 minidisk or the CMSFILES segment.
6. IPL the new CMS.
7. Check the USERS startup parameter in the DMSPARMS file for your server and the MAXCONN value in the server's directory entry against the suggested values found in *z/VM: CMS File Pool Planning, Administration, and Operation*. This can help you avoid potential virtual storage problems.
8. If BACKUP is specified in the DMSPARMS file, issue FILESERV BACKUP to back up control data. You must do this before issuing FILESERV START or you will get message DMS3440E, and FILESERV START will not continue.

   ```
   fileserv backup
   ```
9. Start the SFS file pool servers.

   If FILESERV START is not included in the PROFILE EXEC of the file pool server virtual machine, enter it from the operator's console:

   ```
   fileserv start
   ```
10. If you encounter problems and need to back out to your old system, see "Converting an SFS File Pool Server Back to Your Old Release" on page 438.

User Directory Considerations and Changes

You may have to modify some of the virtual machine definitions for your SFS file pool servers. If you have not done so already, make the following changes:

- **USER or IDENTITY directory control statement:**
  - Make sure the minimum and maximum virtual storage are at least 32 MB.
- **OPTION directory control statement:**
  - Add the NOMDCFS operand to allow the server to use minidisk caching at a rate that is not limited by the Fair Share Limit.

    Note: This is applicable only to SFS file pool servers, not CRR recovery servers.
  - Add the QUICKDSP operand to allow the server to be added to the dispatch list immediately when it has work to do, without waiting in the eligible list.
  - Add the SVMSTAT operand to specify that the virtual machine is a service virtual machine. This causes the server's monitor statistics to be reported separately from end-user virtual machines.
  - Check your MAXCONN value in the server's directory entry against the suggested values found in *z/VM: CMS File Pool Planning, Administration, and Operation*. This can help you avoid potential virtual storage problems.
- **SHARE directory control statement:**
  - Add SHARE REL 1500 to place the server in a more favorable position in the dispatch queue.

- **MACHINE directory control statement:**
  - If not already done, replace 370 with ESA or XC. (XA is accepted for compatibility and is functionally equivalent to ESA.)

    SFS file pool servers should use XC to exploit data spaces. CRR recovery servers do not exploit data spaces and therefore should be set to ESA.
- **XCONFIG directory control statement:**
  - If you designated XC in the MACHINE directory control statement, then add the following XCONFIG directory control statements:

    ```
    XCONFIG ADDRSPACE MAXNUMBER 100 TOTSIZE 8192G SHARE
    XCONFIG ACCESSLIST ALSIZE 1022
    ```

    These statements have appropriate data space values for an SFS file pool server.
- **MINIOPT directory control statement:**
  - Add MINIOPT NOMDC to inhibit expanded storage caching for the following file pool minidisks:
    - For SFS file pool servers:
      - Control minidisk
      - SFS log minidisks
    - For CRR recovery servers:
      - All file pool minidisks

Converting an SFS File Pool Server Back to Your Old Release

z/VM does not support the conversion of SFS file pools and servers back to a previous release. However, IBM does recognize that special situations exist where such a conversion may be desired. The following procedure is intended to provide some guidance. Please read this entire section before beginning the conversion.

If you do the conversion incorrectly, some of the problems you may encounter include:

- Not being able to start the file pool server in your old system
- Loss of data

Procedure

To avoid file pool server problems, you must perform the following steps, where required, in the order indicated:

1. Shut down the new SFS file pool server.

   Stop file pool server multiple user mode processing by entering the STOP operator command. **Do not use STOP IMMEDIATE.** For example, from the server machine console enter:

   ```
   stop
   ```

   Or, from a secondary user console, such as MAINT, enter:

   ```
   #cp send vmservu stop
   ```

   If your SFS file pool server is also a CRR recovery server, this step ensures that all CRR logging activity completes normally. There are no additional migration considerations if the SFS file pool server is also a CRR recovery server. (Your SFS file pool server is a CRR recovery server if the CRR start-up parameter exists in its *serverid* DMSPARMS file.)

Also, if possible, do not change the LU name value on the LUNAME start-up parameter. See *z/VM: CMS File Pool Planning, Administration, and Operation* for LUNAME start-up parameter considerations.

2. Convert your system back to the old release.
3. On the old system:
   a. If you are using SFS control backup, which means you used the BACKUP start-up parameter, execute the FILESERV BACKUP command to create a new control backup file. For example:

      ```
      fileserv backup
      ```

      Note: Control backup files created by one VM release cannot be processed by a different (earlier or later) release.
   b. Start the old-system SFS file pool server for multiple user mode processing.

      If FILESERV START is not included in the PROFILE EXEC of the file pool server virtual machine, enter it:

      ```
      fileserv start
      ```

Your SFS file pool conversion back to the old release is now complete.

Migrating Your BFS Data

Depending on the services and applications used by your installation, you also might need to take specific steps to migrate select byte file system (BFS) file spaces that contain data that has been created or customized for your installation. Such data can and will reside only in the IBM-supplied and defined VMSYS file pool, which by default is intended to maintain BFS files and data. Examples of such data include:

- SSL server data, such as key databases and certificate files
- LDAP server data, such as schema databases, and the change log or checkpoint file

As mentioned in "Migrating Your SFS File Pool Servers" on page 434, you will probably want to have *z/VM: CMS File Pool Planning, Administration, and Operation* available as you do this task.

The suggested method for moving such data is to use the FILEPOOL UNLOAD and FILEPOOL RELOAD commands. The steps below summarize the necessary commands. For complete details and additional command considerations, see the topic on managing users and file spaces in *z/VM: CMS File Pool Planning, Administration, and Operation*.

Follow these steps to move a BFS file space from one system to another, maintaining use of the same storage group:

1. Log on a user ID that is an administrator for the subject file pool (VMSYS, by default) on the system where the BFS data of interest currently exists.
2. Issue a FILEDEF command to identify the file that is to contain the unload data. For ease in identifying the data, in this example the name of the file space (*file_space*) is used as part of the file identifier for a disk-resident file.

   ```
   filedef unload disk file_space unload a
   ```

   By default, SSL key database data is maintained in the GSKSSLDB file space, while LDAP server data is maintained in the LDAPSRV file space.
3. Issue a FILEPOOL UNLOAD command for the file space that is being moved:

```
filepool unload filespace file_space vmsys: ( all
```

4. Transfer the *file_space* UNLOAD file to the new system, using appropriate means (for a second level system, perhaps by copying the file to a system-acccesible minidisk, or by using FTP).
5. Issue a FILEPOOL RELOAD command for the file space that is being moved:

```
filepool reload filespace file_space vmsys:
```

Establishing Connectivity between Your New and Old Systems

Establishing connectivity between your new system and your old system allows you to move data files from the old system to the new system. If you have installed the new system on a separate processor or in another logical partition on the same processor as the old system, you can use real channels to establish connectivity between the two systems. If you have installed the new system as a guest of the old system, you can use virtual channels to establish connectivity between the two systems.

For example, you can set up a virtual channel between the RSCS licensed program installed on the old system and RSCS installed on the new system. A virtual channel for using PVM, VTAM®, TSAF, or AVS could be set up in a similar manner.

To set up RSCS, you make a virtual channel and add the proper definitions in the configuration and control files. RSCS relies on GCS for supervising services. The RSCS virtual machine must be defined as part of the GCS group so that it can share a common storage area for information exchange, multitasking services, and general I/O services.

Figure 4 shows the relationship between the old (first-level) and new (second-level) RSCS machines.

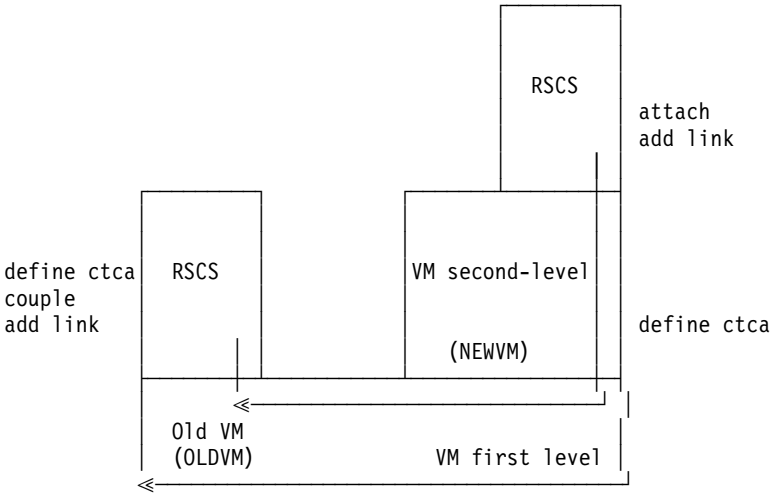

Figure 4. Connectivity between First-Level and Second-Level RSCS Machines

Procedures

On the First-Level System

1. Define a link between the first-level RSCS virtual machine and the user ID that will be running the second-level system.

In this example, the user ID that will be running the second-level system is NEWVM.

The virtual address of the channel that is used for the link is 312.

a. From **both** the first-level RSCS virtual machine and the NEWVM user ID, define the channel using the CP DEFINE command or the SPECIAL user directory control statement.

   For example, using the CP DEFINE command, you can enter:

   ```
   define ctca 312
   ```

   Or, in the user directory entries for NEWVM and the RSCS virtual machine, add the following statement:

   ```
   SPECIAL CTCA 312
   ```

b. Ensure that the channel you specified using DEFINE CTCA is defined in the second-level system's SYSTEM CONFIG file.

   In our example, we are using 312 as the address of the channel that is used for the link.

2. From either the RSCS machine or the NEWVM user ID, couple the addressed links to enable communications over the virtual channel. For example, from the first-level RSCS virtual machine, you can enter:

   ```
   couple 312 newvmesa 312
   ```

3. Bring up the second-level system, including the second-level RSCS machine.

On the Second-Level System

1. From the second-level RSCS virtual machine, attach the channel address defined by NEWVM. For example, enter:

   ```
   attach 312 rscs 312
   ```

How to Back Out of the Migration

You may have to back out of the migration because of either system problems or problems with application programs.

Note: To back out spool files, use the SPXTAPE command.

Backing Out Because of System Problems

How you back out depends on whether you are using separate residence volumes for the old system and the new system or the same residence volume for both.

Using Separate Residence Volumes

If you run into a problem with z/VM V6.4 that forces you to back out, and you have not yet exploited functions specific to z/VM V6.4 or migrated any CMS users, the backout will not affect your end users. However, if you have to back out after migrating some CMS users, you will greatly affect those users. Once users convert their applications to run on the z/VM V6.4 CMS, they may not be able to run these applications on the old CMS. Therefore, you should not migrate users to z/VM V6.4 until you are satisfied with its stability.

Throughout the migration, you should maintain your old system residence volume and CP-owned packs. Then, if you need to back out, you can easily do so.

If you must back out from z/VM V6.4 and IPL an old system, do the following:

1. Dump the spool files with the SPXTAPE DUMP command.
2. Shut down the z/VM V6.4 system.

3. Re-IML the processor to the appropriate mode, if needed.
4. IPL the old system.
5. Load the spool files with the SPXTAPE LOAD command.
6. Enable the terminals.

Using the Same Residence Volume

In preparation for a possible backout:

1. Before going to the new system, use stand-alone DDR to dump the nucleus:
   ```
   ddr dump nuc
   ```
2. When you go to the new system, use the same CP-owned volumes, warmstart area, and checkpoint area. Also, do not overwrite where the old system's nucleus was; keep that space available.

Then, if you have to back out, do the following:

1. Use the SPXTAPE DUMP command to save the spool files, if you want to.
2. IPL the stand-alone DDR program.
3. Use DDR to restore the nucleus to the system residence pack.
4. IPL the system residence pack.
5. Use the SPXTAPE LOAD command to restore any saved spool files.

Backing Out Individual Users Because of Problems with Application Programs

When the system is stable, you can begin to migrate CMS users. To prepare for a possible backout, either dual-path the code or keep copies of the old (pre-migration) versions of your applications.

User Directory Concerns

If a user or group of users has problems running applications on z/VM V6.4 CMS, you will need to backout these users to the old release. To plan for this, keep a user directory entry in an old-level system for each z/VM V6.4 user until you are satisfied that the users are running smoothly. You can use either an old-level guest or a separate old-level system as the backout system. This will be extra work for you but will have the least impact on the end users if they run into problems.

If you are using two source user directories, you must remember to reflect changes made in one user directory to the other user directory.

Installing a Backlevel CMS

Note: If z/VM is installed on an IFL engine, you cannot run CMS Level 16 or earlier.

IBM provides limited support for multiple levels of CMS on z/VM V6.4. You can continue to use your production CMS with the CP component of z/VM V6.4, and then gradually migrate users and applications to the new CMS. Previous releases of the CMS component will be supported for a limited time following the general availability of z/VM V6.4.

There is no intent to retrofit new function onto old releases of CMS. Attempts to use new function on backlevel releases of CMS are unsupported and the results are undefined. New function is defined as any device support, new CMS, or new CP

functions introduced in later releases of the operating system and not retrofitted to the old CMS through the APAR service stream.

IBM will help with problem determination in these mixed environments and will take APARs for problems in older releases of CMS. However, when IBM discontinues service of an old VM release, support for the corresponding CMS under this offering also ends.

Notes:

1. Throughout this discussion, backlevel CMS refers to an older level of CMS, still in service, running on the z/VM V6.4 CP.
2. The z/VM V6.4 CMS does not support back levels of CP.

Before You Install a Backlevel CMS

While you are still running on your old system, there is certain information you should record that will help when you install your backlevel CMS. Record this information in Table 61 on page 444.

- Record the MDISK statement that defines the 190 minidisk on your old system. This can be used as a reference when defining the minidisk to hold the backlevel CMS on z/VM V6.4.
- Record the number of cylinders needed for your old 190 disk. You can get this information from the MDISK statement that defines your old 190 disk, or you can enter:

  ```
  q virtual 190
  ```

- Determine the type of DASD your old CMS will reside on when you install it on your z/VM V6.4 system. While installing a backlevel CMS, you will have to copy the contents of your 190 minidisk from the old system to your z/VM V6.4 system.
  - If you use the same type of DASD to hold the backlevel CMS on your z/VM V6.4 system as you used on your old system, you can use DDR to copy the files over. This is the preferred method.
  - If the DASD type is not the same, you can use the VMFPLC2 command. This method is not as easy as using DDR because you have to rebuild the CMS nucleus.

  Record the type of DASD you will use and how you will copy files over to z/VM V6.4.
- If you are using the VMFPLC2 command to copy files, you have to format the z/VM V6.4 minidisk that holds the backlevel CMS using the FORMAT command with the RECOMP option. The RECOMP option changes the number of cylinders on a disk that are available to you. By using the RECOMP option on the FORMAT command you can leave a number of cylinders on the minidisk available for the CMS nucleus. If you do not FORMAT the minidisk with RECOMP, you cannot build the CMS nucleus.

  To figure out how many cylinders to recomp, from your old system enter:

  ```
  query virtual 190
  ```

  An example of what this command returns is:

  ```
  DASD 0190 3380 SYGEMC R/O     135 CYL  ON DASD
  ```

  Now enter:

  ```
  query disk s
  ```

  An example of what this command returns is:

```
LABEL  VDEV M  STAT  CYL TYPE BLKSIZE  FILES  BLKS USED-(%) BLKS LEFT  BLK TOTAL
CMSOLD 190  S   R/O  120 3390 4096       843      17905-83       3695      21600
Ready;
```

The difference in the number of cylinders displayed by these two commands is the amount of space needed to hold the CMS nucleus. The number of cylinders shown in the QUERY DISK response is the number of cylinders needed to hold the contents for the 190 disk and is also the number of cylinders to recomp. In this example, 120 cylinders are required on a 3390 DASD. By recomping 120 cylinders, fifteen cylinders are available for the CMS nucleus.

If you are using the VMFPLC2 method because your DASD types are not the same, the number of cylinders you have to recomp may not be exactly the number calculated above. This is because the number of pages per cylinder varies from DASD type to DASD type. To make sure you are recomping a sufficient number of cylinders:

1. Figure out how many pages per cylinder there are for each DASD type you are using. You can find this information in the DASD storage capacity tables in *z/VM: CP Planning and Administration*. For example, a 3380-E has 150 pages per cylinder and a 3390-1 has 180 pages per cylinder.
2. Multiply the number of cylinders needed to hold the contents of your old 190 disk (120 in this example) by the pages per cylinder for each DASD.

   For 3380 model E:
   120 * 150 = 18000 pages

   For 3390-1:
   120 * 180 = 21600 pages

   This shows you that you need 18000 pages of storage for the contents of the old 190 disk. If you recomp 120 cylinders of 3390-1 DASD, this gives you 21600 pages of storage. Therefore, 120 cylinders is a sufficient number of cylinders to recomp.
3. Multiply the number of cylinders needed to hold the CMS nucleus (fifteen in the previous example) by the pages per cylinder for each DASD. For example,

   For 3380 model E:
   15 * 150 = 2250 pages

   For 3390-1:
   15 * 180 = 2700 pages

   If you reserve fifteen cylinders on a 3390-1 DASD, it is plenty of space to hold the CMS nucleus that you stored in fifteen cylinders on a 3380 model E.

If you calculated the number of cylinders to recomp on a 3390-1 DASD to be 120 cylinders, you would have to recomp 144 cylinders on a 3380-E DASD to store the same amount of data. This is because there are fewer pages per cylinders on a 3380-E.

Table 61 shows examples of the information you need.

Table 61. Information Needed Before Installing a Backlevel CMS

Type of Information	Example
MDISK statement for old VM system 190 disk	MDISK 190 3380 125 135 CMSRES RR ALL
Number of cylinders for old VM system 190 disk	135
DASD type for old VM system CMS	3380

Table 61. Information Needed Before Installing a Backlevel CMS (continued)

Type of Information	Example
DASD type for backlevel CMS on new system	3390
Method you will use to copy files (DDR or VMFPLC2)	VMFPLC2
If using VMFPLC2, number of cylinders to recomp	120

Steps for Installing a Backlevel CMS

When you are done installing the old CMS, the environment will look like this:

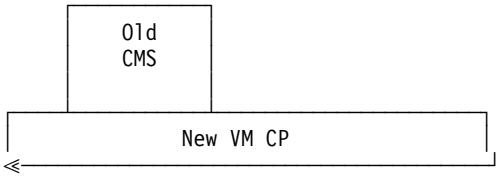

Figure 5. Backlevel CMS Running on a z/VM V6.4 CP

The steps for installing the old CMS on the z/VM V6.4 CP are as follows:

1. Log on to your z/VM V6.4 MAINT user ID.
2. Create a minidisk where the backlevel CMS will reside. You can use the MDISK directory control statement from your old system as a reference. You recorded this in Table 61 on page 444.

   The MDISK directory control statement in your z/VM V6.4 user directory for the backlevel CMS may look like this:

   ```
   MDISK 590 3380 125 135 620RES MR ALL WRITE MULTIPLE
   ```

   - The 590 minidisk eventually becomes the 190 disk for the backlevel CMS. The 590 minidisk can reside on any DASD volume.
   - In this example, 135 cylinders are reserved for the 590 disk starting at cylinder 125. You recorded the number of cylinders needed by your old 190 disk in Table 61 on page 444.
     - If you are using the DDR method later in the procedure, you must make your backlevel CMS minidisk on z/VM V6.4 exactly the same size as your old CMS 190 minidisk.
     - If you are using the VMFPLC2 method, the backlevel CMS minidisk can be the same size or bigger than your old CMS 190 minidisk.
   - Once you know how many cylinders you need for the 590 disk, you have to figure out where to put the minidisk in your z/VM V6.4 configuration. Use the DISKMAP command to find out where to put the backlevel CMS minidisk. In the previous MDISK statement, the starting cylinder is 125. This was obtained by entering:

     ```
     diskmap user direct
     ```

     For more information on DISKMAP, see *z/VM: CP Commands and Utilities Reference*.

     The output file from this command is USER DISKMAP. This file shows you the space you have available. If you cannot find enough contiguous space, you may have to shift other minidisk locations and data around to make room for the backlevel CMS.

Note: If you move a minidisk location, make sure you also move the data on that minidisk. If you decide to shift locations of minidisks used for the Shared File System, there may be additional considerations. See *z/VM: CMS File Pool Planning, Administration, and Operation* for details on replacing the SFS and CRR file pool minidisks.

3. Put the directory with the new statement online by issuing:

   ```
   directxa user direct
   ```

4. Make sure your MAINT user ID is running an XA virtual machine. To check the virtual machine mode, you can use the QUERY SET command. You can use the SET MACHINE XA command or the MACHINE directory control statement to set MAINT's virtual machine to XA.
5. Link the minidisk where the backlevel CMS will reside by entering:

   ```
   link maint 590 590 MR
   ```

6. Format the 590 disk and access it using an unused file mode by entering:

   ```
   format 590 k
   DMSFOR603R Format will erase all files on disk k(590).
   Do you wish to continue? Enter 1 (YES) or 0 (NO).
   1
   DMSFOR605R Enter disk label:
   CMSOLD
   ```

Note to VMFPLC2 Users

If you plan to use the VMFPLC2 method to copy files to z/VM V6.4, skip to step 10 on page 447.

If you are using the DDR method, continue with step 7.

7. Copy the contents of your old 190 minidisk to the backlevel CMS minidisk on z/VM V6.4 590. The following steps show how to DDR the old 190 contents to the z/VM V6.4 590 minidisk.
   a. Log on to your old-system MAINT user ID.
   b. Determine the location and size of the old 190 minidisk:

      ```
      query mdisk 190 location
      ```

      The response from this command would be something like this:

      ```
      TargetID  Tdev  OwnerID   Odev  Dtype Vol-ID  Rdev     StartLoc      Size
      MAINT     0190  MAINT     0190  3390  410SYS  E202         1937       130
      ```

   c. Log on to your z/VM V6.4 MAINT user ID.
   d. Define the 190 location as a minidisk on the new system. For example:

      ```
      define mdisk f190 1937 130 410sys
      ```

   e. Dump the old-system files:

      ```
      ddr
      sysprint cons
      input f190 3390
      output 590 3390
      copy all
      ```

   Now you have the old VM system 190 CMS files on your z/VM V6.4 590 minidisk.
8. Define the 190 minidisk that contains z/VM V6.4 CMS with another address; for example:

   ```
   define 190 as 90
   ```

9. Define the 590 minidisk with the old-system CMS on it as your 190; for example:

   ```
   def 590 190
   ```

Note to DDR Users

DDR users can now skip to step 17 on page 448.

If you are using the VMFPLC2 method, continue with step 10.

10. While you are still on your z/VM V6.4 MAINT user ID, use the FORMAT command with the RECOMP option to change the available cylinders on 590 for disk-resident files. In the previous example, 120 cylinders needed to be recomped. (See Table 61 on page 444.) Remember, this leaves fifteen cylinders for the backlevel CMS nucleus. Use the following FORMAT command:

    ```
    format 590 k 120 (recomp
    ```

    The response from this command would be:

    ```
    LABEL  VDEV M  STAT  CYL TYPE BLKSIZE  FILES  BLKS USED-(%) BLKS LEFT  BLK TOTAL
    CMSOLD 590  K   R/W  120 3390 4096       843      17905-83       3695      21600
    Ready;
    ```

11. Copy the contents of the old-system CMS on your 190 minidisk to the backlevel CMS minidisk on your z/VM V6.4 590. The following steps show how to move the old-system 190 contents to the z/VM V6.4 590 minidisk using VMFPLC2.

    a. Log on to your old-system MAINT user ID.

    b. Create a copy of the CMS nucleus and put it in MAINT's reader by entering:

    ```
    spool punch *
    spool printer *
    vmfload cmsload dmsvm
    ```

    These commands put the load deck in MAINT's reader.

    c. Receive the load deck as a CMS file on your 191 disk. Make sure you record the file name and file type. The example described here uses the file name CMSNUC FILE A.

    d. Attach a tape to the MAINT user ID on your old system:

    ```
    attach rdev to maint as 181
    ```

    e. Access the old-system 190 disk as something other than S. If you access it as S, only files with a file mode of S2 are accessible.

    ```
    access 190 v
    ```

    f. Dump the old-system 190 files and the new copy of the CMS nucleus on the A-disk to tape and write two tape marks:

    ```
    vmfplc2 dump * * v
    vmfplc2 wtm
    vmfplc2 dump cmsnuc file a
    vmfplc2 wtm 2
    detach 181
    ```

    g. Log on to your z/VM V6.4 MAINT user ID.

    h. Attach the tape to the MAINT user ID on your z/VM V6.4 system:

    ```
    attach rdev to maint as 181
    ```

    i. Access the z/VM V6.4 590 disk and load the old-system 190 files off the tape and onto the 590 minidisk:

    ```
    access 590 v
    vmfplc2 load * * v
    ```

    Now you have the old-system CMS files on your 590 minidisk.

    j. Load the CMSNUC FILE on to your 191 A disk:

    ```
    vmfplc2 load cmsnuc file a
    ```

k. Detach the tape:

```
detach 181
```

12. Build your backlevel CMS nucleus on your z/VM V6.4 system.
    a. Enter the following commands:

```
spool punch *
spool printer *
spool reader hold
```

    b. Punch the file containing the CMS load deck, CMSNUC FILE, to the MAINT user ID. Make sure you use the NOHEADER option so a header is not included in the file.

```
punch cmsnuc file a (noh
```

    An example of a response from the previous PUNCH command is:

```
RDR FILE 0002 SENT FROM MAINT    PUN WAS 0002 RECS 025K CPY  001 A NOHOLD NOKEEP
```

    You need the file number of the CMS nucleus for the next step. The file number in this example is 0002.

    c. If the CMS nucleus is not the first file in your reader, order your reader so that the CMS nucleus is the first file processed; for example:

```
order rdr 0002
```

13. Define the 190 minidisk that contains z/VM V6.4 CMS with another address; for example:

```
define 190 as 90
```

14. Define the 590 minidisk with the old-system CMS on it as your 190. For example:

```
def 590 190
```

15. IPL the load deck from MAINT's reader:

```
ipl 00c cl
```

16. When you are sure you are done with the CMSNUC FILE, you may want to erase it; it can be rather big. Keep the tape with the old-system 190 files and the CMSNUC FILE for backup.

Note to Both DDR Users and VMFPLC2 Users
Use the following steps to put CMS into a named saved system.

17. Determine the DEFSYS command to use to define CMSOLD as a named saved system. You can probably just use your old system's DEFSYS command. You may want to add this DEFSYS command to an exec where you define all the other saved segments and saved systems. The sample DEFNSS EXEC shown in the following section defines various CMS named saved systems.
18. Issue the DEFSYS command or use the sample DEFNSS EXEC to define the CMSOLD named saved system; for example:

```
defnss cmsold
```

19. Use the QUERY NSS MAP command to check that CMSOLD exists.
20. IPL the 190 minidisk with CMSOLD on it:

```
ipl 190 parm savesys cmsold
```

Defining Various Levels of CMS

Use the following sample DEFNSS EXEC to define different levels of CMS in a named saved system.

```
/*****************************************************************/
/*                                                               */
/* An EXEC that will define the various levels of CMS            */
```

```
/* NSSs for our new system.                                          */
/*                                                                   */
/* You can pass in the following arguments:                          */
/* o CMSNEW - for the new CMS                                        */
/* o CMSOLD - for the old CMS                                        */
/* o ALL - for both the old CMS and the new CMS                      */
/*                                                                   */
/*******************************************************************/
  Arg parms .
/* Create a list of all NSSs to be defined */
  nssname.1 = 'CMSNEW'
  nssname.2 = 'CMSOLD'
  nssname.0 = 2
  allnss = ''
  Do i = 1 To nssname.0
     allnss = allnss nssname.i
     End

  skip_processing = 'NO'
  Select
/* Define only the new CMS */
     When parms = nssname.1 Then Do
        start = 1
        stop = 1
        End
/* Define only the backlevel CMS */
     When parms = nssname.2 Then Do
        start = 2
        stop = 2
        End
/* If DEFNSS is called with 'ALL' or blank then */
/* define both the new CMS and the backlevel CMS */
     When parms = 'ALL' | parms = '' Then Do
        start = 1
        stop = 2
        End
/* Set up error message information */
     Otherwise Do
        skip_processing = 'YES'
        badnss = allnss
        errors = 'YES'
        End
     End

  If skip_processing = 'NO' Then Do
/* DEFSYS command to define the new CMS */
     nss.1 = 'CP DEFSYS CMSNEW 0-D EW 20-23 EW F00-13FF SR MINSIZE=256K',
     'MACHMODE XA,XC,ESA PARMREGS=0-15'
/* DEFSYS command to define the backlevel CMS */
     nss.2 = 'CP DEFSYS CMSOLD 0-D EW 20-23 EW F00-12FF SR MINSIZE=256K',
     'MACHMODE XA,XC'

/* Default for no errors detected */
     errors = 'NO'
     badnss = ''

/* Define the NSSs */
     Do i = start To stop
        Address COMMAND nss.i
        If rc <> 0 Then Do
           errors = 'YES'
           badnss = badnss nssname.i
           End
        End
     End

/* If any errors were found, show error messages */
```

```
If errors = 'YES' Then Do
   Say 'The following NSSs were not defined because a definition has'
   Say 'not been set up for the NSS:' badnss
   Say
   Say 'A list of all valid NSS follows:'
   Say '     ' allnss
   End
Exit
```

Swapping Between the Backlevel CMS and z/VM V6.4 CMS

If you want to be able to switch between the z/VM V6.4 CMS and the backlevel CMS, consider using the following execs that link to the different system disks. You can make these execs available to your users. The users then need to do a SET MACH XA and IPL the corresponding CMS segment (IPL CMSNEW or IPL CMSOLD).

```
/* This exec swaps a backlevel CMS with a */
/* new CMS.                               */

    Address Command
    'EXECIO 0 CP (STRING DETACH 190'
    'CP LINK MAINT  190 190 RR'
```

Figure 6. SWAP2NEW EXEC

The user should then do a SET MACH XA and IPL CMSNEW.

```
/* This exec swaps a new CMS with a        */
/* backlevel CMS.                          */

    Address Command
    'EXECIO 0 CP (STRING DETACH 190'
    'CP LINK MAINT  590 190 RR'
```

Figure 7. SWAP2OLD EXEC

The user should then issue:

```
set mach xa
ipl cmsold
```

Appendix A. Sample Utilities for Converting to Configuration Files

This section provides reference information for the following sample utility programs:

- HCPDCON
- HCPRDEV
- HCPTRIO
- HCPTSYS

Note: These sample utility programs were provided with z/VM V4.4 and earlier releases. **They are not included with z/VM V5.1 or later.** You must locate and use them on your current VM system. The sample utilities were shipped with a file type of SAMPEXEC; to use them, you must change the file type to EXEC.

Another utility that you might find useful when converting to configuration files is CPSYNTAX, which verifies the syntax of a SYSTEM CONFIG file. CPSYNTAX is described in *z/VM: CP Commands and Utilities Reference*.

HCPDCON

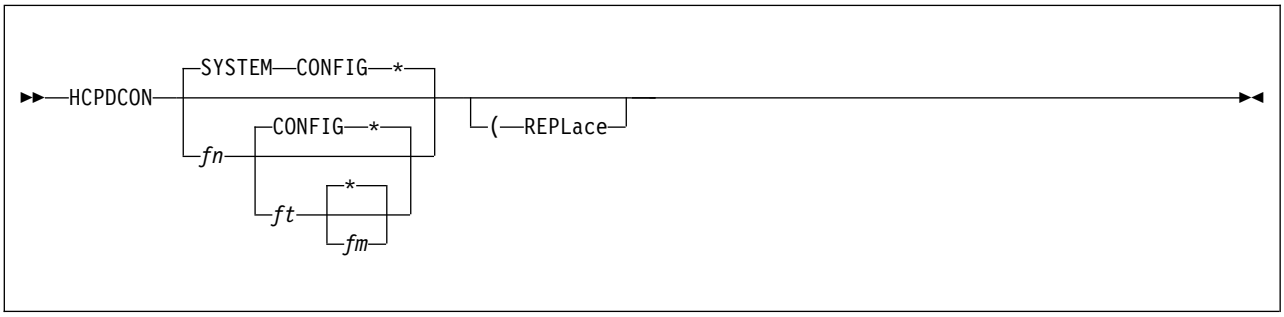

Authorization

Privilege Class: G

Purpose

Use HCPDCON to examine a running system and generate a file of configuration statements.

Operands

fn is the name of the CMS minidisk file that CP should generate. The default is SYSTEM.

ft is the type of file to be generated. The default is CONFIG.

fm is the CMS minidisk on which the file is to reside. This disk must be accessed in write mode. An asterisk (*), the default, causes HCPDCON to use the first CMS minidisk accessed in write mode.

REPLace
replaces the existing file with the new one just generated, if a file with the same name already exists.

HCPDCON sets up a work file with a file type of '$' plus the first seven characters of the file type of the input file. If such a file already exists, HCPDCON erases the file. After the work file is written to the CMS minidisk, HCPDCON erases the original configuration file and renames the work file.

For example, if you have a file called SYSTEM CONFIG A and you enter `HCPDCON`, HCPDCON generates a work file named SYSTEM $CONFIG A. HCPDCON then erases SYSTEM CONFIG A and renames SYSTEM $CONFIG A to SYSTEM CONFIG A.

Examples

1. To examine the running system and generate a configuration file called SYSTEM3 CONFIG A, enter the following:

   ```
   HCPDCON SYSTEM3 CONFIG A
   ```

2. To examine the running system, generate a configuration file called SYSTEM CONFIG, and put this file on the first CMS minidisk accessed in write mode, enter the following:

   ```
   HCPDCON
   ```

3. To examine the running system, generate a configuration file called SYSTEM4 NEWSTMTS, and put this file on the first CMS minidisk accessed in write mode, enter the following:

```
HCPDCON SYSTEM4 NEWSTMTS *
```

Return Codes

Code	Meaning
4	No minidisk accessed in R/W mode
8	User does not have necessary authority
12	Output file already exists and the REPLACE option was not specified
16	PIPE CMS RENAME error
20	I/O error writing to output file
24	Invalid option or too many options
9998	PIPE command not found
9999	Signal on NOVALUE routine exit

HCPRDEVS

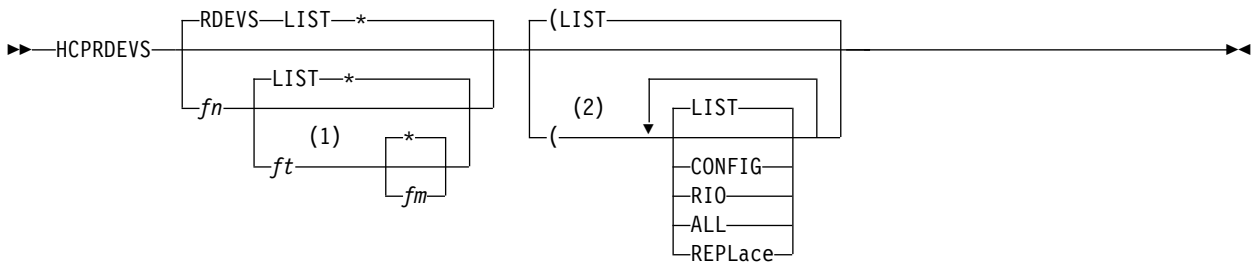

Notes:

1 Default file type is based on the option chosen.

2 You can use the options in any combination, but each option can be used only once.

Authorization

Privilege Class: G

Purpose

Use HCPRDEVS to do any or all of the following:

- Generate a list of all the real devices known to a system.
- Generate a system configuration-like file of those real devices that do not answer a sense ID request or do not return enough information. These devices must be defined to the system.
- Generate a file in an HCPRIO format.

Operands

fn is the name of the CMS minidisk file to be generated. The default is RDEVS.

ft is the type of file to be generated. The default depends upon the type of file requested. The option you choose (with the exception of ALL) denotes the default file type. For example, if the option is CONFIG, the default file type is CONFIG. The use of the option ALL results in the use of each default file type: LIST, RIO, and CONFIG.

fm is the CMS minidisk on which the file is to reside. This disk must be accessed in write mode. An asterisk (*), the default, causes the exec to use the first CMS minidisk accessed in write mode.

LIST
generates a list of real devices.

CONFIG
generates a system configuration-like file that contains statements only for those real devices that cannot be sensed completely.

RIO
generates a list of HCPRIO-like macroinstructions for all real devices on your system.

ALL
generates the information from the LIST, CONFIG, and RIO options. If you use ALL and a specific file type, the following files are generated:

- List of real devices with file type that you specified
- Information generated by the CONFIG option with file type of CONFIG
- Information generated by the RIO option with file type of RIO.

REPLace
replaces the existing file with the new one just generated, if a file with the same name already exists.

HCPRDEVS sets up a work file with a file type of '$' plus the first seven characters of the file type of the input file. If such a file already exists, HCPRDEVS erases the file. After the work file is written to the CMS minidisk, HCPRDEVS erases the original configuration file and renames the work file.

For example, `HCPRDEVS SYSTEM3 RIOLIST A (LIST REPL` generates a work file named SYSTEM3 $RIOLIST A. HCPRDEVS then erases SYSTEM3 RIOLIST A and renames SYSTEM3 $RIOLIST A to SYSTEM3 RIOLIST A.

Examples

1. To generate a file called MYRDEVS CONFIG, composed of configuration file statements for the real devices that do not return enough information to a sense ID request, and to place this file on your A-disk, enter the following:
   ```
   HCPRDEVS MYRDEVS CONFIG A (CONFIG
   ```
2. To create files called RDEVS LIST, RDEVS CONFIG, and RDEVS RIO, and to put them on the first CMS minidisk accessed in write mode, enter the following:
   ```
   HCPRDEVS (ALL
   ```
   RDEVS LIST contains a list of all the real devices known to the running system. RDEVS CONFIG contains a list of configuration file statements only for those real devices that cannot be sensed completely. RDEVS RIO contains a list of HCPRIO-like macroinstructions for real devices.
3. To generate two files, RDEVS LIST and RDEVS RIO, and to put them on the first CMS minidisk that CP has accessed in write mode, enter the following:
   ```
   HCPRDEVS RDEVS (LIST RIO
   ```
   RDEVS LIST contains a list of all the real devices known to the running system. RDEVS RIO contains a list of HCPRIO-like macroinstructions for real devices.

Return Codes

Code	Meaning
4	No minidisk accessed in R/W mode
8	User does not have necessary authority
12	Output file already exists and the REPLACE option was not specified
16	I/O error in creating output file
9998	PIPE command not found
9999	Signal on NOVALUE routine exit

HCPTRIO

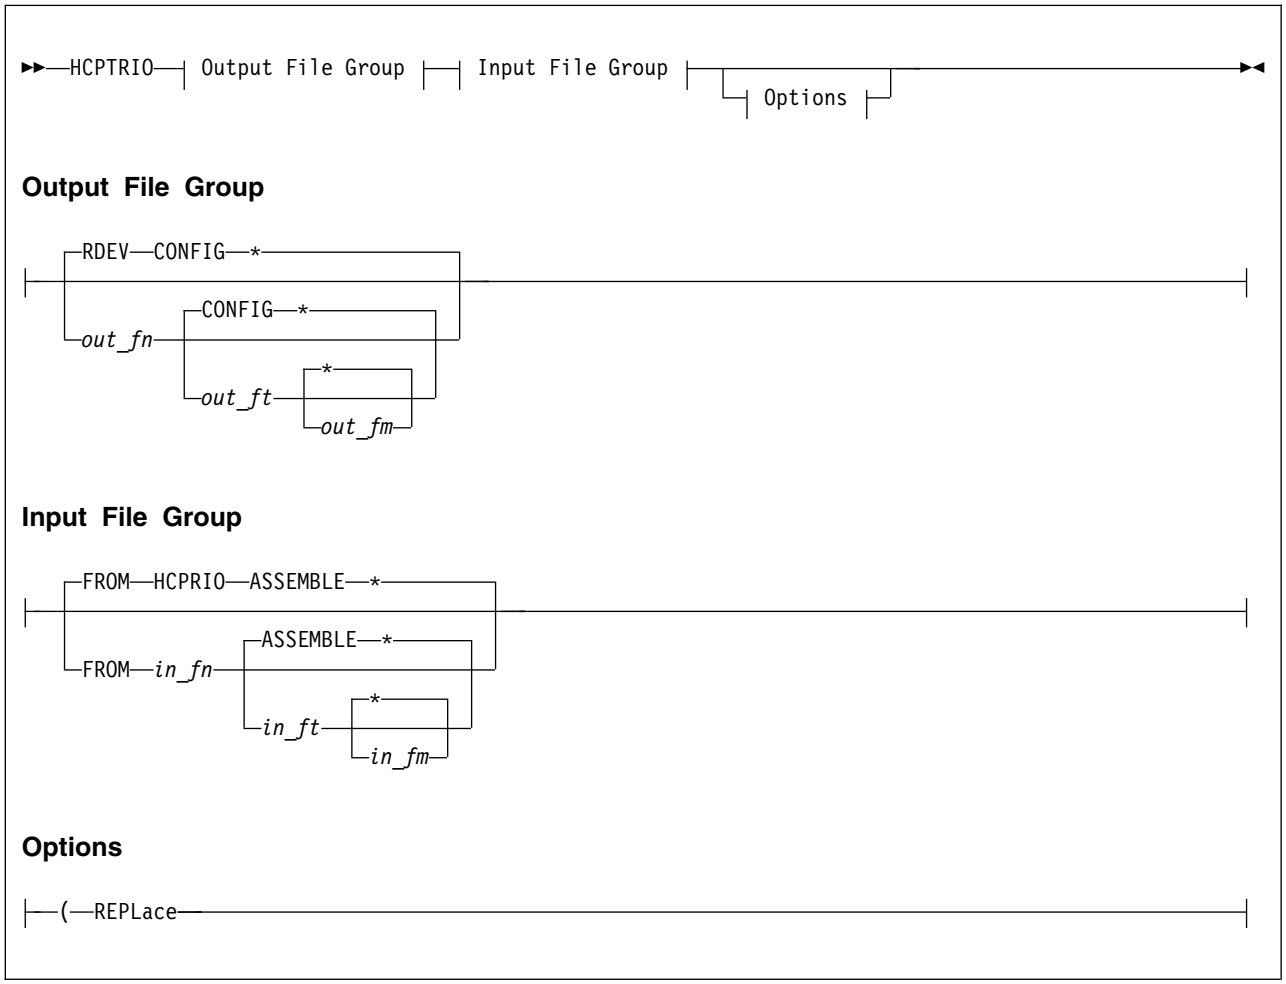

Authorization

Privilege Class: G

Purpose

Use HCPTRIO to generate a system configuration file (SYSTEM CONFIG) from an HCPRIO ASSEMBLE file. The generated file will contain statements for those real devices that do not answer sense ID requests or that do not return enough information.

Operands

out_fn
is the name of the CMS file to be generated. The default is RDEV.

out_ft
is the type of the file to be generated. The default is CONFIG.

out_fm
is the CMS file mode on which the file is to reside. The file mode must be accessed in write mode. An asterisk (*), the default, causes the HCPTRIO to use the first file mode accessed R/W.

FROM
tells HCPTRIO that there is a specific source file it should use.

in_fn
is the name of the CMS file that contains the input. The default is HCPRIO.

in_ft
is the type of the CMS file that contains the input. The default is ASSEMBLE.

in_fm
is the file mode on which the file resides. An asterisk (*), the default, causes HCPTRIO to use the first file mode on which the file resides.

REPLace
indicates that if the output file already exists, it should be replaced.

Usage Notes

1. HCPTRIO uses a work file having a file type of "$" followed by the first seven characters of the file type of the input file. If such a file already exists, HCPTRIO erases it. After writing the work file, HCPTRIO erases the original configuration file and renames the work file.

   Example:

   ```
   hcptrio rdev config a from hcprio assemble b (repl
   ```

   In this case, HCPTRIO writes a work file named RDEV $CONFIG A, erases the old RDEV CONFIG A, and renames RDEV $CONFIG A to RDEV CONFIG A.

Examples

Example 1:

```
hcptrio cambvm3 config a from hcprio assemble d
```

In this example, HCPTRIO uses the file HCPRIO ASSEMBLE D to generate a configuration file named CAMBVM3 CONFIG A.

Example 2:

```
hcptrio
```

In this example, HCPTRIO reads the first file it finds in the CMS search order that is named HCPRIO ASSEMBLE. It creates a configuration file named RDEV CONFIG on the first file mode letter in the CMS search order that is accessed R/W.

Return Codes

Code	Meaning
0	Successful execution
4	Disk not accessed or not accessed R/W
12	Output file already exists and the REPLACE option was not specified
16	Input file to translate was not found
20	I/O error in creating output file or in reading input file

24 Invalid option or too many options

HCPTSYS

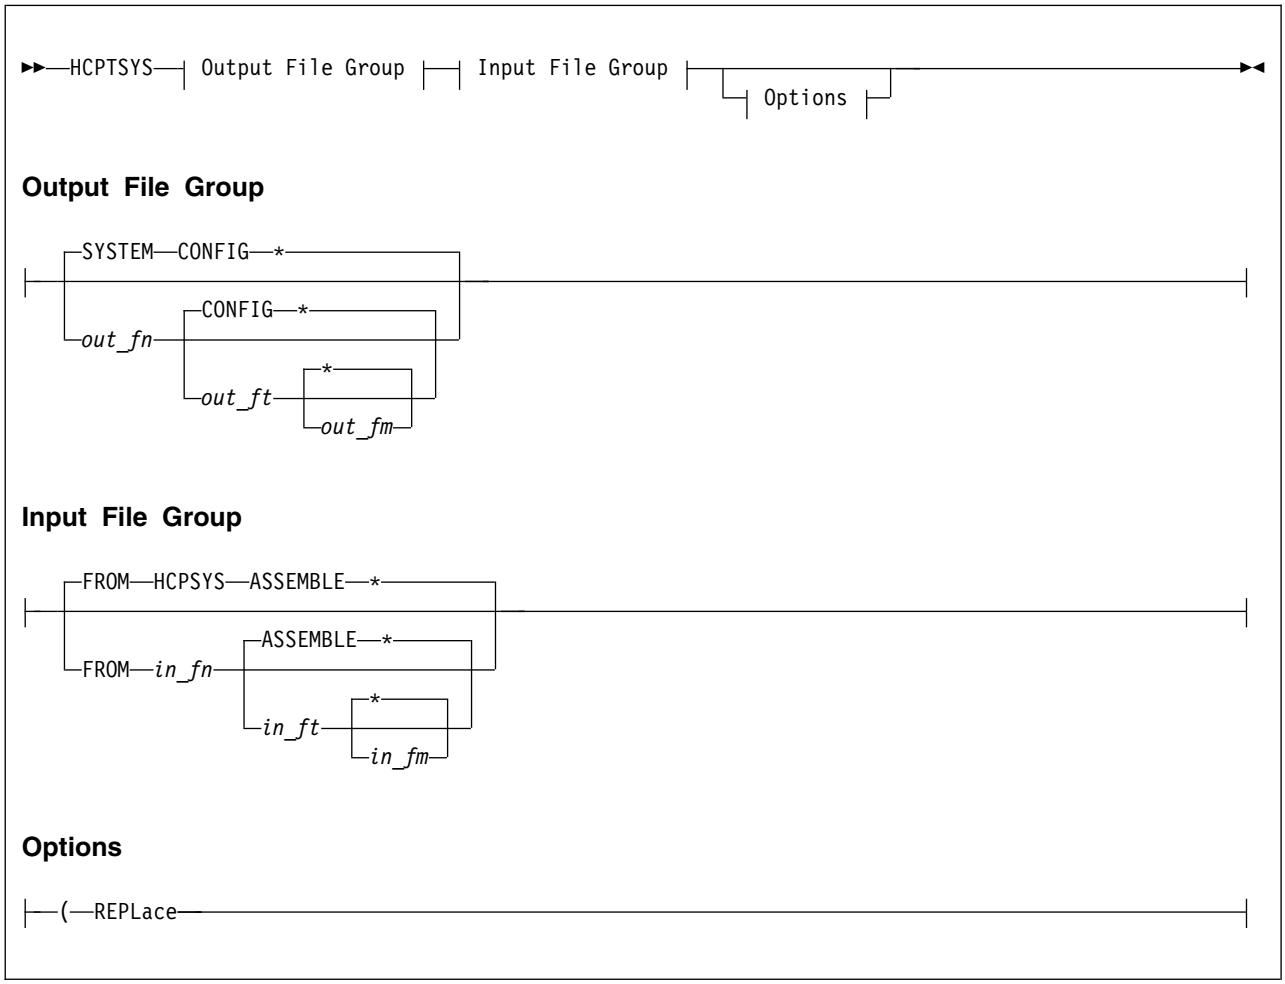

Authorization

Privilege Class: G

Purpose

Use HCPTSYS to generate a system configuration file (SYSTEM CONFIG) from an HCPSYS ASSEMBLE file.

Operands

out_fn
is the name of the CMS file to be generated. The default is SYSTEM.

out_ft
is the type of the file to be generated. The default is CONFIG.

out_fm
is the CMS file mode on which the file is to reside. The file mode must be accessed in write mode. An asterisk (*), the default, causes the HCPTSYS to use the first file mode accessed R/W.

FROM
tells HCPTSYS that there is a specific source file it should use.

in_fn
is the name of the CMS file that contains the input. The default is HCPSYS.

in_ft
is the type of the CMS file that contains the input. The default is ASSEMBLE.

in_fm
is the file mode on which the file resides. An asterisk (*), the default, causes HCPTSYS to use the first file mode on which the file resides.

REPLace
indicates that if the output file already exists, it should be replaced.

Usage Notes

1. HCPTSYS uses a work file having a file type of "$" followed by the first seven characters of the file type of the input file. If such a file already exists, HCPTSYS erases it. After writing the work file, HCPTSYS erases the original configuration file and renames the work file.

   Example:

   ```
   hcptsys system config a from hcpsys assemble b (repl
   ```

   In this case, HCPTRIO writes a work file named SYSTEM $CONFIG A, erases the old SYSTEM CONFIG A, and renames SYSTEM $CONFIG A to SYSTEM CONFIG A.

Examples

Example 1:

```
hcptsys cambvm3 config a from hcpsys assemble d
```

In this example, HCPTRIO uses the file HCPSYS ASSEMBLE D to generate a configuration file named CAMBVM3 CONFIG A.

Example 2:

```
hcptsys
```

In this example, HCPTSYS reads the first file it finds in the CMS search order that is named HCPSYS ASSEMBLE. It creates a configuration file named SYSTEM CONFIG on the first file mode letter in the CMS search order that is accessed R/W.

Return Codes

Code	Meaning
0	Successful execution
4	Disk not accessed or not accessed R/W
12	Output file already exists and the REPLACE option was not specified
16	Input file to translate was not found
20	I/O error in creating output file or in reading input file
24	Invalid option or too many options

Appendix B. [V6.3] Migrating to an Integrated xCAT MN and ZHCP Server in the Same CMA

If you are currently running separate xCAT and ZHCP virtual machines on your z/VM V6.3 system, you need to follow the instructions in this topic to upgrade to using one user ID for both servers in z/VM V6.3 before you migrate to z/VM V6.4. z/VM V6.4 does not support using separate virtual machines for the xCAT and ZHCP servers.

Figure 8 compares the services and virtual machines that run when you use z/VM's xCAT before and after installing VM65780 on z/VM V6.3. CMA and xCAT at earlier service levels correspond to the "before" picture, with the xCAT MN and ZHCP services running in two separate virtual machines. To run the same set of services after installing VM65780, you configure the CMA on z/VM 1 to run in the CMA "mn" role, and the CMA on z/VM 2 to run in the CMA "zhcp" role, as described in *z/VM: Enabling z/VM for OpenStack*.

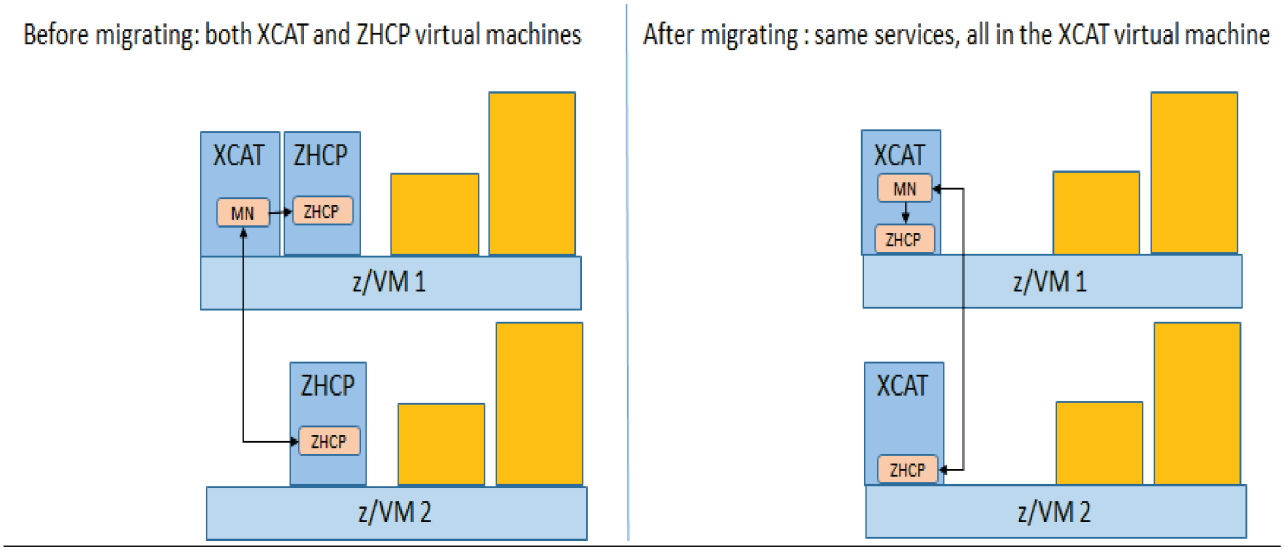

Figure 8. Services and Virtual Machines that Run when Using xCAT

To help with the transition from using two separate z/VM user IDs for xCAT and ZHCP to using one user ID, a Perl script called zxcatexport.pl can be used. This script can be run in the Script panel (see "xCAT Utilities and Common Procedures" in *z/VM: Systems Management Application Programming* for more information. To run zxcatexport, you must install APAR VM65780.

Note: These instructions are not valid for migrating from one level of the CMA (such as the Juno level) to another (such as the Liberty level). For those instructions, please see the appropriate CMAxx FILE that came with the APAR.

The zxcatexport.pl script backs up xCAT tables and other information, then exports the LVM mounted at `/install`. To prepare for using this script, you should make sure no xCAT or OpenStack activity is running. To stop OpenStack from communicating with xCAT, you must stop the compute services. On your compute node, issue:

```
service openstack-nova-compute-host stop
```

Because the zxcatexport.pl script stores some information on the currently active xCAT user ID `/install` LVM, you must have space available to hold the xCAT tables and some small command output. You can use the Linux **df –h /install** command to see how much free space you have. (The shipped xCAT tables with initial entries use about 8000 bytes.)

In the Script window, enter `/usr/local/bin/zxcatexport.pl` and press the Run button. The script completes and you will see output in the status window. The script performs these actions:

1. Creates an xcatmigrate directory with xcattables and fcpconfigs subdirectories.
2. Saves, in files, the volume group, Linux DASD, and z/VM virtual DASD information for the currently active xCAT user ID.
3. Dumps xCAT tables in the xcattables subdirectory in CSV (comma separated value) format. To see all the tables that xCAT has, use the GUI to view Configure-->Tables; this panel contains a short description and a link to display the tables.
4. Saves any ZHCP FCP pool configuration files (found in `/var/opt/zhcp/zfcp/*.conf`) in the fcpconfigs subdirectory.
5. Unmounts the `/install` LVM, makes it inactive, and marks it as exported.

Before Upgrading to Using One User ID for xCAT and ZHCP

Before you upgrade from using two separate z/VM user IDs for xCAT and ZHCP to using one user ID, you should do the following:

- Issue the following commands:

  ```
  SIGNAL SHUTDOWN XCAT WITHIN 60
  SIGNAL SHUTDOWN ZHCP WITHIN 60
  ```

- Change the directory for xCAT to have the current 100 (persistent) disk at a different address that is past 0102. For example, 0110.
- Change the directory for xCAT LVM disks from 200–20* to be past where they are now. For example: 240–24*

  (After you issue the zxcatimport.pl command copy the data from the old LVM disks to the new appliance LVM, you can umount the xCAT LVM and use the disks for something else.)
- You must add the OPTION LNKNOPAS statement to the xCAT directory entry. Before the migration, this statement is in the ZHCP directory entry.
- Install the CMA by following the instructions in APAR VMVM65780.
- Make sure you have new volume ID(s) for the CMA LVM that have as much space as the old xCAT LVM.
- The DMSSICNF COPY should not have the old xCAT LVM volumes in the ISO list. The old LVM is exported and will be added back later when you run zxcatimport.pl.

If you have any problems using zxcatexport.pl, or if you need to get the LVM disk(s) mounted again, issue the following commands in this order:

```
/sbin/vgimport xcat
/sbin/vgchange –a y xcat
mkdir –p /install
mount –t ext3 /dev/xcat/repo /install
```

See the description of the zxcatexport command in *z/VM: Systems Management Application Programming* for more information.

Importing Information from the Old xCAT User ID

After you have the appliance installed on the current xCAT user ID, use zxcatimport.pl to copy data from the old xCAT user ID. For example, you can copy data from the old LVM and persistent disks. See the description of the zxcimport command in *z/VM: Systems Management Application Programming* for more information.

It is recommended that before you modify any new xCAT appliance user ID tables using the zxcatimport.pl command that you back up your new appliance xCAT tables using dumpxCATdb. See the description of the dumpxCATdb command in *z/VM: Systems Management Application Programming* for more information.

Appendix C. [V6.4] Migrating an OpenStack Liberty CMA from z/VM V6.3 to z/VM V6.4

This section provides instructions for migrating a Liberty CMA running in z/VM 6.3 to z/VM 6.4. These steps should be followed just before you do the steps in the "Remove the Obsolete Release" chapter in *z/VM: Installation Guide*.

To ensure that you meet the prerequisites for migrating from a Liberty CMA in z/VM 6.3 to z/VM 6.4, you should do the following:

1. Determine if you are running xCAT or ZHCP outside the CMA, or running the xCAT MN and ZHCP in separate virtual machines. To determine this, on the MAINT 193 disk, in the file DMSSICMO COPY, check the value of the `openstack_system_role` property. If the value is "" (two double quotes), you are running xCAT outside the CMA, or running xCAT MN and ZHCP in separate virtual machines. Otherwise, you are already running a CMA and you can continue with the next step.

   If you are not running a CMA, you must first migrate to a Liberty CMA before you can migrate the Liberty CMA to z/VM 6.4. For instructions, see Appendix B, "[V6.3] Migrating to an Integrated xCAT MN and ZHCP Server in the Same CMA," on page 461. You can also follow the steps in the CMA120 FILE, shipped with the most recent service APARs, to migrate from your current environment to a Liberty CMA, with the CMA configured in the CMA mn role or the CMA zhcp role.

   You can find the current list of APARS at http://www.vm.ibm.com/sysman/osmntlvl.html.
2. Install the most recent Liberty service APARs.
3. If the value of the `openstack_system_role` property in the DMSSICMO COPY file on the MAINT 193 disk indicates that you are running a CMA in CMA controller role, CMA compute role, or CMA compute_mn role, determine the OpenStack release you are using:
   a. ssh into the XCAT virtual machine, using the user ID and password specified in DMSSICNF COPY properties `XCAT_MN_admin` and `XCAT_MN_pw`, respectively.
   b. Issue the following command:

      ```
      cat /opt/ibm/cmo/version
      ```

      You will see output similar to the following:

      ```
      CMO Appliance Version: 1.1.3-20151119 Built on: Thu Nov 19 04:51:02 EST 2015
      ```

      If the "CMO Appliance Version" string starts with a value of 1.2.0 or greater, you are running a Liberty CMA. In this case you must follow the steps in the CMA120 FILE, shipped with the latest service APARs, to install all APARs for the Liberty release. If this string begins with a value equal to or less than 1.1.3, as shown in the example, you are running an earlier version of the CMA. If the value is less than 1.1.3:

      1) In the same ssh session, issue the following command:

         ```
         sudo sh -c "cat /opt/ibm/cmo/version | awk '{print \$4}' > /data/.cmo_data_disk"
         ```

      2) Make sure you install all current APARs for the release you are using, and make sure the version string in `/data/.cmo_data_disk` starts with 1.1.3.

You can find the current list of APARS at http://www.vm.ibm.com/sysman/osmntlvl.html.

3) Follow the steps in the CMA120 FILE shipped with the latest service APARs to apply all the APARs for Liberty release, and to migrate from your current CMA to a Liberty CMA.

Also, note the following:

- Do not change the value of the `XCAT_zvmsysid propety` in the DMSSICNF COPY file.
- In z/VM 6.4, both the xCAT MN and the ZHCP services run within the OPNCLOUD virtual machine. At this point, you should not logon again to the XCAT virtual machine.

Now you are ready to migrate from running a Liberty CMA in z/VM 6.3 to z/VM 6.4. Follow these steps:

1. Make sure the CMA is running the latest level of Liberty service.
2. Shut down the XCAT virtual machine, if it is still running. For more information, see the "Stopping the Server Environment" section in the "Setting Up and Configuring the Server Environment" chapter in *z/VM: Systems Management Application Programming*.
3. Update the OPNCLOUD directory entry to migrate resources defined in the XCAT directory entry to the OPNCLOUD directory entry:
   a. Migrate DASDs defined in the XCAT directory entry to the OPNCLOUD directory entry. The virtual devices defined in the OPNCLOUD directory entry should be the same as those used in XCAT directory entry. This includes the following DASDs:
      - All the root disks in the XCAT virtual machine, including any disks you added manually to extend the root LVM partition.
      - All the LVM data disks in the range 0200-02FF.
      - Any DASDs associated with a SUBCONFIG statement in the XCAT directory entry. Determine which subconfiguration(s) your system is using by checking the existing XCAT directory entry. The entry shows you which subconfigurations are used when this identity virtual server is running on a particular zVM system.

        For example, if the XCAT directory contains the following SUBCONFIG statements:

        ```
        IDENTITY XCAT AUTOONLY 8G 8G G
        BUILD ON member1name USING SUBCONFIG XCAT-1
        BUILD ON member2name USING SUBCONFIG XCAT-2
        ```

        This means that membername1 is using the resources defined in the XCAT-1 subconfiguration, and membername2 is using the resources defined in the XCAT-2 subconfiguration.

        Continuing this example, suppose your OPNCLOUD directory entry contains the following statements:

        ```
        IDENTITY OPNCLOUD AUTOONLY 8G 8G G
        BUILD ON member1name USING SUBCONFIG OPNCLO-1
        BUILD ON member2name USING SUBCONFIG OPNCLO-2
        ```

        You must move the resources defined in the XCAT-1 subconfiguration to the OPNCLO-1 subconfiguration. Similarly, you must move the resources defined in the XCAT-2 subconfiguration to the OPNCLO-2 subconfiguration.

b. Add "OPTION LNKNOPAS" to the OPNCLOUD directory entry. To add this statement, logon to the MAINT user ID and issue the following command:

```
dirm for opncloud setoptn add lnknopas
```

c. Migrate to the OPNCLOUD directory entry any other options you manually added to the XCAT directory entry.

4. Issue the following commands to copy all the manually updated files you have on the xCAT 191 disk to the OPNCLOUD 191 disk:

```
LOGON MAINT
LINK XCAT 191 1191 RR
ACC 1191 O
LINK OPNCLOUD 191 2191 MR
ACC 2191 X
COPY FILENAME FILETYPE O = = X
REL X (DET
REL O (DET
```

5. Update the DMSSISVR NAMES file so that the OPNCLOUD virtual server is defined. Copy the DMSSISVR NAM$$001 file to the existing DMSSISVR NAMES file by issuing these commands:

```
LOGON MAINT or MIGMAINT
LINK MAINT640 493 1493 RR
ACC 1493 O
ACC 193 X
COPY DMSSISVR NAM$$001 O = NAMES X (REPL
REL X (DET
REL O (DET
```

6. Update the VSMWORK1 AUTHLIST file to authorize the OPNCLOUD virtual server to issue SMAPI requests. (For more information, see the "Authorizing API Requests" section in the "Setting Up and Configuring the Server Environment" chapter in *z/VM: Systems Management Application Programming*.) You must add the OPNCLOUD user ID to your VSMWORK1 AUTHLIST file on the VMSYS:VSMWORK1 disk.

   For example, suppose you want to access VMSYS:VSMWORK1 as file mode G, and you want to edit the VSMWORK1 AUTHLIST file in read/write mode. Issue the following commands:

```
ACCESS VMSYS:VSMWORK1. G (FORCERW
XEDIT VSMWORK1 AUTHLIST G
```

   In the VSMWORK1 AUTHLIST file, change "XCAT" to "OPNCLOUD".

7. Update the DMSSICNF COPY file on the MAINT 193 disk. In this file, make sure the value of the `XCAT_User property` is updated to be "OPNCLOUD". When the CMA starts up, it will automatically update the user ID configured on the XCAT node.

8. After you complete these migration steps, see the "Starting and Restarting the Server Environment" section in the "Setting Up and Configuring the Server Environment" chapter in *z/VM: Systems Management Application Programming* for information on restarting the SMAPI servers.

9. Once the CMA is running successfully, clean up any artifacts (files, directory entries, subconfigurations, etc.) you no longer need. For example, you can delete the XCAT directory entry and any subconfigurations it was using.

Appendix D. [V6.4 APAR] xCAT Management Network Changes in the OpenStack Newton CMA

Note: In this section, the term "powered off" refers to the situation where a guest is not running or is logged off.

Prior to the OpenStack Newton CMA, the xCAT Management Node (MN) service always used ssh as the mechanism to interact with and manage deployed guests. Starting with the Newton CMA, the xCAT MN is capable of using IUCV instead of ssh. The xCAT MN attempts to use IUCV when possible because IUCV is usually faster and provides for better separation of the management network from the data network(s). In most common cases the transition is fully automated, and requires no manual intervention.

As a consequence of the transition to IUCV, IBM has deprecated the DMSSICMO COPY file properties `openstack_xcat_mgt_ip` and `openstack_xcat_mgt_mask`. This section helps you understand if and how this affects your installation. The properties were optional, so you may not need to use them at all; for example, if your deployed instances use public IP addresses reachable from the xCAT MN's IP address, then you never had a need to use the deprecated properties.

There are several categories of guests, each of which behaves somewhat differently:

- OpenStack-deployed guests – unless you used Alternative Deployment Provisioning to provision them, they are migrated to IUCV automatically the first time that the xCAT MN interacts with them. Unless the guests are powered off, this typically occurs when the CMA is restarted after upgrading to Newton.

  Guests you deployed using Alternative Deployment Provisioning will continue to use ssh to interact with the xCAT MN. They are not yet eligible to use IUCV.
- Guests deployed using xCAT instead of OpenStack – as long as xCAT can communicate with them using ssh, they are migrated to IUCV automatically the first time that the xCAT MN interacts with them. Unless the guests are powered off or have IP addresses that are unreachable from xCAT, migration typically occurs when the CMA is restarted after upgrading to Newton. This category includes guests that you initially provisioned using xCAT and later made visible to OpenStack using xCAT's discovery function.

You can determine whether the xCAT MN is using ssh or IUCV to manage each guest using the xCAT GUI, or by using ssh commands while logged in as the xCAT MN administrator. In the xCAT GUI, log in and select Configure in the top navigation bar. Then select "zvm" near the bottom of the Tables page, and scroll right (if necessary) to display the status column. If you prefer to use ssh, you should ssh into the IP address specified on the `xcat_mn_ip` property in the DMSSICNF COPY file. Log in as the user ID specified on the `xcat_mn_admin` property in the DMSSICNF COPY file, and run the command `/usr/sbin/tabdump zvm` from the logon shell. Regardless of the method you use, the possible status values displayed are as follows:

- `IUCV=1` indicates that the xCAT MN has successfully used IUCV to manage the guest. The xCAT MN will continue to use IUCV in the future.
- `SSH=1` indicates that the xCAT MN is using ssh to manage the guest, and will continue to do so unless manual intervention occurs. The guest might currently

be ineligible to use IUCV; for example you might have deployed it using xCAT directly rather than by using OpenStack or Alternative Deployment Provisioning.

- The absence of a value means that the xCAT MN has not yet established connectivity with the guest since being upgraded to Newton. For example, the guest might be powered off. xCAT will continue to periodically attempt to establish communications with the guest, starting with using IUCV if the guest is eligible, and falling back to using ssh if necessary. Once connectivity is established, the xCAT will move the guest to one of the other status values (`IUCV=1` or `SSH=1`).

Once all your current guests are using IUCV, and if you plan to deploy only IUCV-eligible categories of guests in the future, you can remove the deprecated DMSSICMO COPY file properties `openstack_xcat_mgt_ip` and `openstack_xcat_mgt_mask`.

Additional information related to this topic, if available, can be found on the Web at http://www.vm.ibm.com/sysman/osmntlvl.html.

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).

Adobe is either a registered trademark or a trademark of Adobe Systems Incorporated in the United States, and/or other countries.

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

Windows is a trademark of Microsoft Corporation in the United States, other countries, or both.

Java is a trademark or registered trademark of Oracle and/or its affiliates.

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.

Bibliography

See the following publications for additional information about z/VM. For abstracts of the z/VM publications, see *z/VM: General Information*, GC24-6193.

Where to Get z/VM Information

z/VM product documentation and other z/VM information is available in IBM Knowledge Center - z/VM (www.ibm.com/support/knowledgecenter/SSB27U).

You can also obtain z/VM product publications from IBM Publications Center (www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss).

z/VM Base Library

Overview

- *z/VM: License Information*, GC24-6200
- *z/VM: General Information*, GC24-6193
- *z/VM: Glossary*, GC24-6195

Installation, Migration, and Service

- *z/VM: Installation Guide*, GC24-6246
- *z/VM: Migration Guide*, GC24-6201
- *z/VM: Service Guide*, GC24-6247
- *z/VM: VMSES/E Introduction and Reference*, GC24-6243

Planning and Administration

- *z/VM: CMS File Pool Planning, Administration, and Operation*, SC24-6167
- *z/VM: CMS Planning and Administration*, SC24-6171
- *z/VM: Connectivity*, SC24-6174
- *z/VM: CP Planning and Administration*, SC24-6178
- *z/VM: Enabling z/VM for OpenStack (Support for OpenStack Liberty Release)*, SC24-6251
- *z/VM: Getting Started with Linux on z Systems*, SC24-6194
- *z/VM: Group Control System*, SC24-6196
- *z/VM: I/O Configuration*, SC24-6198
- *z/VM: Running Guest Operating Systems*, SC24-6228
- *z/VM: Saved Segments Planning and Administration*, SC24-6229
- *z/VM: Secure Configuration Guide*, SC24-6230
- *z/VM: TCP/IP LDAP Administration Guide*, SC24-6236
- *z/VM: TCP/IP Planning and Customization*, SC24-6238
- *z/OS and z/VM: Hardware Configuration Manager User's Guide*, SC34-2670

Customization and Tuning

- *z/VM: CP Exit Customization*, SC24-6176
- *z/VM: Performance*, SC24-6208

Operation and Use

- *z/VM: CMS Commands and Utilities Reference*, SC24-6166
- *z/VM: CMS Primer*, SC24-6172
- *z/VM: CMS User's Guide*, SC24-6173
- *z/VM: CP Commands and Utilities Reference*, SC24-6175
- *z/VM: System Operation*, SC24-6233
- *z/VM: TCP/IP User's Guide*, SC24-6240
- *z/VM: Virtual Machine Operation*, SC24-6241
- *z/VM: XEDIT Commands and Macros Reference*, SC24-6244
- *z/VM: XEDIT User's Guide*, SC24-6245

Application Programming

- *z/VM: CMS Application Development Guide*, SC24-6162
- *z/VM: CMS Application Development Guide for Assembler*, SC24-6163
- *z/VM: CMS Application Multitasking*, SC24-6164
- *z/VM: CMS Callable Services Reference*, SC24-6165
- *z/VM: CMS Macros and Functions Reference*, SC24-6168
- *z/VM: CMS Pipelines User's Guide and Reference*, SC24-6252
- *z/VM: CP Programming Services*, SC24-6179
- *z/VM: CPI Communications User's Guide*, SC24-6180
- *z/VM: Enterprise Systems Architecture/Extended Configuration Principles of Operation*, SC24-6192

- *z/VM: Language Environment User's Guide*, SC24-6199
- *z/VM: OpenExtensions Advanced Application Programming Tools*, SC24-6202
- *z/VM: OpenExtensions Callable Services Reference*, SC24-6203
- *z/VM: OpenExtensions Commands Reference*, SC24-6204
- *z/VM: OpenExtensions POSIX Conformance Document*, GC24-6205
- *z/VM: OpenExtensions User's Guide*, SC24-6206
- *z/VM: Program Management Binder for CMS*, SC24-6211
- *z/VM: Reusable Server Kernel Programmer's Guide and Reference*, SC24-6220
- *z/VM: REXX/VM Reference*, SC24-6221
- *z/VM: REXX/VM User's Guide*, SC24-6222
- *z/VM: Systems Management Application Programming*, SC24-6234
- *z/VM: TCP/IP Programmer's Reference*, SC24-6239
- *Common Programming Interface Communications Reference*, SC26-4399
- *Common Programming Interface Resource Recovery Reference*, SC31-6821
- *z/OS: IBM Tivoli Directory Server Plug-in Reference for z/OS*, SA76-0169
- *z/OS: Language Environment Concepts Guide*, SA22-7567
- *z/OS: Language Environment Debugging Guide*, GA22-7560
- *z/OS: Language Environment Programming Guide*, SA22-7561
- *z/OS: Language Environment Programming Reference*, SA22-7562
- *z/OS: Language Environment Run-Time Messages*, SA22-7566
- *z/OS: Language Environment Writing Interlanguage Communication Applications*, SA22-7563
- *z/OS MVS Program Management: Advanced Facilities*, SA23-1392
- *z/OS MVS Program Management: User's Guide and Reference*, SA23-1393

Diagnosis

- *z/VM: CMS and REXX/VM Messages and Codes*, GC24-6161
- *z/VM: CP Messages and Codes*, GC24-6177
- *z/VM: Diagnosis Guide*, GC24-6187
- *z/VM: Dump Viewing Facility*, GC24-6191
- *z/VM: Other Components Messages and Codes*, GC24-6207
- *z/VM: TCP/IP Diagnosis Guide*, GC24-6235
- *z/VM: TCP/IP Messages and Codes*, GC24-6237
- *z/VM: VM Dump Tool*, GC24-6242
- *z/OS and z/VM: Hardware Configuration Definition Messages*, SC34-2668

z/VM Facilities and Features

Data Facility Storage Management Subsystem for VM

- *z/VM: DFSMS/VM Customization*, SC24-6181
- *z/VM: DFSMS/VM Diagnosis Guide*, GC24-6182
- *z/VM: DFSMS/VM Messages and Codes*, GC24-6183
- *z/VM: DFSMS/VM Planning Guide*, SC24-6184
- *z/VM: DFSMS/VM Removable Media Services*, SC24-6185
- *z/VM: DFSMS/VM Storage Administration*, SC24-6186

Directory Maintenance Facility for z/VM

- *z/VM: Directory Maintenance Facility Commands Reference*, SC24-6188
- *z/VM: Directory Maintenance Facility Messages*, GC24-6189
- *z/VM: Directory Maintenance Facility Tailoring and Administration Guide*, SC24-6190

Open Systems Adapter/Support Facility

- *Open Systems Adapter-Express Customer's Guide and Reference*, SA22-7935
- *Open Systems Adapter-Express Integrated Console Controller User's Guide*, SA22-7990
- *Open Systems Adapter-Express Integrated Console Controller 3215 Support*, SA23-2247
- *Open Systems Adapter-Express3 Integrated Console Controller Dual-Port User's Guide*, SA23-2266

Performance Toolkit for VM

- *z/VM: Performance Toolkit Guide*, SC24-6209
- *z/VM: Performance Toolkit Reference*, SC24-6210

RACF Security Server for z/VM

- *z/VM: RACF Security Server Auditor's Guide*, SC24-6212
- *z/VM: RACF Security Server Command Language Reference*, SC24-6213

- *z/VM: RACF Security Server Diagnosis Guide*, GC24-6214
- *z/VM: RACF Security Server General User's Guide*, SC24-6215
- *z/VM: RACF Security Server Macros and Interfaces*, SC24-6216
- *z/VM: RACF Security Server Messages and Codes*, GC24-6217
- *z/VM: RACF Security Server Security Administrator's Guide*, SC24-6218
- *z/VM: RACF Security Server System Programmer's Guide*, SC24-6219
- *z/VM: Security Server RACROUTE Macro Reference*, SC24-6231

Remote Spooling Communications Subsystem Networking for z/VM

- *z/VM: RSCS Networking Diagnosis*, GC24-6223
- *z/VM: RSCS Networking Exit Customization*, SC24-6224
- *z/VM: RSCS Networking Messages and Codes*, GC24-6225
- *z/VM: RSCS Networking Operation and Use*, SC24-6226
- *z/VM: RSCS Networking Planning and Configuration*, SC24-6227

Prerequisite Products

Device Support Facilities

- *Device Support Facilities: User's Guide and Reference*, GC35-0033

Environmental Record Editing and Printing Program

- *Environmental Record Editing and Printing Program (EREP): Reference*, GC35-0152
- *Environmental Record Editing and Printing Program (EREP): User's Guide*, GC35-0151

Index

Special characters

Numerics

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X

Y

Z

Product Number: 5741-A07

Printed in USA

GC24-6201-12