



IBM Z
Dynamic Partition Manager
(DPM) Guide

SB10-7170-02

Level 02a





IBM Z

Dynamic Partition Manager (DPM) Guide

SB10-7170-02

Level 02a

Note

Before you use this information and the product it supports, read the information in “Safety” on page ix, Appendix B, “Notices,” on page 139 and *IBM Systems Environmental Notices and User Guide*, Z125-5823.

This edition, SB10-7170-02a, applies to IBM Z and IBM LinuxONE (LinuxONE) servers. This edition replaces SB10-7170-02.

There might be a newer version of this document in a **PDF** file available on **Resource Link**. Go to <http://www.ibm.com/servers/resourcelink> and click **Library** on the navigation bar.

© **Copyright IBM Corporation 2017, 2019.**

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

Contents

Figures	v
--------------------------	----------

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

Safety	ix
-------------------------	-----------

Safety notices	ix
World trade safety information	ix
Laser safety information	ix
Laser compliance	ix

About this publication	xi
---	-----------

Intended audience	xi
Prerequisite and related information	xi
Related HMC and SE console information	xii
Related information for Linux on Z	xii
Related information for z/VM	xii
Related information for IBM Secure Service Container	xii
How to use this publication	xii
Accessibility	xiii
How to send your comments	xiii

Summary of changes	xv
-------------------------------------	-----------

Part 1. Introduction to Dynamic Partition Manager.	1
---	----------

Chapter 1. Dynamic Partition Manager: A quicker and easier way to deploy Linux servers on a mainframe	3
--	----------

Chapter 2. Getting started: Creating a new partition and starting your Linux server	5
--	----------

Part 2. Basic concepts and terms for Linux administrators	9
--	----------

Chapter 3. Partitions: Virtual images of a mainframe system	11
--	-----------

Chapter 4. Adapters: Connections to networks, storage, and more	17
--	-----------

Chapter 5. Operating systems and hypervisors.	23
--	-----------

Partition types and supported operating systems and hypervisors	23
Installing an operating system or hypervisor in a partition	24

Working with an installed operating system or hypervisor.	27
--	----

Part 3. Basic tasks for Linux administrators.	29
--	-----------

Chapter 6. Planning checklist for creating a partition	31
---	-----------

Chapter 7. Creating and starting a new partition	35
---	-----------

Selecting which New Partition task mode to use	35
Creating a new partition	37
Starting a partition and its operating system or hypervisor.	48

Chapter 8. Dynamically modifying the resources of an active partition	53
--	-----------

Chapter 9. Summary of tasks for managing systems, adapters, and partitions.	65
--	-----------

Part 4. Topics for system planners	69
---	-----------

Chapter 10. Prerequisites for using Dynamic Partition Manager	71
--	-----------

Chapter 11. Migrating DPM configuration data to a new system	79
---	-----------

Migration instructions for an MES frame upgrade	79
Migration instructions for a system replacement	82

Chapter 12. Upgrading from DPM R3.0 to R3.1.	89
---	-----------

Chapter 13. Adapter configuration	93
--	-----------

Part 5. Topics for network administrators.	99
---	-----------

Chapter 14. Customizing the network boot server environment 101

Chapter 15. Securing FTP operations 105

Chapter 16. Using VLAN IDs to secure network communications 107

Part 6. Topics for storage administrators 111

Chapter 17. Managing storage on a DPM-enabled system 113

A simplified and improved approach to storage management.	113
Setting up storage resources for the first time.	118
Authorizing users of the Configure Storage task	118
Connecting a system to storage	119

Requesting storage resources for partitions to use	123
Managing storage resources	125
Viewing and managing storage groups	126
Attaching storage resources to a partition	130
Modifying the storage configuration of the system	133

Part 7. Appendixes 135

Appendix A. DPM task and resource roles 137

Appendix B. Notices 139

Trademarks	140
Class A Notices.	140

Index 145

Figures

1. New user interface design for the HMC Welcome page	3	27. Sample display for the lscpu command showing configured processor resources	55
2. Portion of the main HMC page, with the Systems Management node expanded, and a system and the New Partition task highlighted	5	28. Sample display for the lsmem command	56
3. Welcome page of the New Partition task in basic mode	6	29. Sample screen of the Memory section in the Partition Details window	56
4. Sample validation window for the New Partition task	7	30. Sample display for the chmem command	57
5. Methods of starting a new partition	7	31. Sample display for the lscss command showing current devices	57
6. Portion of the Start window for a new partition named ASL_Test	8	32. Sample screen of the Network section with the New NIC icon highlighted	58
7. Sample Linux command and display of configured processor resources	8	33. Sample screen of the New Network Interface Card window	58
8. Partitions configured on a DPM-enabled system	12	34. Sample screen of the Network section with the newly added NIC	59
9. Partitions with NICs configured for access to network adapters on a DPM-enabled system	18	35. Sample display for the znetconf command showing current devices	59
10. Sample physical storage configuration with DPM R3.1	19	36. Sample display for the lsqeth command showing details for devices	60
11. Sample storage groups that are attached to partitions on the system	20	37. Sample hardware message that shows the results of a successful DPM conversion of migrated configuration information	80
12. Using HBAs to access storage adapters and the storage network with DPM R3.0 or earlier	21	38. Partition access to storage before and after migration	84
13. Example: Linux servers configured in partitions on a DPM-enabled system	23	39. Partition access to storage before and after upgrading to DPM R3.1	90
14. Welcome page of the New Partition task in basic mode.	38	40. Sample screen of the Manage Adapters task	94
15. Processors page of the New Partition task in basic mode.	39	41. Sample screen of the Adapter Details window, with Name and Description fields highlighted	95
16. Memory page of the New Partition task in basic mode.	40	42. Sample screen of the Manage Adapters task, with the expanded Actions list	95
17. Network page of the New Partition task in basic mode.	41	43. Sample screen of the New HiperSockets Adapter window.	96
18. Boot page of the New Partition task in basic mode	44	44. Sample screen of the new HiperSockets adapter in the Adapter Details window	97
19. Sample validation window for the New Partition task	47	45. Sample VLAN configuration in partitions and guest operating systems	107
20. Sample displays resulting from the lscpu command	50	46. Sample frame and card display on the Configure Storage Cards page	114
21. Sample displays resulting from the lsmem command	51	47. Photograph of a physical adapter card in a system I/O drawer	114
22. Sample displays resulting from the lsqeth and lscss commands	51	48. Sample display of two sites with physical storage hardware configured for the primary site	115
23. Sample displays resulting from the lschp command	52	49. Sample storage groups that are attached to partitions on the system	117
24. Sample display for the lscpu command showing current processor resources	54	50. Sample frame and card display on the Configure Storage Cards page	120
25. Sample screen of the Processors section in the Partition Details window	54	51. Sample display of two sites with physical storage hardware configured for the primary site	121
26. Sample display for the lscpu command showing additional processor resources	55		

Tables

- | | | | |
|--|----|--|-----|
| 1. Planning checklist for defining your new partition | 31 | 5. Summary of DPM releases and their associated EC MCLs | 72 |
| 2. Comparison of editable partition properties in the basic and advanced New Partition task modes | 35 | 6. Supported migration paths for exporting and importing a DPM configuration | 82 |
| 3. Summary of key HMC/SE tasks and displays for working with DPM-enabled systems, partitions, and adapters | 65 | 7. Sample default names for installed adapters | 93 |
| 4. Required feature codes for IBM Dynamic Partition Manager (DPM) | 71 | 8. Affect of fulfillment status on the Volume tab display. | 128 |
| | | 9. DPM task roles mapping legend | 137 |
| | | 10. DPM task roles mapping. | 137 |

Safety

Safety notices

Safety notices may be printed throughout this guide. **DANGER** notices warn you of conditions or procedures that can result in death or severe personal injury. **CAUTION** notices warn you of conditions or procedures that can cause personal injury that is neither lethal nor extremely hazardous. **Attention** notices warn you of conditions or procedures that can cause damage to machines, equipment, or programs.

World trade safety information

Several countries require the safety information contained in product publications to be presented in their translation. If this requirement applies to your country, a safety information booklet is included in the publications package shipped with the product. The booklet contains the translated safety information with references to the US English source. Before using a US English publication to install, operate, or service this product, you must first become familiar with the related safety information in the *Systems Safety Notices*, G229-9054. You should also refer to the booklet any time you do not clearly understand any safety information in the US English publications.

Laser safety information

All IBM® Z® (Z) and IBM LinuxONE™ (LinuxONE) models can use I/O cards such as FICON®, Open Systems Adapter (OSA), InterSystem Channel-3 (ISC-3), zHyperLink Express, or other I/O features which are fiber optic based and utilize lasers (short wavelength or long wavelength lasers).

Laser compliance

All lasers are certified in the US to conform to the requirements of DHHS 21 CFR Subchapter J for Class 1 or Class 1M laser products. Outside the US, they are certified to be in compliance with IEC 60825 as a Class 1 or Class 1M laser product. Consult the label on each part for laser certification numbers and approval information.

CAUTION: Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, never look into the end of an optical fiber cable or open receptacle. (C027)

CAUTION: This product contains a Class 1M laser. Do not view directly with optical instruments. (C028)

About this publication

This book contains general concepts, planning guidelines, instructions, and reference information for creating and managing partitions on systems that are running in Dynamic Partition Manager (DPM) mode.

DPM is a feature of the following IBM Z[®] and IBM[®] LinuxONE (LinuxONE) servers:

- IBM z14[™] (z14) (machine type 3906 or 3907)
- IBM z13[®] (z13[®]) or IBM z13s[®] (z13s[®])
- IBM LinuxONE Emperor II[™] (Emperor II), or IBM LinuxONE Rockhopper II (Rockhopper II)
- IBM LinuxONE Emperor[™] (Emperor) or IBM LinuxONE Rockhopper[™] (Rockhopper) machine type 2965

This book describes DPM Version 3, which is available on the z14, Emperor II, and Rockhopper II systems, starting with Hardware Management Console / Support Element (HMC/SE) Version 2.14.0.

- For additional details about DPM Version 3, see Chapter 10, “Prerequisites for using Dynamic Partition Manager,” on page 71.
- For earlier versions of DPM on earlier systems and HMC/SE versions, see *IBM Dynamic Partition Manager Guide*, SB10-7168.

Partitions on a DPM-enabled system support the following operating systems and hypervisors:

- Various Linux distributions, which are listed on the IBM tested platforms page for Linux environments. These distributions include supported versions of Red Hat Enterprise Linux (RHEL), SUSE Linux Enterprise Server (SLES), and Ubuntu Server (KVM or LPAR DPM).
- z/VM[®] 6.4 or later. z/VM is an IBM virtualization hypervisor on which you can run multiple Linux images.

DPM also supports IBM Secure Service Container partitions.

Figures included in this document illustrate concepts and are not necessarily accurate in content, appearance, or specific behavior.

Intended audience

The primary audience for this book is system administrators with knowledge of virtualization concepts but limited mainframe skills. These administrators are responsible for creating and managing partitions on a DPM-enabled system, to support business applications.

Other IT personnel who might benefit from using this book include mainframe operators, automation engineers, and experienced mainframe system administrators who need to learn how to use DPM tasks that are available starting with Hardware Management Console / Support Element (HMC/SE) Version 2.14.0.

Prerequisite and related information

This book describes how to plan for, create, and manage partitions on DPM-enabled systems. To create and manage partitions and their resources, system administrators use specific tasks on the HMC. These tasks can be accomplished through a program as well, with the HMC Web Services application programming interfaces (APIs) for DPM.

- For more information about a specific system, see the appropriate system overview on the IBM Redbooks® website at <http://www.redbooks.ibm.com/> (for example, for the IBM z14 (z14), see *IBM z14 Technical Guide*, SG24-8451).
- For information about the DPM APIs, see the appropriate edition of *Hardware Management Console Web Services API*, which is available through the Library link on IBM Resource Link® at <http://www.ibm.com/servers/resourcelink>

HMC users also can monitor and manage systems and partitions through the IBM HMC Mobile for Z and LinuxONE (HMC Mobile) mobile app for iOS and Android. The systems can either run in standard mode (that is, with Processor Resource/System Manager or PR/SM™), or run with DPM enabled. The HMC Mobile app provides system and partition views, status monitoring, hardware messages, operating system messages, and the ability to receive push notifications from the HMC, using the existing support server connection. For more information, see the HMC online help for the **HMC Mobile Settings** task.

Related HMC and SE console information

Hardware Management Console (HMC) and Support Element (SE) information can be found on the console help system.

Related information for Linux on Z

For information about installing and running a Linux distribution on an IBM Z or LinuxONE server, see the Linux on IBM Systems topics in IBM Knowledge Center, at: https://www.ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz_r_lib.html

Another useful source of information is the IBM developerWorks® site for Linux at http://www.ibm.com/developerworks/linux/linux390/documentation_dev.html

Related information for z/VM

For information about installing and running z/VM as a hypervisor for Linux guests on an IBM Z or LinuxONE server, see the documentation for the version of z/VM that you are installing. The z/VM library is available at <http://www.vm.ibm.com/library/>

Related information for IBM Secure Service Container

For information about working with Secure Service Container partitions and the appliances they support, see *Secure Service Container User's Guide*, SC28-6978, which is available on <http://www.ibm.com/servers/resourcelink>.

How to use this publication

This book provides an overview of DPM, lists the system requirements for its use, and contains step-by-step instructions for system administrators who create, start, and manage partitions on a DPM-enabled system.

Topics are organized in parts, for the audience who is most likely to use the information:

Part 1: Introduction to the IBM Dynamic Partition Manager (DPM)

Topics in this part describe DPM, explain the benefits of using it, and provide a basic set of instructions for creating and starting a new partition on a DPM-enabled system. These topics are appropriate for readers who need a quick introduction to DPM.

Part 2, “Basic concepts and terms for Linux administrators,” on page 9

Topics in this part explain mainframe and DPM concepts and terms that might be unfamiliar to system administrators who have little or no experience working in a mainframe environment.

Part 3, “Basic tasks for Linux administrators,” on page 29

Topics in this part provide a planning checklist and step-by-step instructions for creating and starting a partition, and its operating system or hypervisor, on a DPM-enabled system. Also included is a summary of HMC tasks for monitoring and managing partitions, adapters, and other resources on a DPM system. These topics are appropriate for any administrator who creates or manages DPM partitions.

Part 4, “Topics for system planners,” on page 69

Topics in this part provide the prerequisites for enabling DPM on a mainframe system, information about supported functions, and the engineering changes (ECs) or machine change levels (MCLs) for upgrading to the latest DPM version. This part also includes migration instructions and information about I/O adapter configuration. These topics are appropriate for experienced system planners and other administrators who are familiar with mainframe systems.

Part 5, “Topics for network administrators,” on page 99

Topics in this part provide information about network-related concepts and tasks that are specific to working with DPM-enabled systems.

Part 6, “Topics for storage administrators,” on page 111

Topics in this part provide information about storage-related concepts and tasks that are specific to working with DPM-enabled systems.

Appendixes

Topics in this part include the default task roles for the HMC tasks for DPM, and trademarks.

Accessibility

IBM strives to provide products with usable access for everyone, regardless of age or ability.

Accessible publications for this product are offered in EPUB format and can be downloaded from Resource Link at <http://www.ibm.com/servers/resourcelink>.

If you experience any difficulty with the accessibility of any IBM Z® (Z) and IBM LinuxONE (LinuxONE) information, go to Resource Link at <http://www.ibm.com/servers/resourcelink> and click **Feedback** from the navigation bar on the left. In the **Comments** input area, state your question or comment, the publication title and number, choose **General comment** as the category and click **Submit**. You can also send an email to reslink@us.ibm.com providing the same information.

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

How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information. Send your comments by using Resource Link at <http://www.ibm.com/servers/resourcelink>. Click **Feedback** on the navigation bar on the left. You can also send an email to reslink@us.ibm.com. Be sure to include the name of the book, the form number of the book, the version of the book, if applicable, and the specific location of the text you are commenting on (for example, a page number, table number, or a heading).

Summary of changes

For the most recent edition only, technical changes to the text are indicated by a vertical bar (|) to the left of the change.

Summary of changes for SB10-7170-02a

The topic “Boot options for z/VM partitions” on page 25 has been updated to include the ISO image option, which is available with the engineering change (EC) machine change levels (MCLs) that are included in bundle S10.

Summary of changes for SB10-7170-02

This edition contains the following updates. If a specific update is available through an engineering change (EC) machine change level (MCL) for the host system, the MCL or bundle number is provided in Chapter 10, “Prerequisites for using Dynamic Partition Manager,” on page 71.

New information

The topic Chapter 10, “Prerequisites for using Dynamic Partition Manager,” on page 71 has been updated to include the following information.

- A list of DPM enhancements that are available with Hardware Management Console / Support Element (HMC/SE) Version 2.14.1. See Table 5 on page 72.
- The requirement for storage area networks to support N Port Identifier Virtualization (NPIV) mode for FCP adapters. See “About storage adapters” on page 76.
- The required partition boot options and installation methods for z/VM V6.4 or later. See the z/VM information in “Requirements for the hypervisor or operating system” on page 76.

Changed information

The topics in Chapter 17, “Managing storage on a DPM-enabled system,” on page 113 have been updated to include information about the following enhancements:

- Storage group templates: To make storage requests even easier to complete, administrators can create a master copy, or template, for a FICON or FCP storage group. Templates can reflect typical usage patterns (production, staging, test); standardize storage for specific user groups; or document company requirements or restrictions that might be in place for storage use.
- Connection Report for FICON storage groups: Use the **Connection Report** to view the logical paths that are in use for a FICON storage group.
- HISTORY tab on the Storage Group details page: The **HISTORY** tab lists the actions that users have taken for this storage group. The most recent action is listed at the top of the History table. Information in the ACTION column not only briefly describes the activity, but also preserves details such as requests that were sent to storage administrators for fulfillment.

Topics throughout this book also include minor editorial and technical updates.

Summary of changes for SB10-7170-01c

This edition contains the following updates.

Changed information

The topic Chapter 10, “Prerequisites for using Dynamic Partition Manager,” on page 71 has been updated with the bundle that is required for using the DPM R3.1 storage management feature.

Summary of changes for SB10-7170-01b

This edition contains the following updates.

New information

Chapter 8, “Dynamically modifying the resources of an active partition,” on page 53 contains a new section, “Instructions: Detaching and reattaching an FCP storage group” on page 61.

Changed information

- “Migration instructions for a system replacement” on page 82 contains updated information in the About this task and What to do next sections.
- Chapter 12, “Upgrading from DPM R3.0 to R3.1,” on page 89 updated information in the About this task and What to do next sections.

Summary of changes for SB10-7170-01a

This edition contains the following updates.

New information

The topic Chapter 10, “Prerequisites for using Dynamic Partition Manager,” on page 71 has been updated with the bundle that is required for using the DPM R3.1 storage management feature.

Changed information

- “Migration instructions for a system replacement” on page 82 contains updated information in the Results and What to do next sections.
- Chapter 12, “Upgrading from DPM R3.0 to R3.1,” on page 89 clarifies that the upgrade process is concurrent, and contains updated information in the Results and What to do next sections.

Summary of changes for SB10-7170-01

This edition contains the following updates.

New information

- Chapter 11, “Migrating DPM configuration data to a new system,” on page 79 contains a new step-by-step procedure for importing an existing DPM configuration to a new system.
- Chapter 12, “Upgrading from DPM R3.0 to R3.1,” on page 89 contains a new step-by-step procedure for upgrading to the DPM R3.1 storage management feature.
- Chapter 17, “Managing storage on a DPM-enabled system,” on page 113 describes the DPM R3.1 storage management feature, which provides a simplified, visual, and automated user interface for configuring and managing storage resources, without requiring extensive knowledge of the mainframe. Through the Hardware Management Console (HMC) **Configure Storage** task:
 - Administrators can connect a single system to storage devices, rather than using either the Hardware Configuration Definition (HCD) or the Input/Output Configuration Program (IOCP) to define storage devices in an input/output configuration data set (IOCDs) for the system.
 - Administrators also use this task to request and fulfill storage resources for use by partitions.

DPM R3.1 requires the installation of an enablement MCL that your service representative can supply.

Changed information

- The following topics have been updated to include the **Configure Storage** task, which is available only on a system that has the DPM R3.1 storage management feature applied.
 - Chapter 10, “Prerequisites for using Dynamic Partition Manager,” on page 71 lists the MCLs that are required for the DPM R3.1 storage management feature.
 - Appendix A, “DPM task and resource roles,” on page 137 lists the authorizations that are required to use the **Configure Storage** task.
- All chapters in the following parts have been updated to note the differences between the DPM R3.1 storage management feature and storage management on earlier releases of DPM.
 - Part 2, “Basic concepts and terms for Linux administrators,” on page 9

- Part 3, “Basic tasks for Linux administrators,” on page 29
- The following topics have been updated to include new support for storage devices. With DPM R3.1 or later, you can access Fibre Connection (FICON) extended count key data (ECKD™) direct-access storage devices (DASD), and Fibre Channel Protocol (FCP) Small Computer System Interface (SCSI) storage devices. DPM R3.0 and earlier releases support only FCP devices.
 - Chapter 3, “Partitions: Virtual images of a mainframe system,” on page 11
 - Chapter 4, “Adapters: Connections to networks, storage, and more,” on page 17
- Chapter 14, “Customizing the network boot server environment,” on page 101 has been updated to note the support for Dynamic Host Configuration Protocol (DHCP) option 93: this **client-arch** option has a fixed value of 32. You can use this option to tailor the processing related to a request from an IBM Z system.

Summary of changes for SB10-7170-00b

This edition contains the following updates.

New information

Various topics throughout this book include the following new servers as systems on which DPM can be enabled:

- IBM z14 (z14) machine type 3907
- IBM LinuxONE Rockhopper II (Rockhopper II)

Deleted information

References to KVM for IBM z Systems® have been removed from the list of hypervisors that can run in a partition on a DPM-enabled system. IBM withdrew support for KVM for IBM z Systems, effective March 31, 2018.

Summary of changes for SB10-7170-00a

The topic Chapter 10, “Prerequisites for using Dynamic Partition Manager,” on page 71 has been updated with the engineering change (EC) machine change level (MCL) that is required for using the Integrated ASCII Console through a new DPM API: Get ASCII Console Websocket URI.

Part 1. Introduction to Dynamic Partition Manager

Topics in this part describe DPM, explain the benefits of using it, and provide a basic set of instructions for creating and starting a new partition on a DPM-enabled system. These topics are appropriate for readers who need a quick introduction to DPM.

Topics covered in this part are:

- Chapter 1, “Dynamic Partition Manager: A quicker and easier way to deploy Linux servers on a mainframe,” on page 3
- Chapter 2, “Getting started: Creating a new partition and starting your Linux server,” on page 5

Chapter 1. Dynamic Partition Manager: A quicker and easier way to deploy Linux servers on a mainframe

Linux servers and applications have run on mainframe systems for years, but configuration and setup are fairly complicated and requires the use of several specific tools. However, with IBM Dynamic Partition Manager (DPM), system administrators now have a quicker and easier way to deploy Linux servers, using only the Hardware Management Console (HMC).

DPM is a configuration manager that is designed for setting up and managing Linux servers that run on a mainframe system. On a DPM-enabled system, the runtime environment for your Linux server is called a *partition*. On other platforms, a partition is a portion of the system hard disk that you create to run different operating systems on the same disk, or to give the appearance of separate hard disks for multiple users or other purposes. On a mainframe system, a partition is a virtual representation of all of the physical hardware resources of that system, which include processors, memory, and input/output (I/O) adapters. On mainframe systems, as on other platforms, an *adapter* is a physical device that connects the system to other computers or devices.

To make use of DPM, system administrators select specific tasks in the HMC; the user interface of these tasks has a design that is similar to the tools that system administrators use on other platforms. These specific HMC tasks are available only on a DPM-enabled system. Figure 1 shows the new design of the HMC Welcome page, which provides a visual summary of the number and status of elements that are managed through the HMC: DPM-enabled systems, the partitions that run on those systems, and the adapters that are configured for those systems.




Figure 1. New user interface design for the HMC Welcome page

When your company orders an IBM system with the required DPM features, IBM service representatives install the system and enable DPM, so it is ready for use when the system is powered on. DPM is available on the following systems:

- IBM z14 (z14) (machine type 3906 or 3907)
- IBM z13 (z13) or IBM z13s (z13s)
- IBM LinuxONE Emperor II (Emperor II), or IBM LinuxONE Rockhopper II (Rockhopper II)
- IBM LinuxONE Emperor (Emperor) or IBM LinuxONE Rockhopper (Rockhopper) machine type 2965

When you use the HMC to configure the running environment for your Linux server, DPM automatically discovers and displays the system resources that are available for your use, and indicates how your selections might affect other servers and applications that are already defined or running on the same system. After your Linux server is up and running, you can use DPM to:

- Modify system resources without disrupting running workloads.
- Monitor sources of system failure incidents and conditions or events that might lead to workload degradation.
- Create alarms so that you can be notified of specific events, conditions, and changes to the state of system resources.
- Update individual partition resources to adjust capacity, redundancy, availability, or isolation.

Chapter 2. Getting started: Creating a new partition and starting your Linux server

This topic provides a quick review of the procedure for creating and starting a new partition to host a single image of the Linux operating system. Links to more detailed information are included in the procedure steps.

Before you begin

- Verify the intended use of DPM-enabled systems with your system planner, so you know which system is the appropriate one for you to use for your Linux server and the business applications that it will support. You also need to know which of the system features and resources (adapters, processors, memory, and so on) that your server and its business applications require.
- After you know which system to use, and complete capacity planning for the applications that you intend to run on the Linux server, you can list the system resources that you need to assign to your partition. Depending on the IT roles and operating procedures at your installation, you might have to work with the network administrator, storage administrator, or security administrator before you create the partition.

Procedure

1. Log in to the Hardware Management Console (HMC) with either the default SYSPROG user ID or a user ID that a system administrator authorized to this task through customization controls in the **User Management** task.
2. On the main HMC page, expand the **Systems Management** node to view managed systems, and select the DPM-enabled system on which you want to create a partition for your Linux server.




Figure 2. Portion of the main HMC page, with the Systems Management node expanded, and a system and the New Partition task highlighted

3. Use the **New Partition** wizard to create a partition. You can access this task from the main HMC page by selecting the Systems Management node, by selecting a specific DPM-enabled system, or by selecting the task in the Tasks index.

For more detailed guidance and instructions, see the following topics:

- Chapter 6, “Planning checklist for creating a partition,” on page 31

- “Creating a new partition” on page 37
- a. Open the **New Partition** task.

This action opens the New Partition window, which is shown in Figure 3.



Figure 3. Welcome page of the New Partition task in basic mode

The **New Partition** task offers two modes through which you can create a partition: basic and advanced. For a comparison of the two modes and the implications of switching between them, see “Selecting which New Partition task mode to use” on page 35.

- Complete the required fields on the task pages to create the new partition, depending on the requirements of the applications that your Linux server will host.
In Figure 3, pages that contain required fields are denoted with an asterisk. For step-by-step instructions for using the basic mode to complete the information on these pages, see “Creating a new partition” on page 37.

- Depending on the task mode you are using, click **Finish** or **OK** to create the partition. A progress indicator is displayed until DPM finishes creating the partition.

When it finishes creating the partition definition, DPM opens the validation window, a sample of which is shown in Figure 4 on page 7. The validation window displays a message indicating that your partition has been created, and lists more tasks that you can use to work with the new partition.




Figure 4. Sample validation window for the New Partition task

Note: The tasks that are listed in the validation window vary, depending on whether the DPM R3.1 storage management feature is applied to your system.

4. Use the **Start** task to start the partition. You can open the **Start** task by clicking the link on the Validation window, or through two other methods shown in Figure 5.




Figure 5. Methods of starting a new partition

- One way to access the Start task is to select the new partition and click the double-arrow icon in the Name field to display the cascading task menu. Then expand the **Daily** group, and click **Start**.
- Another way to access the Start task is to select the new partition, expand the Daily category in the Tasks area, and click **Start**.

A new window opens to display the progress of the start operation. Figure 6 shows a portion of the Start window, with its progress indicators. In addition to displaying a progress indicator, the window also contains a Details column with messages that are updated as the start process continues. These messages indicate the progress of configuring partition resources and initializing the operating system or hypervisor to run in the partition.




Figure 6. Portion of the Start window for a new partition named ASL_Test

When the start process completes, the Details column contains a success message, along with an **Open Console** link through which you can enter operating system commands.

- Click the **Open Console** link to open the **Integrated ASCII Console** task, through which you can log in to the operating system that is running on the partition.

After you log in, you can enter commands to view the processor, memory, and other resources that you configured for the partition. For example, Figure 7 shows a sample Linux **lscpu** command and the resulting display of configured processor resources.

```
Linux localhost 3.10.0-123.20.1.el7_0.zfpc5_0.33.s390x #1 SMP Mon Feb 8 12:13:30
EST 2016 s390x s390x s390x GNU/Linux
[root@localhost ~]# lscpu
Architecture:          s390x
CPU op-mode(s):        32-bit, 64-bit
Byte Order:            Big Endian
CPU(s):                1
On-line CPU(s) list:   0
Thread(s) per core:    1
Core(s) per socket:    8
Socket(s) per book:    3
Book(s):               8
Vendor ID:             IBM/S390
BogoMIPS:              20325.00
Hypervisor:            PR/SM
Hypervisor vendor:     IBM
Virtualization type:   full
Dispatching mode:      horizontal
L1d cache:             128K
L1i cache:             96K
L2d cache:             2048K
L2i cache:             2048K
[root@localhost ~]# lsmem
```

Figure 7. Sample Linux command and display of configured processor resources

Results

The partition and the Linux server are ready to support business applications.

Part 2. Basic concepts and terms for Linux administrators

Topics in this part explain mainframe and DPM concepts and terms that might be unfamiliar to system administrators who have little or no experience working in a mainframe environment.

Topics covered in this part are:

- Chapter 3, “Partitions: Virtual images of a mainframe system,” on page 11
- Chapter 4, “Adapters: Connections to networks, storage, and more,” on page 17
- Chapter 5, “Operating systems and hypervisors,” on page 23

Chapter 3. Partitions: Virtual images of a mainframe system

A partition is a virtual representation of the hardware resources of an IBM Z or LinuxONE system. A partition hosts an operating system and applications, which are sometimes called the *workload*.

The system planners at your company order and configure mainframe systems according to their plan for the business applications that each system will support. This plan determines the system on which you configure your Linux server and its workload, and determines which system resources are available when you configure a partition.

The following operating systems and hypervisors can run on a partition on a DPM-enabled system:

- Various Linux distributions, which are listed on the IBM tested platforms page for Linux environments. These distributions include supported versions of Red Hat Enterprise Linux (RHEL), SUSE Linux Enterprise Server (SLES), and Ubuntu Server (KVM or LPAR DPM).
- z/VM 6.4 or later. z/VM is an IBM virtualization hypervisor on which you can run multiple Linux images.

DPM also supports Secure Service Container, which is a container technology through which you can more quickly and securely deploy firmware and software appliances. Unlike most other types of partitions, a Secure Service Container partition contains its own embedded operating system, security mechanisms, and other features that are specifically designed for simplifying the installation of appliances, and for securely hosting them.

Figure 8 on page 12 illustrates the physical and virtual resources of a mainframe system, along with the firmware components that are used to manage these resources. Systems can be configured to run in either standard Processor Resource/Systems Manager™ (PR/SM) mode or IBM Dynamic Partition Manager (DPM) mode. DPM uses PR/SM functions but presents a simplified user interface for creating partitions and managing system resources through tasks in the Hardware Management Console (HMC) / Support Element (SE).




Figure 8. Partitions configured on a DPM-enabled system

In Figure 8, several partitions are configured on a DPM-enabled system. Each partition hosts either a hypervisor or an operating system, and has virtual system resources that represent its share of physical resources: processors, memory, and adapters.

- Partitions A through C each host one Linux operating system image.
- Partitions D and E each host one z/VM image and its multiple Linux guests.
- Partition F is a Secure Service Container partition that hosts a supported firmware or software appliance.

Partition properties and configuration settings

A *partition definition* contains the specific properties and configuration settings for one partition on a DPM-enabled system. You use the **New Partition** task to create a partition definition; through that task, you specify how many processors, how much memory, and which adapters to use.

When you use the **New Partition** task to create a partition definition, DPM indicates which system resources are available for your partition to use, and also shows the current usage or reservation of system resources by active (started) partitions or by partitions with reserved resources. You may define more resources than are currently available, and you can specify whether DPM is to reserve those resources for the partition. When you specify that the system resources for a partition are to be reserved, DPM does not allocate them to any other partitions. This reservation means that your partition is guaranteed to be startable; in contrast, partitions without reserved resources might fail to start, if sufficient resources are not available.

The following list describes key properties and configuration settings of partitions on a DPM-enabled system. The list labels correspond to navigation labels or individual fields in the **New Partition** task, and

the **Partition Details** task, through which you can modify an existing partition definition. For a complete list of the partition properties and settings, see the online help for either task.

Name A partition name must uniquely identify the partition from all other partitions defined on the same system. On a DPM-enabled system, you can define a name for your partition that is 1 - 64 characters in length. Supported characters are alphanumeric, blanks, periods, underscores, dashes, or at symbols (@). Names cannot start or end with blank characters. This partition name is shown in HMC task displays that contain information about system partitions.

A partition also has a short name, which is a name by which the operating system can identify the partition. By default, DPM automatically generates a partition short name that you can modify.

Partition type

Administrators can choose one of the following partition types for a new partition. Through the partition type, DPM can optimize the partition configuration for a specific hypervisor or operating system.

Linux In this type of partition, you can install and run a Linux on Z distribution.

z/VM In this type of partition, you can install and run z/VM as a hypervisor for multiple Linux guests.

Secure Service Container

This type of partition is a Secure Service Container, in which you can run only specific software appliances that the Secure Service Container supports.

Processors

Most DPM-enabled systems support one type of processor: Integrated Facility for Linux (IFL), which is designed specifically for the Linux operating system. In some cases, a system might also support an additional type: Central Processor (CP), which also supports the Linux operating system.

Each partition on a system can either have exclusive use of a specific number of physical processors installed on the system, or can share processor resources from the pool of physical processors that are not dedicated to other partitions on the same system. The number of available processors is limited to the number of entitled processors on the system. Entitled processors are processors that are licensed for use on the system; the number of entitled processors might be less than the total number of physical processors that are installed on the system.

When you create a new partition on a DPM-enabled system:

- You can select which processor type to use only if both types are installed on the system. Generally, IFLs are the most appropriate choice for Linux servers. If you want to enable simultaneous multithreading for this partition, you must select the IFL processor type.
- You can specify the number of processors to assign to the partition, and view how your selection affects the processing resources of other partitions on the system. The number of processors that you can assign ranges from a minimum value of 1 to a maximum value of the total number of entitled processors on the system.

Memory

Each partition on a DPM-enabled system has exclusive use of a user-defined portion of the total amount of entitled memory that is installed on the system. Entitled memory is the amount of memory that is licensed for use, which might be less than the total amount of memory that is installed on the system. The amount of memory that a specific partition requires depends on the storage limits of the operating system that will run on it, on the storage requirements of the applications that run on the operating system, and on the size of the I/O configuration.

When you define the amount of memory to be assigned, or allocated, to a specific partition, you specify an initial amount of memory, and a maximum amount that must be equal to or greater than the initial amount. The partition receives its initial amount when it is started. If the

maximum amount of memory is greater than the initial amount, you can add memory up to this maximum to the active partition, without stopping and restarting it.

Secure Service Container partitions require a minimum amount of 4096 MB (4 GB).

Network

Network interface cards (NICs) provide a partition with access to internal or external networks that are part of or connected to a system. Each NIC represents a unique connection between the partition and a specific network adapter that is defined or installed on the system.

You need to define a NIC for each network connection that is required for the operating system or hypervisor that runs on this partition, or for the applications that the operating system or hypervisor supports. DPM supports several types of network adapters, including Open Systems Adapter-Express (OSA-Express) features, IBM HyperSockets™, and Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) Express features.

Secure Service Container partitions require at least one NIC for communication with the Secure Service Container web interface.

Storage

The way in which partitions access storage depends on whether the DPM-enabled system has the DPM R3.1 storage management feature applied.

- With the feature applied, you select one or more storage groups to attach to a partition. A *storage group* is a logical group of storage volumes that share certain attributes. For example, one attribute is shareability: you can define a storage group as either dedicated for use by only one partition, or shared by multiple partitions. When administrators request and fulfill storage groups for use by partitions, DPM automatically generates the world wide port names (WWPNs) that are allocated to virtual storage resources when the storage group is attached to a partition.
- Without the feature applied, host bus adapters (HBAs) provide a partition with access to external storage area networks (SANs) and devices that are connected to a system. Each HBA represents a unique connection between the partition and a physical FICON channel that is configured on the system. When using HBAs, you need to define an HBA for each storage area network that is required for the applications that run on this partition.

DPM supports various FICON Express features, which are adapter cards that are installed on the system and are attached by cables to storage devices. The types of storage devices that you can use depend on the version and release of DPM that is enabled on your system. With DPM R3.1 or later, you can select both Fibre Connection (FICON) extended count key data (ECKD) direct-access storage devices (DASD), and Fibre Channel Protocol (FCP) Small Computer System Interface (SCSI) storage devices for your partition to use. With DPM R3.0 or earlier, you can select only FCP storage devices for use.

The storage requirements for a Secure Service Container partition depend on the type of appliance that it hosts.

Accelerators

Accelerators are adapters that provide specialized functions to improve performance or use of computer resources. One supported accelerator is the IBM zEnterprise® Data Compression (zEDC) feature, which provides hardware-based acceleration for data compression and decompression.

Accelerators are optional features and, therefore, might not be installed on the system. If one is installed, your decision to enable your partition to access it depends on the workload that your partition will support. Your system planner can advise you about the use of available accelerators.

Cryptos

The term *cryptos* is a commonly used abbreviation for adapters that provide cryptographic processing functions. DPM supports various Crypto Express features.

Crypto features are optional and, therefore, might not be installed on the system. If these features are installed, your decision to enable your partition to access them depends on your company's security policies, and the workload that your partition will support. Your system planner or security administrator can advise you about the use of available crypto features.

Boot options

When you define a partition with a type of **Linux** or **z/VM**, you can specify the boot option through which DPM locates and installs the executables for the hypervisor or operating system to be run on the partition. You can choose one of several different options, including booting from a storage device, network server, FTP server (with your choice of protocol), and Hardware Management Console removable media.

DPM automatically sets the boot option for the first-time start of Secure Service Container partitions.

Creating, starting, and managing a partition

To create a partition, you use the **New Partition** task, through which you define the hardware resources that the partition can use: processors, memory, adapters, and so on. The end result of the task is a partition definition, which you can modify through the **Partition Details** task, or use to start the partition through the **Start** task. When you start a partition, DPM uses the partition definition to determine which hardware resources to allocate to the partition, and starts the initialization process.

After the partition definition exists, you can use the **Partition Details** task to modify it; note that you cannot change the partition type after you create the partition definition. You can also use the **Stop** task to stop a partition, or the **Delete Partition** task to delete it. You can accomplish these tasks programmatically as well, through the Hardware Management Console Web Services application programming interfaces (APIs) for DPM.

You can create as many partition definitions as you want, but only a specific number of partitions can be active at any given time. The system limit determines the maximum number of concurrently active partitions; for example, 85 partitions can be active on a z13, while 40 can be active on a z13s. Practical limitations of memory size, I/O availability, and available processing power usually reduce the number of concurrently active partitions to less than the system maximum. In fact, conditions on the system might prevent a partition from successfully starting, or change its status after it has successfully started. You can view the status of a partition through the **Partition Details** task or use the **Monitor System Events** task to set notifications for specific partition events, such as a change in status.

For more details about working with partitions, see Part 3, “Basic tasks for Linux administrators,” on page 29.

Chapter 4. Adapters: Connections to networks, storage, and more

Adapters on a system fall into four categories: Network, Storage, Accelerators, and Cryptos. Each adapter type plays a specific role in communication, or data transfer, for partitions and the applications that run in them.

Most adapters are installed in the I/O cage or drawer of a physical processor frame. Depending on your company's planned use of specific systems, each system might have a different combination of installed adapters.

When adapters are installed in the processor frame, the adapters are configured using default settings. DPM automatically discovers these adapters and assigns names to them, using a standard naming convention. You can change the name and other default adapter settings through the **Manage Adapters** task, to conform with conventions that your company uses, or to provide more easily recognizable names for monitoring purposes.

To make use of the adapters configured on a DPM-enabled system, you select them when you use the **New Partition** task to create a new partition. Factors that determine your selections include:

- The specific adapters that are actually configured on the system.
- The requirements of the operating system and applications, which are sometimes called the *workload* that your new partition will support.
- Any requirements or restrictions that your company has for the use of specific adapters. For example, your company might recommend selecting several adapters of the same type to maximize efficiency and provide redundancy.

Each partition on the system can be configured to access any combination of the I/O adapters that are either installed or configured on the system. Partitions can share all types of supported adapters, up to specific limits. The following list provides a description of adapter types, by category.

Network

Several types of network adapters enable communication through different networking transport protocols. These network adapters are:

- Open Systems Adapter-Express (OSA-Express) features, which provide direct, industry-standard Ethernet LAN connectivity through various operational modes and protocols. OSA adapters can provide connectivity between partitions on the same system, as well as connectivity to external LANs.

The supported OSA adapters vary, depending on the system configuration; for example, the z14 supports the OSA-Express6S, OSA-Express5S, and OSA-Express4S adapters.

- HiperSockets, which provide high-speed communications between partitions within a single system, without the need for any physical cabling or external networking connections.
- Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) Express features. These features provide similar benefits for system-to-system communication as the HiperSockets technology provides for communication between partitions on the same system.

The supported RoCE adapters vary, depending on the system configuration; for example, the z14 supports the 10 Gigabit Ethernet (GbE) RoCE Express2 and 10 GbE RoCE Express adapters.

For a list of the network adapters that are supported on a specific system, see the appropriate system overview on the IBM Redbooks web site at <http://www.redbooks.ibm.com/>.

DPM automatically discovers OSA and RoCE adapters because they are physical cards that are installed on the system. In contrast, HiperSockets are not physical adapters; you must configure

them if you want to use them on your system. To create HiperSockets on a DPM-enabled system, use the **Create HiperSockets Adapter** task, which is available through the **Actions** list on the **Adapters** tab of the **Manage Adapters** task.

Network interface cards (NICs) provide a partition with access to internal or external networks that are part of or connected to a system. Each NIC represents a unique connection between the partition and a specific network adapter that is defined or installed on the system.

Most systems have OSA adapters installed, and you will probably define a NIC to connect your partition to at least one of those OSA network connections. Your system planner or network administrator can advise you on which network connections to use for the workload that your partition supports. Figure 9 provides a conceptual illustration of partitions that are configured to use different types of network adapters for access to internal and external networks.

- Partitions A and B are both connected to the HiperSockets network within the DPM-enabled system, and to an OSA card for Ethernet access to an external network.
- Partition C has only one NIC defined, for access to a RoCE adapter, which provides Ethernet access between the DPM-enabled system and one other system in the network.




Figure 9. Partitions with NICs configured for access to network adapters on a DPM-enabled system

Storage

Fibre Channel connections (FICON) provide high-speed data transfer between systems and storage devices. Fibre Channel networks consist of servers, storage controllers, and devices as end nodes, which are interconnected by Fibre Channel switches, directors, and hubs. Switches and directors are used to build Fibre Channel networks or fabrics. Through cables, FICON adapter cards connect the DPM-enabled system to the devices in this storage area network (SAN).

FICON adapter cards operate in different modes, which determine the type of storage devices that you can access. Typically, storage administrators configure the mode in which each FICON adapter card operates.

- Fibre Channel Protocol (FCP) mode provides access to Small Computer System Interface (SCSI) devices, through single- or multiple-channel switches. Note that, because DPM runs all FCP adapters in N Port Identifier Virtualization (NPIV) mode, the SAN must support NPIV. DPM senses but does not display FCP tape devices, so administrators cannot manage them through DPM tasks on the HMC; however, operating systems that are hosted by partitions on a DPM-enabled system can access those tape devices. Support for FCP mode is available with all DPM releases.

- FICON native (FC or FICON) mode provides access to extended count key data (ECKD) devices, and tape devices, through point-to-point (direct) connections, or single- or multiple-channel switches. ECKD devices are more commonly known as direct-access storage devices (DASD). DPM supports access to ECKD devices, but does not support FICON channel-to-channel (CTC) or FICON tape devices. Also, DPM supports the use of parallel access volumes, but only through the optional HyperPAV feature on the IBM System Storage® DS8000® series. Support for FICON mode is available only with DPM R3.1 or later.

Figure 10 shows a sample physical FICON network that can be configured for a system with DPM R3.1, which supports both FCP storage disks and FICON DASD.




Figure 10. Sample physical storage configuration with DPM R3.1

In Figure 10:

- Partitions can access FICON storage devices, FCP storage devices, or both: Partitions A and C can access both, while Partition B accesses only FCP storage devices.
- Adapter cards 01 and 42 provide access through a single switch (A) to FICON DASD on Control Unit 01.
- Adapter cards 22 and 31 provide access through cascaded switches (B and C) to FCP storage disks on Controllers 1 and 2.

DPM automatically discovers any FICON adapter cards that are configured on the system. These storage adapter cards are FICON Express features, which enable multiple concurrent I/O operations at various data transmission rates in gigabytes-per-second (Gbps), using Fibre Channel connections. The supported FICON Express adapter cards vary, depending on the system configuration; for example, the z14 supports the FICON Express16S+, FICON Express16S, and FICON Express8S adapter cards. For a list of the storage adapters that are supported on a specific system, see the appropriate system overview on the IBM Redbooks web site at <http://www.redbooks.ibm.com/>.

The way in which partitions access storage depends on whether the DPM-enabled system has the DPM R3.1 storage management feature applied.

- With the feature applied, you select one or more storage groups to attach to a partition. When administrators request and fulfill storage groups for use by partitions, DPM automatically

generates the world wide port names (WWPNs) that are allocated to virtual storage resources when the storage group is attached to a partition. During the attachment process, DPM generates the virtual storage resources (host bus adapters or FICON subchannels) that are required for partitions, and the operating systems that they host, to access the physical storage volumes. For more information, see “Using storage groups (DPM R3.1 or later)”

- Without the feature applied, you create host bus adapters (HBAs), which provide a partition with access to external storage area networks (SANs) and devices that are connected to a system. Each HBA represents a unique connection between the partition and a physical FICON channel that is configured on the system. For more information, see “Using HBAs (DPM R3.0 or earlier)” on page 21

Using storage groups (DPM R3.1 or later)

With DPM R3.1, system and storage administrators use the **Configure Storage** task to connect a system to devices in the SAN. They define the protocol mode of system adapters as either FCP or FICON, and they build a visual copy of the FICON connections of the SAN to which the system is attached. Through this process, DPM generates the virtual configuration that is required to connect the system to the physical SAN hardware.

After completing the initial configuration through the **Configure Storage** task, and after attaching physical cables to connect the system to storage devices, administrators use the **Configure Storage** task to request and fulfill storage groups for partitions to use.

A *storage group* is a logical group of storage volumes that share certain attributes. For example, one attribute is shareability: you can define a storage group as either dedicated for use by only one partition, or shared by multiple partitions. When an administrator submits a request for one or more storage groups, DPM automatically generates the world wide port names (WWPNs) that are allocated to virtual storage resources when the storage group is attached to a partition. The request goes to one or more storage administrators, who fulfill the request through tools for managing storage subsystems.

To access storage, you attach one or more storage groups to a partition, through either the **New Partition** task or **Partition Details** task. Figure 11 shows several storage groups with key attributes, and information about the physical devices that an administrator has configured to fulfill them. These physical devices are the same as the devices shown in Figure 10 on page 19.

DPM-enabled system




Figure 11. Sample storage groups that are attached to partitions on the system

In Figure 11 on page 20, the ovals in the partitions indicate which storage groups are attached to each partition.

- Partition A can access storage group SG01, which consists of a set of FICON base volumes that are dedicated for use by only one partition. Partition A accesses these volumes through FICON-mode adapter card 01 and switch A.
- Partitions A and B share storage group SG02, which consists of FCP disks that both partitions can access through FCP-mode adapter card 22 and switch B.
- Similarly, partitions B and C share storage group SG03, which consists of FCP disks that both partitions can access through FCP-mode adapter card 31 and cascaded switches B and C.

Using HBAs (DPM R3.0 or earlier)

When using HBAs, you need to define an HBA for each storage area network that is required for the applications that run on this partition.

Whether you need to define an HBA for your partition depends the following factors:

- The location of the code for the hypervisor or operating system that will run on the partition.
- The storage requirements of the workload that your partition will support. Your system planner or storage administrator can advise you on which storage connections to use for the workload that your partition supports.

Typically, you need to define an HBA for each storage area network that is required for the applications that run on this partition.

Figure 12 provides a conceptual illustration of partitions that are configured to use HBAs to access a storage area network.

DPM-enabled system




Figure 12. Using HBAs to access storage adapters and the storage network with DPM R3.0 or earlier

Accelerators

Accelerators are adapters that provide specialized functions to improve performance or use of computer resources. DPM automatically discovers accelerators that are installed on the system. One supported accelerator is the IBM zEnterprise Data Compression (zEDC) feature, which provides hardware-based acceleration for data compression and decompression.

zEDC features are installed in the Peripheral Component Interconnect Express (PCIe) I/O drawer. For each feature installed in the PCIe I/O drawer, one adapter/coprocessor compresses data according to the Internet Engineering Task Force (IETF) DEFLATE Compressed Data Format Specification, RFC 1951.

An accelerator virtual function provides a partition with access to zEDC features that are installed on a system. Each virtual function represents a unique connection between the partition and a physical feature card that is configured on the system.

Accelerators are optional features and, therefore, might not be installed on the system. If one is installed, your decision to enable your partition to access it depends on the workload that your partition will support. Your system planner can advise you about the use of available accelerators.

Cryptos

The term *cryptos* is a commonly used abbreviation for adapters that provide cryptographic processing functions. Industry Public Key Cryptography Standards (PKCS) and the Common Cryptographic Architecture (CCA) define various cryptographic functions, external interfaces, and a set of key cryptographic algorithms. These specifications provide a consistent, end-to-end cryptographic architecture across supported operating systems.

The use of the IBM cryptographic architecture is enabled through Crypto Express features, which provide a secure hardware and programming environment for cryptographic processes. Crypto Express features are installed in the Peripheral Component Interconnect Express (PCIe) I/O drawer.

The supported Crypto Express features vary, depending on the system configuration; for example, the z14 supports the Crypto Express6S and Crypto Express5S features. For a list of the cryptographic adapters that are supported on a specific system, see the appropriate system overview on the IBM Redbooks web site at <http://www.redbooks.ibm.com/>.

DPM automatically discovers cryptographic features that are installed on the system. Each Crypto Express adapter can be configured in one of the following modes.

- Secure CCA coprocessor (CEX4C) for Federal Information Processing Standard (FIPS) 140-2 Level 4 certification.
- IBM Enterprise PKCS#11 (EP11) coprocessor (CEX4P) for an industry-standardized set of services that adhere to the PKCS #11 specification v2.20 and more recent amendments.
- Accelerator (CEX5A) for acceleration of public key and private key cryptographic operations that are used with Secure Sockets Layer/Transport Layer Security (SSL/TLS) processing.

Additionally, you can enable or disable the key import functions that are available through the CP Assist for Cryptographic Functions (CPACF) feature. CPACF supports clear and protected key encryption based on the Advanced Encryption Standard (AES) algorithm and the Secure Hash Algorithm (SHA) with the Data Encryption Standard (DES) algorithm. For operating systems and applications to take advantage of key encryption support, the partition in which they run must be configured to permit AES and DES protected key import functions.

Crypto features are optional and, therefore, might not be installed on the system. If these features are installed, your decision to enable your partition to access them depends on your company's security policies, and the workload that your partition will support. Your system planner or security administrator can advise you about the use of available crypto features.

Linux servers require specific device drivers to use the adapters that are defined for a partition. For more information about the device drivers through which Linux servers can use adapters, see the *Linux on Z Device Drivers, Features, and Commands* documentation for the Linux kernel version that you are using. This documentation, which also describes commands and parameters for configuring Linux on Z, is available in IBM Knowledge Center at http://www.ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz_r_dd.html

Chapter 5. Operating systems and hypervisors

Partitions on a DPM-enabled system can host a single hypervisor or operating system. DPM supports the Linux operating system and several specialized hypervisors.

Partition types and supported operating systems and hypervisors

Administrators can choose one of the following partition types for a new partition. Through the partition type, DPM can optimize the partition configuration for a specific hypervisor or operating system.

Linux In this type of partition, you can install and run various Linux distributions, which are listed on the IBM tested platforms page for Linux environments. These distributions include supported versions of Red Hat Enterprise Linux (RHEL), SUSE Linux Enterprise Server (SLES), and Ubuntu Server (KVM or LPAR DPM).

z/VM In this type of partition, you can install and run z/VM 6.4 or later. z/VM is an IBM virtualization hypervisor on which you can run multiple Linux images.

Secure Service Container

This type of partition is a Secure Service Container, in which you can run only specific software appliances that the Secure Service Container supports.

Figure 13 shows a sample configuration of Linux servers and the types of workloads that they might support.




Figure 13. Example: Linux servers configured in partitions on a DPM-enabled system

- Partitions A, B, and C each support a single Linux server that provides specific functions or applications within the environment: firewall protection, domain name server (DNS), and data warehouse.
- Partition D is configured to run a z/VM hypervisor that hosts multiple Linux web servers.
- Partition E is configured to run a z/VM hypervisor that hosts multiple Linux images running various applications.
- Partition F is configured as a Secure Service Container that hosts a software appliance.

For details about the specific versions of operating systems and hypervisors that DPM supports, see “Requirements for the hypervisor or operating system” on page 76.

Installing an operating system or hypervisor in a partition

When you define a partition, you can specify the boot option through which DPM locates and installs the executables for the hypervisor or operating system to be run on the partition. You can select a boot option through the Boot section of the **New Partition** task or, for an existing partition, the **Partition Details** task. If you are using the advanced mode of the **New Partition** task, or using the **Partition Details** task, you also have the option of setting a time limit for the boot operation. By default, the time-out setting has a value of 60 seconds.

Select a boot option that is appropriate for the type of partition that you are defining.

- “Boot options for Linux partitions”
- “Boot options for z/VM partitions” on page 25
- “Boot option for Secure Service Container partitions” on page 27

Boot options for Linux partitions

For partitions with a type of **Linux**, you can choose one of several different options, including booting from a storage device, network server, and Hardware Management Console (HMC) removable media. For more detail about each option, see step 9 on page 44 in “Creating a new partition” on page 37, or use the online help for the **New Partition** task or the **Partition Details** task.

None Select this option if you want to start a partition without a hypervisor or operating system. Although the partition can be started, it is not in a usable state. This option is the default for partitions with a partition type of **Linux** and **z/VM**.

Storage device (SAN)

Select this option when the hypervisor or operating system executables reside on a device in a storage area network connected to the DPM-enabled system. This option is available only when storage groups or host bus adapters (HBAs) are defined for the partition.

When you specify this option, you need to either select a boot volume in an attached storage group or, to boot from a volume accessed through an HBA, provide the 64-bit worldwide port number (WWPN) of the boot image and the 64-bit hexadecimal logical unit number (LUN) that corresponds to the location of the boot image.

Network server (PXE)

Select this option when you want to use a preboot execution environment (PXE) on a network server. This option is available only if a network interface card (NIC) for either an OSA port or HiperSockets switch is defined for the partition.

FTP server

Select this option if you want to use FTP to boot an image that is located on a different system. To specify this option, you need to provide the host name, user name, and password associated with the FTP server, as well as an .INS file, which maps image components (for example, kernel, ramdisk, parameter file) to the appropriate storage addresses in main memory.

FTPS server

Select this option if you want to use the FTP Secure (FTPS) protocol to boot an image that is located on a different system. FTPS uses the Secure Socket Layer (SSL) protocol to secure data. With this option, you need to supply a host name, user ID, password, and .INS file, as described for the **FTP server** boot option.

SFTP server

Select this option if you want to use the Secure File Transfer Protocol (SFTP) to boot an image that is located on a different system. SFTP uses the Secure Shell (SSH) protocol to secure data. With this option, you need to supply a host name, user ID, password, and .INS file, as described for the **FTP server** boot option.

Hardware Management Console removable media

Select this option if you want to use an INS file from a media drive that is connected to the HMC. The media drive must be available when you are creating the partition definition and when the partition is started.

ISO image

Select this option when you want to upload an ISO file that is located on your workstation file system. This option is available only when you are connecting to the HMC through a remote browser. An ISO file is a collection of files and metadata for installing software.

Boot options for z/VM partitions

For partitions with a type of **z/VM**, you can choose one of several boot options that are described in step 4 of the following procedure, which provides an overview of the DPM-specific portion of a traditional, first-level z/VM installation process. For complete instructions, use the *Installation Guide* for the version of z/VM that you are installing. The z/VM library is available at <http://www.vm.ibm.com/library/>

1. Prepare to access the installation files.

- If you are using the **Hardware Management Console removable media** boot option, load the z/VM product DVD in the HMC DVD drive. Note that, when you select this option, installation can take up to several hours to complete.
- If you are using any of the FTP boot options, upload the contents of each DVD or electronic deliverable to a new directory on the FTP server. After the contents of the z/VM product deliverable have been uploaded, upload the contents of the installation recommended service upgrade (RSU) deliverable to the same directory, overwriting any duplicate files. (For details about directory path names and required space, see the *z/VM Installation Guide*.)
- If you are using the **ISO image** boot option, complete the following steps.
 - a. Upload the contents of each DVD or electronic deliverable to a new directory on your hard drive.
 - b. After the contents of the z/VM product deliverable have been uploaded, upload the contents of the installation recommended service upgrade (RSU) deliverable to the same directory, overwriting any duplicate files. (For details about directory path names and required space, see the *z/VM Installation Guide*.)
 - c. Using the tool of your choice, create an ISO image file of the directory.

2. On the main HMC page, expand the **Systems Management** node to view managed systems, and select the DPM-enabled system on which the partition resides. On the **Partitions** tab for the selected system, find the table entry for the z/VM partition, and select it.
3. Open an integrated 3270 console for the partition in which you are installing z/VM. You can open the **Integrated 3270 Console** task through the cascading task menu next to the partition name, or through the Recovery category in the task area.
4. Open the **Partition Details** task and navigate to the **Boot** section.
 - a. Specify one of the following boot options, and select the INS file for the version of z/VM that you are installing; for example, 710VM.ins

Hardware Management Console removable media

Select this option if you want to use an INS file from a media drive that is connected to the HMC. The media drive must be available when you are creating the partition definition and when the partition is started.

When you select this option:

- 1) Identify the type of media drive on which the INS file resides by clicking the appropriate radio button.
- 2) Either enter the fully qualified name (relative to the mount point) of an INS file, or click **Browse** to start a search on the target media drive to retrieve a list of INS files. Any INS files found are displayed in a separate window. Select only one INS file and click **OK** to close the Browse Removable Media window.

FTP server

Select this option if you want to use FTP to boot an image that is located on a different system. Provide the following information:

Host name

Enter either the fully qualified domain name of the FTP server, or its IP address.

User name

Enter the user name on the target FTP server.

Password

Enter the password associated with the user name on the target FTP server.

INS file

Either click **Browse** to retrieve a list of INS files from the target FTP server and select one file, or enter the fully qualified name (relative to FTP root) of an INS file.

Depending on the size of the FTP site, browsing might require more time than manually entering the full path and name of the INS file. Also note that the browsing function returns INS files found in the user's home directory or its subdirectories. Because you cannot select a starting directory, or navigate to a directory above the user's home directory, manually entering the full path and name of the INS file might be more expedient.

If you click **Browse**, a separate window displays the user's home directory and its subdirectories. Select one INS file, and click **OK** to close the Browse FTP Server window.

FTPS server

Select this option if you want to use the FTP Secure (FTPS) protocol to boot an image that is located on a different system. FTPS uses the Secure Socket Layer (SSL) protocol to secure data. With this option, you need to supply a host name, user ID, password, and .INS file, as described for the **FTP server** boot option.

SFTP server

Select this option if you want to use the Secure File Transfer Protocol (SFTP) to boot an image that is located on a different system. SFTP uses the Secure Shell (SSH) protocol to secure data. With this option, you need to supply a host name, user ID, password, and .INS file, as described for the **FTP server** boot option.

ISO image

Select this option when you want to upload an ISO file that is located on your workstation file system. This option is available only when you are connecting to the HMC through a remote browser.

When you select this option:

- 1) Click **Browse** to find the ISO image file on your workstation file system. You cannot select an ISO image from an HMC media drive. As soon as you select an ISO image file, DPM starts to upload the file, and displays a progress indicator for the upload operation.
- 2) After the upload operation completes, click **Browse** to search the ISO image file for the INS file that you want to use. Any INS files found are displayed in a separate window. Select only one INS file and click **OK** to close the Browse ISO Image window.
- b. Select **OK** to save the boot option in the partition definition, and close the Partition Details window.
5. Using one of the following methods, open the **Start** task to start the partition.
 - a. From the entry for the selected partition on the **Partitions** tab, select the double-arrow icon in the Name field to display the cascading task menu, and select **Start**.
 - b. Expand the **Daily** category in the task area, and select **Start**.

A new window opens to display the progress of the start operation. In addition to displaying a progress indicator, this window also contains a Details column with messages that are updated as the start process continues. These messages indicate the progress of configuring partition resources and initializing the z/VM RAMDISK. If the start process is successful, the z/VM system loads with the MAINT user ID logged on.
6. Through the Integrated 3270 Console window, run the **DVDPRIME** command with the *dasdtype* and *source* that you are using to install.

dasdtype Specify either **3390** or **FBA**.

source Specify the *source* value that corresponds to the boot option that you specified.

 - For the **Hardware Management Console removable media** boot option, specify **dvd** as the source.
 - For the FTP boot options, specify **server** as the source.
 - For the **ISO image** boot option, specify **dvd** as the source.
7. Complete the installation by following the instructions in the z/VM *Installation Guide* for a non-Single System Image (SSI) traditional installation.

Boot option for Secure Service Container partitions

DPM automatically sets the boot option for Secure Service Container partitions to **Secure Service Container**. This boot option cannot be changed unless you first change the partition type.

For more information about creating and using Secure Service Container partitions, see *Secure Service Container User's Guide*, SC28-6978, which is available on Resource Link: <http://www.ibm.com/servers/resourcelink>

Working with an installed operating system or hypervisor

When you start a partition, DPM configures partition resources and initializes the operating system or hypervisor, according to the information you supply through the **New Partition** or **Partition Details** task. After the start process completes and the partition becomes active, you can open one of the following tasks to log in to the operating system or hypervisor.

- To log in to a Linux system, use the **Operating System Messages** task or the **Integrated ASCII Console** task. The **Integrated ASCII Console** task must be enabled through the operating system before you can use it.
- To log in to a z/VM hypervisor that is hosting multiple Linux systems, use the **Integrated 3270 Console** task.

For additional information about a specific operating system or hypervisor, see the appropriate resource:

- For Linux distributions, see the Linux on IBM Systems topics in IBM Knowledge Center, at: https://www.ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz_r_lib.html
- For the z/VM hypervisor, see the documentation for the version of z/VM that you are installing. The z/VM library is available at <http://www.vm.ibm.com/library/>

Note: DPM supports z/VM with the limitations listed in “Requirements for the hypervisor or operating system” on page 76, so not all z/VM functions described in the z/VM library are available for use.

- For more information about creating and using Secure Service Container partitions, see *Secure Service Container User's Guide*, SC28-6978, which is available on Resource Link: <http://www.ibm.com/servers/resourcelink>

Part 3. Basic tasks for Linux administrators

Topics in this part provide a planning checklist and step-by-step instructions for creating and starting a partition, and its operating system or hypervisor, on a DPM-enabled system. Also included is a summary of HMC tasks for monitoring and managing partitions, adapters, and other resources on a DPM system. These topics are appropriate for any administrator who creates or manages DPM partitions.

Topics covered in this part are:

- Chapter 6, “Planning checklist for creating a partition,” on page 31
- Chapter 7, “Creating and starting a new partition,” on page 35
- Chapter 8, “Dynamically modifying the resources of an active partition,” on page 53
- Chapter 9, “Summary of tasks for managing systems, adapters, and partitions,” on page 65

On a system that has the DPM R3.1 storage management feature applied, system administrators can also use the **Configure Storage** task to configure, view, and manage storage resources for the partitions that they create and manage. For more information, see Chapter 17, “Managing storage on a DPM-enabled system,” on page 113.

Chapter 6. Planning checklist for creating a partition

Before you can use the **New Partition** task to create a new partition, you need to know some specifics about your company's IT environment and about the type of applications that your new partition will support. Use the checklist in this topic to help you prepare to fill in information on the pages or in the sections of the **New Partition** task, which you can use in either basic or advanced mode. Although this checklist does not cover every field or selection that you might make, it covers key decisions and functions that enable you to determine which **New Partition** task mode to use, and helps you identify other IT personnel who might need to supply you with information.

Before you can use the **New Partition** task to create a new partition, you need to know the following information. You might need to consult with your company's system planner, network administrator, storage administrator, or security administrator for some of this information.

- The resource requirements of the operating system and applications, which are sometimes called the *workload*, that your new partition will support. Specific resources include processor capacity, memory, network or storage connections, and security requirements.
- The requirements or restrictions, if any, that your company has for the use of the adapters that are installed and configured on the system.
- The naming and numbering conventions, if any, that your company uses for its hardware resources, such as partitions and storage devices.

Note: When you create a network or storage connection for the partition, you also select an adapter from a list that DPM displays in the **New Partition** task. Although these adapters are configured on the system, they might not be attached by cables to external systems or devices. Make sure that you consult with the network administrator or storage administrator to determine which adapters are cabled (not all adapters require cables to external devices).

Table 1 lists the key decisions and information that you might need to supply while using the **New Partition** task.

- If any terms in the checklist are unfamiliar, see the topics in Part 2, “Basic concepts and terms for Linux administrators,” on page 9, or use the online help for the **New Partition** task.
- After you have completed the checklist, see the information in “Creating a new partition” on page 37.

Table 1. Planning checklist for defining your new partition

✓	Planning item or decision	Your notes:
	Partition characteristics: Choose the system on which your new partition will reside, and choose a unique name for your partition. You can also decide whether you want to provide a short name for the operating system to use, provide a unique partition ID that is assigned to your partition each time it is started, select a type, and reserve partition resources.	System name: Partition name: Partition ID: Partition type: Linux z/VM Secure Service Container Reserve resources: Yes No
	Status: Define the acceptable availability status values for the partition, based on the importance of its workload. Setting status values is useful if you want to know when various conditions have affected the operation of your partition. The default value is Active, which means that any other condition is flagged with an exception icon.	For a list and explanation of status values for partitions, see the online help for the Status section of the New Partition task. This section is available only in advanced task mode.

Table 1. Planning checklist for defining your new partition (continued)

✓	Planning item or decision	Your notes:
	Controls: Use the Controls section to enable or disable partition access to various controls. By default, all settings are unchecked. The major control categories are: • Partition access • Counter facility authorization • Sampling facility authorization	For a complete list and explanation of available controls, see the online help for the Controls section of the New Partition task. This section is available only in advanced task mode.
	Processors: Based on your knowledge of the workload this partition will support, determine the amount of processor resource it requires. If you have run this workload in another environment, you might already know its capacity requirements; if not, you can use the default value or specify a value, and make adjustments through the Partition Details task after starting the partition and observing the workload performance. If you want to enable simultaneous multithreading for this partition, you must select the IFL processor type.	Amount: _____ Type: IFL CP Shared: Yes No
	Memory: Based on your knowledge of the workload this partition will support, determine the amount of memory it requires. Specify a maximum amount if you want the partition to have access to additional memory resources without having to stop and restart it. You can specify initial values and use the Partition Details task to make adjustments, if necessary. If you are creating a Secure Service Container partition, you must specify a minimum amount of 4096 MB (4 GB).	Amount: _____ MB GB TB Maximum amount: _____ MB GB TB
	Network connections: Create a network interface card (NIC) for each network to which the partition needs access. For each NIC, provide a name and optional device number, and select one network adapter from the list of available adapters. For availability, select at least two network adapters of the same type, and create a NIC for each one. For partitions with a type of Linux or z/VM only, you also can optionally specify a virtual LAN (VLAN) identifier only if you plan to select an OSA-Express or HiperSockets adapter. For any type of partition, you can optionally specify a media access control (MAC) address, also only if you plan to select an OSA-Express or HiperSockets adapter. If you are creating a Secure Service Container partition, you must specify at least one NIC for communication with the Secure Service Container web interface. Although you can specify a VLAN ID, note that DPM does not provide VLAN enforcement for Secure Service Container partitions.	NIC name: NIC device ID: VLAN ID: MAC address: Adapter type: HiperSocket OSA RoCE NIC name: NIC device ID: VLAN ID: MAC address: Adapter type: HiperSocket OSA RoCE NIC name: NIC device ID: VLAN ID: MAC address: Adapter type: HiperSocket OSA RoCE

Table 1. Planning checklist for defining your new partition (continued)

✓	Planning item or decision	Your notes:
	Storage: The way in which partitions access storage, and the types of storage they can access, depend on whether the system has the DPM R3.1 storage management feature applied. With the feature applied, you attach one or more storage groups to a partition to access FCP disks or FICON DASD. Without the feature applied, you create host bus adapters (HBAs) to access FCP disks only. Attaching FCP or FICON storage groups (DPM R3.1 or later) Select one or more storage groups to attach to the partition. Use the attributes, capacity, and other information to determine which storage groups to select. You can select FCP storage groups, FICON storage groups, or both types of storage group. Creating HBAs to access FCP disks (DPM R3.0 or earlier) For each HBA, provide a name and optional device number, and select one storage adapter from the list of available adapters. For availability, select at least two storage adapters of the same type, and create an HBA for each one. Note: With DPM R3.0 or earlier, after you create HBAs and save the partition definition, you need to use the Export WWPNS task before you start the partition. Your storage administrator needs this list of WWPNS to perform zoning and LUN masking so your partition can successfully access the storage devices. You do not need to use the Export WWPNS task if you are working on a system that has the DPM R3.1 storage management feature applied, because WWPNS are assigned when FCP storage groups are configured.	HBA name: HBA device ID: HBA name: HBA device ID: HBA name: HBA device ID: Storage group name: Type: FCP FICON Storage group name: Type: FCP FICON Storage group name: Type: FCP FICON
	Operating system / Hypervisor: For a list of supported operating systems and hypervisors, see Chapter 5, “Operating systems and hypervisors,” on page 23. When you know which operating system or hypervisor that you want to run in your partition, determine the location where the executables reside, and the name of the ISO image file or .INS file. The location determines which boot option you select. Only one boot option, ISO image, requires the ISO file name.	Operating system Hypervisor: Location: SAN Net boot FTP HMC local drive ISO file name: .INS file name:
	Boot option: Select a boot option through which DPM can locate and upload the required files to initialize the hypervisor or operating system when the partition itself is started. With all options except None, you can specify a time limit (in seconds) for the load operation. With the Storage device (SAN) option, you also can specify additional information, such as OS load parameters. DPM automatically sets the boot option for the first-time start of a Secure Service Container partition. For a list and descriptions of the boot options, see Chapter 5, “Operating systems and hypervisors,” on page 23 or the online help for the New Partition task.	Boot loader time-out value: _____ seconds Storage device (SAN): WWPN: LUN: <i>or</i> Storage group boot volume: FTP server: Host name: User name / password: Protocol: FTP FTPS SFTP HMC removable media: Type of media drive: CD/DVD USB

Chapter 7. Creating and starting a new partition

To create a partition on a IBM Dynamic Partition Manager (DPM)-enabled system, use the **New Partition** task. For more information, see the following topics:

- “Selecting which New Partition task mode to use”
- “Creating a new partition” on page 37
- “Starting a partition and its operating system or hypervisor” on page 48

Selecting which New Partition task mode to use

The **New Partition** task offers two modes through which you can create a partition: basic and advanced. Basic is the default mode, but you have the option of setting advanced as the default mode.

Basic The basic task, which is presented the first time that you open the **New Partition** task, provides a quick, guided method of creating a partition; DPM either provides default values or automatically generates many of the values for partition properties that are required to successfully start a partition. Some of these properties are not displayed or editable in the basic task mode. To navigate through the task, use the **Next** and **Back** buttons. When you have finished entering values in the required fields, click **Finish** to create the partition definition.

Advanced

The advanced task, which you can launch from the basic task, enables experienced users to view all partition properties and to change any default values. To access each section in the advanced task, click the appropriate link in the navigation pane, or scroll down the main page and expand or collapse each section as necessary. When you have finished entering values in the required fields, click **OK** to create the partition definition.

To use the **New Partition** task in either mode, you need to use either the default SYSPROG user ID or a user ID that a system administrator authorized to this task through customization controls in the **User Management** task.

Comparing the task modes

Table 2 lists key partition properties, and indicates whether you can edit those properties using the **New Partition** task in either basic or advanced mode.

- A dash (—) indicates a property that you cannot edit in the basic task mode. DPM either provides default values or automatically generates values for these properties.
- A check mark (✓) indicates a property that you can edit.

Table 2. Comparison of editable partition properties in the basic and advanced **New Partition** task modes

Partition property	Basic mode	Advanced mode
Partition name	✓	✓
Partition short name and ID	—	✓
Partition type	✓	✓
Reserved resources	—	✓
Acceptable partition status values	—	✓

Table 2. Comparison of editable partition properties in the basic and advanced **New Partition** task modes (continued)

Partition property	Basic mode	Advanced mode
Controls: • Partition access • Counter facility authorization • Sampling facility authorization	—	✓ (editing requires SYSPROG or SERVICE user ID)
Shared processors	✓	✓
Dedicated processors	—	✓
Processing weights and capping	—	✓
Memory (initial allocation)	✓	✓
Maximum memory (dynamic allocation)	✓	✓
Network interface cards (NICs)	✓	✓
VLAN ID and MAC address for NICs	—	✓
Storage (storage groups or HBAs)	✓	✓ (ability to edit device numbers and to change adapters for FCP storage groups in this mode only)
Accelerators (virtual functions)	✓ (if installed on system)	✓ (if installed on system)
Cryptos (security)	✓ (if installed on system)	✓ (if installed on system) Permitting AES or DEA protected key import is available in this mode only.
Boot options	✓	✓

Switching between task modes

You have the option of switching between the basic and advanced task modes, and the option of setting the advanced mode as the default mode whenever you subsequently launch the **New Partition** task. To switch from the basic mode to the advanced mode, click **Advanced**, which is located in the lower left corner of the New Partition window. Clicking **Advanced** opens a confirmation dialog through which you can set the advanced mode as the default mode whenever you launch the **New Partition** task.

If you start in basic mode and switch to advanced mode

- If you edited any fields in the basic mode and then switch to the advanced mode, your changes are automatically carried over into the advanced mode. For example, if you entered a name for your new partition on the **Name** page of the basic task, that name is displayed on the **General** page of the advanced task.
- To switch back to the basic task mode, click **Basic**, which is located in the lower left corner of the New Partition window.
 - Clicking **Basic** opens a confirmation dialog through which you can set the basic mode as the default mode whenever you launch the **New Partition** task.
 - If you edited any fields in the advanced mode, those changes are not preserved when you switch back to the basic mode. However, any edits that you originally made in the basic mode are preserved.

If you start in advanced mode and switch to basic mode

If you edited any fields in the advanced mode and then switch to the basic mode, your changes are discarded, even if the partition property is available for editing in the basic mode.

Creating a new partition

This procedure provides step-by-step instructions for using the **New Partition** task to create a new partition.

Before you begin

The **New Partition** task offers two modes through which you can create a partition: basic and advanced. This procedure provides instructions only for the basic mode of the **New Partition** task. Because the advanced mode is similar, however, you can use these instructions for the advanced mode as well. Note that some pages, or sections, of the advanced mode might have slightly different names and additional content, compared to the basic mode. If you want to review the differences between the two modes, see “Selecting which New Partition task mode to use” on page 35.

- Make sure you have the appropriate authorization to use the **New Partition** task. You need to use either the default SYSPROG user ID or a user ID that a system administrator authorized to this task through customization controls in the **User Management** task.
- Use the online help for the **New Partition** task together with these instructions; the online help explains the page elements and functions in more detail. To access the online help, click **Help** on the **New Partition** task window. Note that the basic and advanced modes of the task have separate online help; to access the help for the advanced mode, switch to that mode and then click **Help**.
- If you are creating a partition only to familiarize yourself with the process, you can accept default or automatically generated property values or settings. After you successfully complete the **New Partition** task, you can use the **Partition Details** task to modify the partition definition to conform to your company's conventions and planned use for this system. Note that you cannot change the partition type through the **Partition Details** task.
- If you want to supply your own property values or configuration settings rather than accepting default values, use the checklist in Chapter 6, “Planning checklist for creating a partition,” on page 31 to gather the information that you need to select or fill in values in the **New Partition** task.

About this task

The basic mode of the **New Partition** task provides a quick, guided method of creating a partition; DPM either provides default values or automatically generates many of the values for partition properties that are required to successfully start a partition. Some of these properties are not displayed or editable in the basic task mode.

Some of the following individual steps are marked as required, which indicates that the corresponding task page contains fields for which you need to supply a value or make a selection. The end result of the task is a partition definition, which you can modify through the **Partition Details** task, or use to start the partition through the **Start** task.

Although the steps in this procedure include information that is specific to creating a Secure Service Container partition, you can find more detailed instructions either in the online help for the **New Partition** task or in the *Secure Service Container User's Guide*, SC28-6978, which is available on <http://www.ibm.com/servers/resource link>.

Procedure

1. Open the **New Partition** task. You can access this task from the main HMC page by selecting the Systems Management node, by selecting a specific DPM-enabled system, or by selecting the task in the Tasks index. For example:

- a. Select a DPM-enabled system listed under the Systems Management node.
- b. From the Configuration task group, click the link for the **New Partition** task. The New Partition window opens, with an overlay that highlights key task controls on the window.
- c. Click the **Okay, got it** button to remove the overlay. The Welcome page is displayed.




Figure 14. Welcome page of the New Partition task in basic mode

- d. On the Welcome page, you can use two controls to modify the page display.
 - Click the **Show this welcome page next time** check box if you want to see the Welcome page the next time that you open this task. By default, the check box is not selected.
 - Click the icon at the end of the check box label if you want to restore the page overlay.
 - e. Click **Next** to navigate to the next page in the task.
2. Required: Use the Name page to enter the name of the new partition, the partition type, and an optional description. A partition name must uniquely identify the partition from all other partitions defined on the same system.
 - a. Specify the name of the new partition, which can be 1 - 64 characters in length. Supported characters are alphanumeric, blanks, periods, underscores, dashes, or at symbols (@). Names cannot start or end with blank characters.
 - b. Optionally, specify a description for the partition. The description can be up to 1024 characters in length.
 - c. Specify one of the following values that identifies the type of partition that you are creating.

Linux In this type of partition, you can install and run a Linux on Z distribution.

z/VM In this type of partition, you can install and run z/VM as a hypervisor for multiple Linux guests.

Secure Service Container
This type of partition is a Secure Service Container, in which you can run only specific software appliances that the Secure Service Container supports.

If you select the partition type **Secure Service Container**, the page display includes fields for entering a master user ID and password to secure access to the Secure Service Container web interface. If you need help supplying values for these fields, see the online help.

- d. When you have finished, click **Next** to navigate to the next page in the task.
3. Required: Use the Processors page to define the number of shared virtual processors for the partition, and to view various charts that are based on your selections. The virtual processors are allocated from physical processors of the selected type.




Figure 15. Processors page of the New Partition task in basic mode

- a. If the Processors type field is displayed, select a value. If you want to enable simultaneous multithreading for this partition, you must select the IFL processor type.
- b. Review the Processors bar chart to determine how many processors are available on this system, and how many are already in use or reserved for other partitions.
- c. Select the number of processors that you want to assign to your new partition. If you are creating a partition only to familiarize yourself with the process, you can accept the default value. Otherwise, base your selection on your knowledge of the processing requirements of the operating system and applications that you plan to run on this new partition.
- d. Review the Processors bar chart and pie chart to understand how your selection affects the availability of processing resources on the system. Although you can select a number of processors greater than the number that is currently available, your new partition will not start unless currently active, unreserved partitions are stopped or more processors are added to the system.
- e. When you have finished, click **Next** to navigate to the next page in the task.
4. Required: Use the Memory page to define the initial and maximum amounts of memory to be assigned to the new partition.

When you define the amount of memory to be assigned, or allocated, to a specific partition, you specify an initial amount of memory, and a maximum amount that must be equal to or greater than the initial amount. If you are creating a partition only to familiarize yourself with the process, you can accept the default values for both the Memory and Maximum Memory fields. Otherwise, base your selection on your knowledge of the memory requirements of the operating system and applications that you plan to run on this new partition.




Figure 16. Memory page of the New Partition task in basic mode

- a. Review the Installed Memory bar chart to determine how much memory is available on this system, and how much is already in use or reserved for other partitions.
 - b. Select the amounts of initial and maximum memory that you want to assign to your new partition. If you are creating a Secure Service Container partition, you must specify a minimum amount of 4096 MB (4 GB).
 - c. Review the Installed Memory bar chart again, to understand how your selection affects the availability of memory resources on the system.
 - d. When you have finished, click **Next** to navigate to the next page in the task.
5. Use the Network page to define network interface cards (NICs) to enable the new partition to access specific networks.

If you are creating a partition only to familiarize yourself with the process, you do not need to create any NICs unless you want to do something more than simply start the partition when you have finished creating it. If you are only trying out the process, skip to step 6 on page 42 of these instructions. Otherwise, you need to create a NIC for each network connection that is required for the operating system or hypervisor that runs on this partition, or for the applications that the operating system or hypervisor supports.

- For availability, select at least two network adapters of the same type, and create a NIC for each one.
- If you are creating a Secure Service Container partition, you must specify at least one NIC for communication with the Secure Service Container web interface. Note that, if you select an OSA-Express adapter for this NIC, you must select one that is using port 0.




Figure 17. Network page of the New Partition task in basic mode

When you first use the **New Partition** task, the Network display contains an empty NICs table. To create a NIC, complete the following steps.

- a. From the Actions list in the NICs table, select **New** to open the New Network Interface Card window.
- b. Enter a unique, meaningful name and, optionally, a description of the new NIC.
- c. If you are creating a Secure Service Container partition, supply values for the additional information about the network connection that is required to access the Secure Service Container web interface. This information includes an optional, virtual local area network (VLAN) identifier, the required IP address and type, and a mask / prefix.

If you need more detailed descriptions as you provide these configuration values, see the online help.

- d. Review the entries in the Adapter Ports and Switches table to determine which network adapters are configured on the system.
 - 1) Check the percentages listed in the Uplink Utilization and Adapter NIC Allocation columns. If the percentage in either column is high (for example, 90%) for a specific port or switch, consider selecting a different port or switch on the same network.
 - 2) Look for a warning icon next to the name in the Adapter Name column; if the warning icon is displayed for a specific port or switch, select a different one on the same network.
 - 3) Select one port or switch by clicking the radio button in the Select column. Note that, if you select an OSA-Express adapter port other than port 0, you need to manually specify the relative port number through a Linux qeth device driver command, before entering the Linux command to bring the device online. Instructions for specifying a port number are provided at the end of the procedure in “Starting a partition and its operating system or hypervisor” on page 48.
- e. Click **OK** to create the new NIC and close the New Network Interface Card window.
- f. Check the entry for the new NIC that is displayed in the NICs table on the Network page. Change the device number if your company uses a specific numbering convention for its networks.
- g. If you are creating a Secure Service Container partition, provide the network settings that are displayed after the NICs table. These settings include a host name, default gateway, and DNS servers.

If you need more detailed descriptions as you provide these configuration values, see the online help.

- h. Repeat the preceding steps, as necessary, to create a new NIC for each network connection that your new partition requires.
 - i. When you have finished, click **Next** to navigate to the next page in the task.
6. Use the Storage page to attach storage groups or create host bus adapters (HBAs) that enable the partition to access storage networks and hardware that is connected to the DPM-enabled system.

If you are creating a partition only to familiarize yourself with the process, you do not need to attach any storage groups or create any HBAs unless you want to boot a specific hypervisor or operating system image that resides on a storage device. If you are only trying out the process, skip to step 7 on page 43 of these instructions. Otherwise, continue with this step.

The Storage section contains either a Storage Groups table or an HBAs table with controls that you can use to attach storage groups or create HBAs. The Storage Groups table is displayed only when the DPM R3.1 storage management feature is applied on the system. Follow the instructions that correspond to the type of table displayed on the page.

- “Attach FCP or FICON storage groups” (DPM R3.1 or later)
- “Access FCP disks through HBAs” (DPM R3.0 or earlier)

Attach FCP or FICON storage groups

System administrators create storage groups to enable partitions (and the operating systems and applications that they host) to use physical storage hardware that is connected to the system. A *storage group* is a logical group of storage volumes that share certain attributes. The volumes reside on either Fibre Connection (FICON) extended count key data (ECKD) direct-access storage devices (DASD), or Fibre Channel Protocol (FCP) Small Computer System Interface (SCSI) storage devices. Storage groups can be shared by multiple partitions, and multiple storage groups can be attached to one partition.

To attach one or more storage groups to the partition, complete the following steps.

- a. When you first use the **New Partition** task, the Storage display contains an empty Storage Groups table. Select the plus icon in the table toolbar to open the Attach Storage Groups window.
- b. On the Attach Storage Groups window, select one or more storage groups listed in the Storage Groups table to attach to this partition. You can select FCP storage groups, FICON storage groups, or both types of storage group.
 - The suggested practice is to select storage groups that are in the Complete fulfillment state, but you can select any storage group except for those with a fulfillment state of Incomplete, or those that are already attached to the maximum number of partitions. If you do select groups in states other than Complete, some storage might not be available for use when you start the partition.
 - Use the additional information in the Storage Groups table, as necessary, to decide which storage groups to attach.
- c. When you have finished selecting storage groups to attach, select **OK** to close the Attach Storage Groups window.
- d. Check the entries for the storage groups that you selected, which are now displayed in the Storage Groups table in the Storage section. If necessary, you can use the minus icon in the table toolbar to remove a storage group from the table.
- e. When you have finished, review another section or click **OK** to save the partition definition.

Access FCP disks through HBAs

Host bus adapters (HBAs) provide a partition with access to external storage area networks (SANs) and devices that are connected to a system. Each HBA represents a unique connection between the partition and a physical FICON channel that is configured on the

system. When you create an HBA, you can select the adapter that you want to use from a list of all of the storage adapters that are currently configured on the system.

- For availability, select at least two storage adapters of the same type, and create an HBA for each one.
- If you are creating a Secure Service Container partition to install a software appliance, define at least one HBA to access the storage device on which the appliance installation image resides.

To work with the information on the Storage page, complete the following steps.

- a. When you first use the **New Partition** task, the Storage display contains an empty HBAs table. From the Actions list in the HBAs table, select **New** to open the New Host Bus Adapter window.
- b. On the New Host Bus Adapter window, define an HBA for each storage area network that is required for the applications that run on this partition. For each HBA that you define, complete the following steps. For more detailed descriptions of the New Host Bus Adapter window elements, see the online help.
 - 1) Enter a unique, meaningful name and, optionally, a description of the new HBA.
 - 2) Review the entries in the Adapter Ports table to determine which storage adapters are configured on the system.
 - a) Check the percentage listed in the Adapter HBA Allocation column. If the percentage is high (for example, 90%) for a specific port, consider selecting a different port.
 - b) Look for a warning icon next to the name in the Adapter Name column; if the warning icon is displayed for a specific port, select a different one.
 - c) Select one port by clicking the radio button in the Select column.
 - 3) Click **OK** to create the new HBA and close the New Host Bus Adapter window.
- c. Check the entry for the new HBA that is displayed in the HBAs table in the Storage section. If your company uses a specific numbering convention for its storage networks, you can change a system-generated device number by selecting the **Details** action and editing the HBA device number.
- d. Repeat the preceding steps, as necessary, to create a new HBA for each storage area network that your new partition requires.
- e. When you have finished, click **Next** to navigate to the next page in the task.

The next page to open might be either Accelerators, Cryptos, or Boot, depending on the system configuration. If the system does not have any configured accelerators or cryptographic features, the Accelerators or Cryptos page cannot be accessed.

7. If the system has configured accelerators, use the Accelerators page to enable the new partition to use accelerators that it requires.

Accelerators are optional features and, therefore, might not be installed on the system. If none are installed, the Accelerators page is disabled, and you can skip to step 8 of these instructions.

If accelerators are installed on the system but you are creating a partition only to familiarize yourself with the process, you can skip to step 8. Otherwise, use instructions in the online help to enable your partition to use accelerators.

8. If the system has configured cryptographic features, use the Cryptos page to enable the new partition to use the cryptographic features that it requires.

Crypto features are optional and, therefore, might not be installed on the system. If none are installed, the Cryptos page is disabled, and you can skip to step 9 on page 44 of these instructions.

If cryptographic features are installed on the system but you are creating a partition only to familiarize yourself with the process, you can skip to step 9 on page 44. Otherwise, use instructions in the online help to enable your partition to use cryptographic adapters.

9. Required: Use the Boot page to select the location of the executables for the hypervisor or operating system to be run on this partition, or to upload the required files to initialize the hypervisor or operating system when the partition itself is started.

Some of these boot options require that you find and select an ISO image file, which is a collection of files and metadata for installing software, and an .INS file, which maps image components (for example, kernel, ramdisk, parameter file) to the appropriate storage addresses in main memory.

The “Boot from” menu lists the boot options that are available for the hypervisor or operating system. If an option in the list is disabled, hover your cursor over that option to display additional information for that option. If necessary, take appropriate action to make that selection available; for example, if you want to use the Storage device (SAN) option, return to the Storage page to attach a storage group with a boot volume, or to define an HBA.

If you have selected the partition type **Secure Service Container**, note that option set in the “Boot from” menu is also **Secure Service Container**. This boot option cannot be changed unless you first change the partition type.

To define a boot option for other types of partitions, complete the following steps.

- Click the down arrow to display the available options in the “Boot from” list.
- Choose one of the available options and provide any additional information that is required.




Figure 18. Boot page of the New Partition task in basic mode

When you select a specific boot option, the display shows editable fields and other information related to the selected option. The following list describes each boot option, and provides instructions for providing any required information.

None Select this option if you want to start a partition without a hypervisor or operating system. Although the partition can be started, it is not in a usable state. This option is the default for partitions with a partition type of **Linux** and **z/VM**.

Storage device (SAN)

Select this option when the hypervisor or operating system executables reside on a device in a storage area network connected to the DPM-enabled system. This option is available only when storage groups or host bus adapters (HBAs) are defined for the partition.

When you select this option, the Boot section contains either a Storage Groups table or an HBA table. The Storage Groups table is displayed only when the DPM R3.1 storage management feature is applied on the system. Follow the instructions that correspond to the type of table displayed on the page.

Boot from a boot volume in an FCP or FICON storage group

The Storage Groups table displays the available FCP or FICON storage groups

that contain a boot volume. To view the available boot volumes, expand any table entry by selecting the storage group. The expanded table display contains a Boot Volume table that lists all available boot volumes that the storage group contains. The Boot Volume table content and Advanced Boot Volume Settings fields vary, depending on the storage group type.

- For each boot volume in an FCP storage group, the Boot Volume table provides the universally unique identifier (UUID) and capacity of the volume, along with a user-supplied description, if any.
- For each boot volume in a FICON storage group, the Boot Volume table provides the name of the storage subsystem in which the volume resides, along with the volume ID, capacity, type, and device number. If a user-supplied description is available, it is also displayed in the table.

Select the boot volume on which the operating system or hypervisor image resides.

Boot from a boot volume accessed through an HBA

The HBA table displays the available host bus adapters. Select the HBA connected to the storage subsystem that hosts the boot volume, provide the 64-bit worldwide port number (WWPN) of the storage subsystem, and provide the 64-bit hexadecimal logical unit number (LUN) of the volume that contains the boot image. For example:

Target WWPN: 50:0a:09:85:87:09:68:ad or 500a0985870968 (hexadecimal)

Target LUN: 4021400000000000

Network server (PXE)

Select this option when you want to use a preboot execution environment (PXE) on a network server. This option is available only if a network interface card (NIC) for either an OSA port or HiperSockets switch is defined for the partition.

When you select this option, the NIC table displays the available network interface cards. Select the NIC for the adapter that connects the partition to the network on which the network boot server resides.

FTP server

Select this option if you want to use FTP to boot an image that is located on a different system. Provide the following information:

Host name

Enter either the fully qualified domain name of the FTP server, or its IP address.

User name

Enter the user name on the target FTP server.

Password

Enter the password associated with the user name on the target FTP server.

INS file

Either click **Browse** to retrieve a list of INS files from the target FTP server and select one file, or enter the fully qualified name (relative to FTP root) of an INS file.

Depending on the size of the FTP site, browsing might require more time than manually entering the full path and name of the INS file. Also note that the browsing function returns INS files found in the user's home directory or its subdirectories. Because you cannot select a starting directory, or navigate to a directory above the user's home directory, manually entering the full path and name of the INS file might be more expedient.

If you click **Browse**, a separate window displays the user's home directory and its subdirectories. Select one INS file, and click **OK** to close the Browse FTP Server window.

FTPS server

Select this option if you want to use the FTP Secure (FTPS) protocol to boot an image that is located on a different system. FTPS uses the Secure Socket Layer (SSL) protocol to secure data. With this option, you need to supply a host name, user ID, password, and .INS file, as described for the **FTP server** boot option.

SFTP server

Select this option if you want to use the Secure File Transfer Protocol (SFTP) to boot an image that is located on a different system. SFTP uses the Secure Shell (SSH) protocol to secure data. With this option, you need to supply a host name, user ID, password, and .INS file, as described for the **FTP server** boot option.

Hardware Management Console removable media

Select this option if you want to use an INS file from a media drive that is connected to the HMC. The media drive must be available when you are creating the partition definition and when the partition is started.

When you select this option:

- 1) Identify the type of media drive on which the INS file resides by clicking the appropriate radio button.
- 2) Either enter the fully qualified name (relative to the mount point) of an INS file, or click **Browse** to start a search on the target media drive to retrieve a list of INS files. Any INS files found are displayed in a separate window. Select only one INS file and click **OK** to close the Browse Removable Media window.

ISO image

Select this option when you want to upload an ISO file that is located on your workstation file system. This option is available only when you are connecting to the HMC through a remote browser.

When you select this option:

- 1) Click **Browse** to find the ISO image file on your workstation file system. You cannot select an ISO image from an HMC media drive. As soon as you select an ISO image file, DPM starts to upload the file, and displays a progress indicator for the upload operation.
- 2) After the upload operation completes, click **Browse** to search the ISO image file for the INS file that you want to use. Any INS files found are displayed in a separate window. Select only one INS file and click **OK** to close the Browse ISO Image window.

c. When you have finished, click **Next** to navigate to the next page in the task.

10. Use the Summary page to review the properties for the new partition.

You might need to vertically scroll the page to view all of the partition properties. If necessary, click **Back** to return to a particular page to change a property value or setting.

11. Required: On the Summary page, click **Finish** to save the partition definition. A progress indicator is displayed until DPM finishes creating the partition.

Results

DPM opens the validation window when it finishes creating the partition definition. The validation window displays a message indicating that your partition has been created, and lists additional tasks that you can use to work with the new partition.




Figure 19. Sample validation window for the New Partition task

Note: The tasks that are listed in the validation window vary, depending on whether the DPM R3.1 storage management feature is applied to your system.

What to do next

To work with the partition, click any of the links on the validation window to open a related task in a separate window.

Export WWPNs

Opens the **Export WWPNs** task, through which you can obtain a list of worldwide port names of the host bus adapters that are defined for a partition. If you defined HBAs for this partition, or specified a boot option that requires access to a storage device through an HBA, use the **Export WWPNs** task before you start the partition. Your storage administrator needs this list of WWPNs to complete zoning and LUN masking so your partition can successfully access the storage devices.

Note: You do not need to use the **Export WWPNs** task if you are working on a system that has the DPM R3.1 storage management feature applied, because WWPNs are assigned when FCP storage groups are configured.

Start the partition

Opens the **Start** task, with this partition selected as the partition to start.

Setup auto-start

Opens the Start Options section in the **System Details** task, through which you can specify that the partition is to be started automatically, when the system is started. You can specify the order in which the partition is started, relative to other partitions on the system, or add the partition to an auto-start group.

Setup scheduled operations

Opens the **Customize Scheduled Operations** task, through which you can automate start and stop operations for a partition by scheduling them to be run on specific dates and times.

Create monitors

Opens the Monitor System Events task, through which you can define specific events for which you want to be notified, if these events occur. For example, you can request to be notified when network traffic on a system, or processor utilization on a partition, reaches a specific threshold that you set.

When you are finished reviewing the information on the validation window or using the provided links to related tasks, click **Close** to close the validation window.

Starting a partition and its operating system or hypervisor

This procedure provides step-by-step instructions for starting a partition with a type of **Linux** or **z/VM**, and its operating system or hypervisor.

Before you begin

- For partitions with a type of **Secure Service Container**, see the *Secure Service Container User's Guide*, SC28-6978, for information about starting and managing Secure Service Container partitions and their appliances. This book is available on <http://www.ibm.com/servers/resourcelink>.
- Make sure that you log in to the Hardware Management Console (HMC) with a user ID that has authorization to use the **Start** task to start a partition. You can use either a user ID that a system administrator authorized to this task through customization controls in the **User Management** task, or one of the default user IDs listed for the **Start** task in Appendix A, "DPM task and resource roles," on page 137.

Procedure

- On the main HMC page, expand the **Systems Management** node to view managed systems, and select the DPM-enabled system on which the partition resides.
- On the **Partitions** tab for the selected system, find the table entry for the partition that you want to start, and select it.
- Select the double-arrow icon in the Name field to display the task menu, and select **Partition Details** to verify the boot option for this partition.
 - If you have not already done so, use the Boot page to select the location of the executables for the hypervisor or operating system to be run on this partition, or to upload the required files to initialize the hypervisor or operating system when the partition itself is started.
 - Select **OK** to apply any changes and close the Partition Details window.
- Use the **Start** task to start the selected partition.
 - From the entry for the selected partition on the **Partitions** tab, select the double-arrow icon in the Name field to display the cascading task menu.
 - Expand the **Daily** group, and select **Start**.

If the DPM R3.1 storage management feature has been applied to the system, and one or more of the partitions to be started have attached storage groups that are being configured or modified, a warning message is displayed. The warning message includes the name of the affected partitions. The Start task does not continue until you make a selection.

- Select **YES** to allow the affected partitions to be started.
- Select **NO** to cancel the start operation for only the affected partitions.

A new window opens to display the progress of the start operation. In addition to displaying a progress indicator, this window also contains a Details column with messages that are updated as the start process continues. These messages indicate the progress of configuring partition resources and initializing the operating system or hypervisor to run in the partition.

Results

The Details column contains one of the following labels, with a clickable **Details** link that identifies the failed partitions and provides a message that explains the failure.

- Success** indicates partitions that have started.
- Failed** indicates partitions that failed to start.
- Cancelled** indicates partitions for which the start operation was canceled.

If the result is anything other than successful, use the information in the Details column to diagnose and correct the problem.

What to do next

Open the appropriate console task through which you can log in to the operating system or hypervisor that is running on the partition.

- To log in to a Linux system, use the **Operating System Messages** task or the **Integrated ASCII Console** task. The **Integrated ASCII Console** task must be enabled through the operating system before you can use it.
- To log in to a z/VM hypervisor that is hosting multiple Linux systems, use the **Integrated 3270 Console** task.

Configuring partition resources on the operating system

If this time is the first time that you have started this partition, you need to configure the partition resources (processors, memory, and adapters) through configuration files on the operating system. The suggested practice is to open the **Partition Details** task and use it as a reference as you create or modify the appropriate configuration files on the operating system.

Depending on the version of Linux that you have installed, the operating system might automatically configure some resources.

- Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) devices are automatically configured when you are running any of the minimum supported Linux versions. For recommended Linux on Z distribution levels on IBM Z and LinuxONE servers, see the IBM tested platforms at: <http://www.ibm.com/systems/z/os/linux/resources/testedplatforms.html>
- Auto-configuration of other devices requires a version of the Linux operating system that supports auto-configuration. These other devices include Fibre Channel connections (FICON) in Fibre Channel Protocol (FCP) mode, IBM HiperSockets, and Open Systems Adapter-Express (OSA-Express) devices. See the Red Hat, SUSE, or Ubuntu product information page to determine which RHEL, SLES, or Ubuntu Server version provides this support.

Ensuring that storage groups are visible to the Linux operating system

After you start the new partition, you might need to enter Linux commands to make the storage groups available to the operating system that the partition hosts. The actions required depend on the fulfillment state and type of the storage group, and whether the storage group contained the boot volume for the operating system. Typically, the operating system stores the FCP HBA or FICON volume configuration so it can automatically bring the devices online on the next reboot, so you need to take action only for the initial boot of the operating system.

When attaching a storage group in Complete state

- For an FCP storage group:
 - If the storage group contained the boot volume, the operating system brings online all of the HBAs for this storage group, and all volumes in the storage group are available. No action is required unless you have attached other storage groups.
 - If the storage group does not contain the boot volume, and the operating system is not configured to bring HBAs online automatically, you need to issue the **chccwdev** command to bring online all of the HBAs.
- For a FICON storage group, the operating system brings online only the boot volume. You need to issue the **chccwdev** command to bring online all of the remaining volumes in the storage group that contains the boot volume, as well as the volumes in any other storage groups that you attached.

When attaching an unfulfilled storage group that becomes Complete as the partition is running

- For an FCP storage group:
 - If adapters were assigned to HBAs while the partition is running, you need to use the **chchp** command to activate the channel paths for those new adapters.

- To access the volumes in the storage group, you need to issue the **chccwdev** command to bring online all of the HBAs.
- For a FICON storage group:
 - If the adapters connecting the storage group to the storage subsystem were assigned while the partition is running, use the **chchp** command to activate the channel paths for those new adapters.
 - All volumes are offline. You need to issue the **chccwdev** command to bring online all of the volumes in the storage group.

To find the IDs that you need to use for the Linux commands, use the following tasks.

- HBA device numbers are available in the Host Bus Adapters (HBA) table when you expand the storage group table entry in the Storage section of the **Partition Details** task.
- Channel path IDs for FCP adapters are shown in the Host Bus Adapters (HBA) table when you expand the storage group table entry in the Storage section of the **Partition Details** task.
- Channel path IDs for FICON adapters are shown on the **ADAPTERS** tab of the Storage Group details; open the **Configure Storage** task and select the storage group in the **Storage Overview** to open the Storage Group details page.
- FICON volume device numbers are shown on the **VOLUMES** tab of the Storage Group details page; open the **Configure Storage** task and select the storage group in the **Storage Overview** to open the Storage Group details page.

Verifying that the partition resources are online

To verify that the partition resources are online, use the appropriate Linux commands, samples of which are displayed in the following list.

- To display information about processor resources, use the **lscpu** command. The following screen shows a sample display that results from entering this command.

```
[root@lbskvm4 ~]# lscpu
Architecture:          s390x
CPU op-mode(s):        32-bit, 64-bit
Byte Order:            Big Endian
CPU(s):                 8
On-line CPU(s) list:   0-7
Thread(s) per core:    2
Core(s) per socket:    8
Socket(s) per book:    3
Book(s):                8
Vendor ID:              IBM/S390
BogoMIPS:               7936.00
Hypervisor:             PR/SM
Hypervisor vendor:      IBM
Virtualization type:    full
Dispatching mode:       horizontal
L1d cache:              128K
L1i cache:              96K
L2d cache:              2048K
L2i cache:              2048K
[root@lbskvm4 ~]#
```

Figure 20. Sample displays resulting from the **lscpu** command

- To display information about memory resources, use the **lsmem** command. The following screen shows a sample display that results from entering this command.

```
[root@lbskvm4 ~]# lsmem
```

Address Range	Size (MB)	State	Removable	Device
0x0000000000000000-0x000000001fffffffff	8192	online	no	0
0x0000000020000000-0x0000000f5fffffffff	999424	online	yes	1-122
0x0000000f60000000-0x0000000fffffffff	40960	online	no	123-127
0x0000010000000000-0x0000030dffffffff	2154496	offline	-	128-390

```

Memory device size : 8192 MB
Memory block size : 8192 MB
Total online memory : 1048576 MB
Total offline memory: 2154496 MB
[root@lbskvm4 ~]#

```

Figure 21. Sample displays resulting from the **lsmem** command

- To display information about adapters, use the appropriate command for the device type. For example, to view Open Systems Adapter-Express (OSA-Express) features, use the **lsqeth**, **lscss**, and **lschp** device driver commands. The following screens show sample displays that result from entering these commands.

```
0.0.0005 host1
[root@lbskvm4 ~]# lsqeth
```

```
Device name : encw0.0.0001
```

```

card_type : OSD_1000
cdev0 : 0.0.0001
cdev1 : 0.0.0002
cdev2 : 0.0.0003
chpid : 13
online : 1
portname : DUMMY
portno : 0
state : UP (LAN ONLINE)
priority_queueing : always queue 0
buffer_count : 128
layer2 : 1
isolation : none

```

```
[root@lbskvm4 ~]# lscss
```

Device	Subchan.	Devtype	CU	Type	Use	PIM	PAM	POM	CHPIDs
0.0.0001	0.0.0000	1732/01	1731/01	yes	80	80	ff	13000000	00000000
0.0.0002	0.0.0001	1732/01	1731/01	yes	80	80	ff	13000000	00000000
0.0.0003	0.0.0002	1732/01	1731/01	yes	80	80	ff	13000000	00000000
0.0.0006	0.0.0003	1732/03	1731/03		80	80	ff	39000000	00000000
0.0.0005	0.0.0004	1732/03	1731/03	yes	80	80	ff	2a000000	00000000
0.0.0004	0.0.0005	1732/03	1731/03	yes	80	80	ff	2c000000	00000000
0.0.0007	0.0.0006	1732/03	1731/03		80	80	ff	01000000	00000000

```
[root@lbskvm4 ~]#
```

Figure 22. Sample displays resulting from the **lsqeth** and **lscss** commands

```
[root@lbskvm4 ~]# lschp
CHPID  Vary  Cfg.  Type  Cmg  Shared  PCHID
=====
0.01   1     1     25    2     1     01f0
0.13   1     1     11    2     1     017c
0.2a   1     1     25    2     1     0121
0.2e   1     1     25    2     1     0109
0.39   1     1     25    2     1     0144
[root@lbskvm4 ~]#
```

Figure 23. Sample displays resulting from the **lschp** command

Specifying the relative port number of an OSA device

If the partition is connected to a network through an OSA-Express adapter port other than port 0, you need to manually specify the relative port number through a Linux qeth device driver command, before entering the Linux command to bring the device online. The following sample commands show how to create a device group, to specify the relative port number and layer mode, and to bring the group of devices online. The highlighted command (the second line) specifies the port number; that command contains 1 for the port number, along with the attribute portno.

```
echo 0.0.1100,0.0.1101,0.0.1102 > /sys/bus/ccwgroup/drivers/qeth/group
echo 1 > /sys/bus/ccwgroup/drivers/qeth/0.0.1100/portno
echo 1 > /sys/bus/ccwgroup/drivers/qeth/0.0.1100/layer2
echo 1 > /sys/bus/ccwgroup/drivers/qeth/0.0.1100/online
```

Finding additional information about operating system or hypervisor commands

- For more information about using Linux commands to work with partition resources and adapters, see the *Linux on Z Device Drivers, Features, and Commands* documentation for the Linux kernel version that you are using. This documentation, which also describes commands and parameters for configuring Linux on Z, is available in IBM Knowledge Center at http://www.ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz_r_dd.html
- For information about using z/VM commands to work with partition resources and adapters, see the *z/VM: CP Commands and Utilities Reference* for the z/VM version that you are using. The z/VM library is available at <http://www.vm.ibm.com/library/>

Chapter 8. Dynamically modifying the resources of an active partition

You can use the **Partition Details** task to add processors, memory, and devices to an active partition, without stopping and restarting it. When you click **Apply** to save any changes you have made through the **Partition Details** task, DPM updates the partition definition but does not bring any of the new resources or devices online. To do so, you must use the appropriate operating system or hypervisor commands. Note that you cannot change the partition type through the **Partition Details** task.

Depending on the version of Linux that you have installed, the operating system might automatically configure some partition resources:

- Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) devices are automatically configured when you are running any of the minimum supported Linux versions. For recommended Linux on Z distribution levels on IBM Z and LinuxONE servers, see the IBM tested platforms at: <http://www.ibm.com/systems/z/os/linux/resources/testedplatforms.html>
- Auto-configuration of other devices requires a version of the Linux operating system that supports auto-configuration. These other devices include Fibre Channel connections (FICON) in Fibre Channel Protocol (FCP) mode, IBM HiperSockets, and Open Systems Adapter-Express (OSA-Express) devices. See the Red Hat, SUSE, or Ubuntu product information page to determine which RHEL, SLES, or Ubuntu Server version provides this support.

To manually configure partition resources on the HMC and on the operating system that a partition hosts, use the following examples as models for your changes. Note that the content of the screen captures in these topics might differ from what you see on the HMC, depending on the DPM version and the host system that you are using.

- “Example: Adding more processor resources”
- “Example: Adding more memory resources” on page 55
- “Example: Adding a new network device” on page 57

If the system has the DPM R3.1 storage management feature installed and activated, use the following instructions to dynamically add storage resources for a partition to use.

- “Instructions: Attaching storage groups” on page 60
- “Instructions: Detaching and reattaching an FCP storage group” on page 61
- “Instructions: Changing the adapters that are assigned to an FCP storage group” on page 63

Example: Adding more processor resources

Consider the display in the following screen, which shows that a sample partition has two threads per core, with eight processors online.

```

[root@lpskvm4 /]# lscpu
Architecture:          s390x
CPU op-mode(s):        32-bit, 64-bit
Byte Order:            Big Endian
CPU(s):                 8
On-line CPU(s) list:   0-7
Thread(s) per core:    2
Core(s) per socket:    8
Socket(s) per book:    3
Book(s):               8
Vendor ID:             IBM/S390
BogoMIPS:              7936.00
Hypervisor:            PR/SM
Hypervisor vendor:     IBM
Virtualization type:   full
Dispatching mode:      horizontal
L1d cache:             128K
L1i cache:             96K
L2d cache:             2048K
L2i cache:             2048K

```

Figure 24. Sample display for the **lscpu** command showing current processor resources

To bring more processors online, complete the following steps.

1. Use the **Partition Details** task to increase the number of processors for this partition from four to six cores, and click **Apply** to save your changes and close the Partition Details window.




Figure 25. Sample screen of the Processors section in the Partition Details window

2. Through the Linux **lscpu** command, verify that the number of processor cores has increased.

```
[root@lbskvm4 /]# lscpu
Architecture:          s390x
CPU op-mode(s):        32-bit, 64-bit
Byte Order:            Big Endian
CPU(s):                 12
On-line CPU(s) list:    0-7
Off-line CPU(s) list:   8-11
Thread(s) per core:     2
Core(s) per socket:     8
Socket(s) per book:     3
Book(s):                8
Vendor ID:              IBM/S390
BogoMIPS:               7936.00
Hypervisor:             PR/SM
Hypervisor vendor:      IBM
Virtualization type:    full
Dispatching mode:       horizontal
L1d cache:              128K
L1i cache:              96K
L2d cache:              2048K
L2i cache:              2048K
```

Figure 26. Sample display for the **lscpu** command showing additional processor resources

3. Use the Linux **chcpu** command to configure and then enable the additional processor cores.

```
[root@lbskvm4 /]# chcpu -c 8-11
CPU 8 configured
CPU 9 is already configured
CPU 10 configured
CPU 11 is already configured
[root@lbskvm4 /]# chcpu -e 8-11
CPU 8 enabled
CPU 9 enabled
CPU 10 enabled
CPU 11 enabled
[root@lbskvm4 /]#
```

Figure 27. Sample display for the **lscpu** command showing configured processor resources

Example: Adding more memory resources

Consider the display in the following screen, which shows that a sample partition has 1 terabyte (TB) of memory online, and 2 TB offline. It also shows that the memory size is in 8 gigabyte (GB) increments.

```
[root@lbskvm4 ~]# lsmem
```

Address Range	Size (MB)	State	Removable	Device
0x0000000000000000-0x00000001ffffffff	8192	online	no	0
0x0000000200000000-0x000000f5ffffffff	999424	online	yes	1-122
0x000000f600000000-0x000000ffffffff	40960	online	no	123-127
0x0000010000000000-0x0000030dffffffff	2154496	offline	-	128-390

```

Memory device size : 8192 MB
Memory block size : 8192 MB
Total online memory : 1048576 MB
Total offline memory: 2154496 MB
[root@lbskvm4 ~]#

```

Figure 28. Sample display for the **lsmem** command

To bring more memory online, complete the following steps.

1. Use the **Partition Details** task to increase the amount of memory for this partition from 1 TB to 2 TB, and click **Apply** to save your changes and close the Partition Details window.




Figure 29. Sample screen of the Memory section in the Partition Details window

2. Then, through the Linux **chmem** command, increase the amount of online memory in 8 GB increments.

```
[root@lbskvm4 /]# time chmem -e 8g
```

```
real    0m0.691s
user    0m0.157s
sys     0m0.553s
```

```
[root@lbskvm4 /]# lsmem
```

Address Range	Size (MB)	State	Removable	Device
0x0000000000000000-0x00000001ffffffff	8192	online	no	0
0x0000000200000000-0x000000f5ffffffff	999424	online	yes	1-122
0x000000f600000000-0x000000ffffffff	40960	online	no	123-127
0x0000010000000000-0x00000103ffffffff	16384	online	yes	128-129
0x0000010400000000-0x0000030dffffffff	2138112	offline	-	130-390

```
Memory device size : 8192 MB
Memory block size : 8192 MB
Total online memory : 1064960 MB
Total offline memory: 2138112 MB
[root@lbskvm4 /]#
```

Figure 30. Sample display for the **chmem** command

Example: Adding a new network device

In this example, the Linux system currently has seven network devices configured, as shown in the following display.

```
[root@lbskvm4 ~]# lscss
```

Device	Subchan.	DevType	CU	Type	Use	PIM	PAM	POM	CHPIDs
0.0.0001	0.0.0000	1732/01	1731/01	yes	80	80	ff	13000000	00000000
0.0.0002	0.0.0001	1732/01	1731/01	yes	80	80	ff	13000000	00000000
0.0.0003	0.0.0002	1732/01	1731/01	yes	80	80	ff	13000000	00000000
0.0.0006	0.0.0003	1732/03	1731/03		80	80	ff	39000000	00000000
0.0.0005	0.0.0004	1732/03	1731/03	yes	80	80	ff	2a000000	00000000
0.0.0004	0.0.0005	1732/03	1731/03	yes	80	80	ff	2e000000	00000000
0.0.0007	0.0.0006	1732/03	1731/03		80	80	ff	01000000	00000000

```
[root@lbskvm4 ~]#
```

Figure 31. Sample display for the **lscss** command showing current devices

1. Use the **Partition Details** task to add a new network interface card (NIC).




Figure 32. Sample screen of the Network section with the New NIC icon highlighted

2. Select the adapter for the new network connection.

New Network Interface Card

Provide a name and description for the new NIC, and then select the backing adapter port or switch.

* Name:

Description:

Device Number:

Adapter Ports and Switches

Adapter Name	Adapter Port	Card Type	Uplink Utilization ?	Adapter NIC Allocation ?	Location	Description
☐ OSD 017C A13B-38	0	OSA-Express5s 1000Base-T	0%	0%	A13B-D138J.01-D238J.01	
☐ OSD 017C A13B-38	1	OSA-Express5s 1000Base-T	0%	0%	A13B-D138J.01-D238J.01	
☒ OSD 0188 A06B-03	0	OSA-Express5s 1000Base-T	0%	0%	A06B-D103J.01-D203J.01	
☐ OSD 0188 A06B-03	1	OSA-Express5s 1000Base-T	0%	0%	A06B-D103J.01-D203J.01	
☐ OSM 0100 A13B-01	0	OSA-Express5s	0%	0%	A13B-D101J.01-	

Total: 6 Selected: 1

Figure 33. Sample screen of the New Network Interface Card window

3. Click **OK** to save the new NIC definition, and check the updated NICs table in the Network section of the Partition Details task.




Figure 34. Sample screen of the Network section with the newly added NIC

- Click **Apply** to save your changes and close the Partition Details window.
- Then, through the Linux **znetconf** command, define the new devices 0.0.0008 through 0.0.000a.

```
[root@lbskvm4 /]# znetconf -u
Scanning for network devices...
Device IDs          Type      Card Type      CHPID Drv.
-----
0.0.0008,0.0.0009,0.0.000a 1731/01 OSA (ODIO)      14 qeth
[root@lbskvm4 /]# znetconf -a 0.0.0008 -o layer2=1
Scanning for network devices...
Successfully configured device 0.0.0008 (enccw0.0.0008)
[root@lbskvm4 /]# znetconf -c
Device IDs          Type      Card Type      CHPID Drv. Name      State
-----
0.0.0001,0.0.0002,0.0.0003 1731/01 OSD_1000      13 qeth enccw0.0.0001  online
0.0.0008,0.0.0009,0.0.000a 1731/01 OSD_1000      14 qeth enccw0.0.0008  online
[root@lbskvm4 /]#
```

Figure 35. Sample display for the **znetconf** command showing current devices

Note: If the display for the **znetconf -u** command does not list the newly added device, the adapter might be offline. Use the **lschp** command to determine the current state of the new adapter, and use the **chchp** command to bring the adapter online. Then issue the **znetconf** command again.

Issuing an **lsqeth** command displays the successfully defined network device, which can be defined to the Linux TCP/IP stack or passed to a virtual switch. Figure 36 on page 60 illustrates a sample display for the **lsqeth** command.

```

[root@lbskvm4 /]# lsqeth
Device name           : enccw0.0.0001
-----
card type             : OSD 1000
cdev0                 : 0.0.0001
cdev1                 : 0.0.0002
cdev2                 : 0.0.0003
chpid                 : 13
online                : 1
portname              : DUMMY
portno                : 0
state                 : UP (LAN ONLINE)
priority_queueing     : always queue 0
buffer_count          : 128
layer2                : 1
isolation              : none

Device name           : enccw0.0.0008
-----
card_type             : OSD_1000
cdev0                 : 0.0.0008
cdev1                 : 0.0.0009
cdev2                 : 0.0.000a
chpid                 : 14
online                : 1
portno                : 0
state                 : UP (LAN ONLINE)
priority_queueing     : always queue 0
buffer_count          : 64
layer2                : 1
isolation              : none

[root@lbskvm4 /]#

```

Figure 36. Sample display for the `lsqeth` command showing details for devices

Instructions: Attaching storage groups

Through the Storage section of the **Partition Details** task, you can attach one or more storage groups to an existing partition. When the DPM R3.1 storage management feature is applied on the system, the Storage section contains a Storage Groups table that lists any currently attached storage groups. To attach new storage groups to the partition, complete the following steps.

1. In the Storage Group table toolbar, select the plus icon to open the Attach Storage Groups window.
2. On the Attach Storage Groups window, select one or more storage groups listed in the Storage Groups table to attach to this partition. You can select FCP storage groups, FICON storage groups, or both types of storage group.
 - The suggested practice is to select storage groups that are in the Complete fulfillment state, but you can select any storage group except for those with a fulfillment state of Incomplete, or those that are already attached to the maximum number of partitions. If you do select groups in states other than Complete, some storage might not be available for use immediately.
 - Use the additional information in the Storage Groups table, as necessary, to decide which storage groups to attach. For descriptions of the columns in the Storage Groups table, see the online help.
3. When you have finished selecting storage groups to attach, select **OK** to close the Attach Storage Groups window.

4. Check the entries for the storage groups that you selected, which are now displayed in the Storage Groups table in the Storage section. If necessary, you can use the minus icon in the table toolbar to remove a storage group from the table.
5. When you have finished, click **OK** to save the partition definition.

If the partition is running, or when you restart a stopped partition, you might need to enter Linux commands to make any newly attached storage groups available to the operating system that the partition hosts. The actions required depend on the fulfillment state and type of the storage group, and whether the storage group contained the boot volume for the operating system.

When attaching a storage group in Complete state when the partition is stopped

- For an FCP storage group:
 - If the storage group contained the boot volume, the operating system brings online all of the HBAs for this storage group, and all volumes in the storage group are available. No action is required unless you have attached other storage groups.
 - If the storage group does not contain the boot volume, and the operating system is not configured to bring HBAs online automatically, you need to issue the **chccwdev** command to bring online all of the HBAs.
- For a FICON storage group, the operating system brings online only the boot volume. You need to issue the **chccwdev** command to bring online all of the remaining volumes in the storage group that contains the boot volume, as well as the volumes in any other storage groups that you attached.

When attaching a Complete storage group to a running partition, or attaching an unfulfilled storage group that becomes Complete as the partition is running

- For an FCP storage group:
 - If adapters were assigned to HBAs while the partition is running, you need to use the **chchp** command to activate the channel paths for those new adapters.
 - To access the volumes in the storage group, you need to issue the **chccwdev** command to bring online all of the HBAs.
- For a FICON storage group:
 - If the adapters connecting the storage group to the storage subsystem were assigned while the partition is running, use the **chchp** command to activate the channel paths for those new adapters.
 - All volumes are offline. You need to issue the **chccwdev** command to bring online all of the volumes in the storage group.

To find the IDs that you need to use for the Linux commands, use the following tasks.

- HBA device numbers are available in the Host Bus Adapters (HBA) table when you expand the storage group table entry in the Storage section of the **Partition Details** task.
- Channel path IDs for FCP adapters are shown in the Host Bus Adapters (HBA) table when you expand the storage group table entry in the Storage section of the **Partition Details** task.
- Channel path IDs for FICON adapters are shown on the **ADAPTERS** tab of the Storage Group details; open the **Configure Storage** task and select the storage group in the **Storage Overview** to open the Storage Group details page.
- FICON volume device numbers are shown on the **VOLUMES** tab of the Storage Group details page; open the **Configure Storage** task and select the storage group in the **Storage Overview** to open the Storage Group details page.

Instructions: Detaching and reattaching an FCP storage group

When you detach an FCP storage group from a partition, DPM does not preserve the HBAs, device numbers, or backing adapters that were in use for that storage group. Consequently, if you reattach the storage group to the same partition, the device numbers and the backing adapters of the HBAs are not guaranteed to be the same as they were before the detachment. The same condition is true if you attach

the FCP storage group to a new partition: the device numbers and the backing adapters of the HBAs are not guaranteed to be the same as they were when the storage group was attached to a different partition.

If the operating system image for the partition resides on a boot volume in the storage group that you reattach, the operating system might not start when the partition is restarted. (Operating systems that are started from a storage volume usually have a preconfigured device number and a path to the volume.) To avoid this situation, administrators must review the device numbers and the backing adapters of the HBAs, after reattaching the storage group and before restarting the partition. If necessary, the administrators must change the device number to match the preconfigured device number for the operating system, and make sure that the preconfigured device number is assigned to the backing adapter that was assigned to the HBA before the detachment.

Notes:

- If the storage group that is being reattached contains the boot volume for the operating system, you must stop the partition before detaching and reattaching the storage group. Running workloads can be disrupted if you detach and reattach a storage group that contains boot volumes while the partition is active. (If a storage group contains only data volumes, you do not have to stop the partition.)
- When you detach and reattach a FICON storage group, the device numbers of storage volumes do not change because they are associated with the storage group volumes.

Use the following instructions to detach an FCP storage group that contains a boot volume, and reattach it to the same partition.

1. If the partition is active, use the **Stop** task to stop the partition.
2. Open the **Partition Details** task for the partition.
 - a. Go to the **Boot** section. If the **Storage device (SAN)** boot option is selected, and a boot volume in an FCP storage group is selected, note the name of the storage group.
 - b. Go to the **Storage** section, and select the FCP storage group that contains the boot volume for the operating system.
 - 1) Expand the storage group table entry to show the Host Bus Adapters (HBA) table, and record the device number and the assigned adapter.
 - 2) In the Storage Group table toolbar, select the minus icon to detach the storage group.
3. When you reattach the storage group, make sure that the partition is stopped and, if necessary, reopen the **Partition Details** task to the **Storage** section.
 - a. In the Storage Group table toolbar, select the plus icon to open the Attach Storage Groups window.
 - b. On the Attach Storage Groups window, select the storage group.
 - c. Select **OK** to close the Attach Storage Groups window.
 - d. Select and expand the entry for the storage group, which is listed in the Storage Groups table in the **Storage** section.
 - 1) Look for the adapter that you recorded before detaching the storage group. If it is listed, continue to the next step (check the device number for the HBA). Otherwise, complete the following steps:
 - a) Select **Change Adapters** to open the FCP adapter assignment window.
 - b) In the Assigned Adapters table, select an existing adapter, and select **UNASSIGN**.
 - c) In the Adapter Candidates table, find the adapter that you recorded, and select **ASSIGN**.
 - d) Select **SAVE** to save your changes and return to the **Storage** section of the **Partition Details** task.
 - 2) Check the device number for the HBA that is associated with the recorded adapter. If necessary, change the device number to match the preconfigured device number that the operating system uses (the device number that you recorded before detaching the storage group).

- 3) Select **Apply** to save your changes.
- e. Go to the **Boot** section of the **Partition Details** task.
 - 1) Select **Storage device (SAN)** as the Boot from option.
 - 2) Use the radio button in the Select column to select the reattached FCP storage group that contains the boot volume for the operating system.
- f. Select **OK** to save your changes and close the Partition Details window.
4. Use the **Start** task to restart the partition.

Instructions: Changing the adapters that are assigned to an FCP storage group

When you open the Storage Group details page for an FCP storage group, you can review the assigned adapters and remove or replace them with other adapters that are available for use by a partition. However, if an FCP adapter is configured while the storage group is attached to an active partition, DPM cannot detect and list the new adapter as available for use by any partition. To make sure that you can choose a new adapter from a complete list of available adapters, perform the following steps on the HMC.

1. Open the **Configure Storage** task to the Storage Overview page, and select the table row for the FCP storage group. The Storage Details page opens.
2. Select the **PARTITIONS** tab to display a list of the partitions to which this storage group is attached. If any partitions are listed as Active, use the **Stop** task to stop them. After all of the active partitions are stopped, continue to the next step.
3. Select the **Connection Report** icon to open the Connection Report page for the FCP storage group. On the Connection Report page, select the **Update report** icon (🔄) to start a new background check of the available connections for the storage group. When the check is complete, the Report Date field is updated.
4. After the background check completes, verify that the new adapter is listed in the Mainframe section of the connection report, and that no errors are associated with the adapter. Close the Connection Report page.
5. On the **PARTITIONS** tab, expand the table row for one of the partition for which you want to use the new adapter, and select **Change Adapter**. On the FCP adapter assignment window, complete the following steps.
 - a. Review the entries in the Assigned Adapters table and, if necessary, select an existing adapter, and select **UNASSIGN** to remove it from the table.
 - b. In the Adapter Candidates table, find the table entry for the new adapter.
 - c. In the adapter table entry, select **ASSIGN** to assign the new adapter.
 - d. Select **SAVE** to return to the Storage Group details page.

Repeat these steps for each partition for which you want to use the new adapter.

6. Check the fulfillment state of the storage group. When the state is Complete, use the **Start** task to restart the partitions.

Chapter 9. Summary of tasks for managing systems, adapters, and partitions

DPM tasks are available starting with HMC/SE Version 2.13.1. Chapter 9, “Summary of tasks for managing systems, adapters, and partitions” provides an alphabetical summary of DPM tasks and other tasks that you might use to work with a DPM-enabled system, partitions, storage groups, and adapters.

- For more information about each task, see the online help on the HMC or SE.
- To use specific tasks, you need to log into the HMC or SE with a specific default user ID, or with a user ID that a system administrator has authorized to the task through customization controls in the **User Management** task. For information about authorization requirements for DPM tasks, see Appendix A, “DPM task and resource roles,” on page 137.
- These tasks can be accomplished programmatically as well, through the HMC Web Services application programming interfaces (APIs) for DPM. For information about the DPM APIs, see the appropriate edition of *Hardware Management Console Web Services API*, which is available through the Library link on IBM Resource Link at <http://www.ibm.com/servers/resourcelink>

Table 3. Summary of key HMC/SE tasks and displays for working with DPM-enabled systems, partitions, and adapters

Task name	Icon	On HMC or SE	Description
Configure Storage		HMC	Use the Configure Storage task to initially connect a system to storage devices, request storage resources for one or more partitions to use, and to view or modify the current storage configuration. This task is available for use only when the DPM R3.1 storage management feature is applied.
Customize Scheduled Operations		Both	Use the Customize Scheduled Operations task to customize a schedule for selected DPM-enabled systems. Scheduled operations are helpful for situations where automatic, delayed, or repetitious processing of system operations is necessary. A scheduled operation is started at a specified time, without operator assistance to perform the operation. A schedule can be set for one operation or repeated many times.
Delete Partition		HMC	Use the Delete Partition task to delete the definition associated with one or more selected partitions on a DPM-enabled system.
Disable Dynamic Partition Manager		SE only	Use the Disable Dynamic Partition Manager task to disable DPM, remove all partitions, and unconfigure all adapters from the system.
Dump (Partition)		HMC	Use the Dump task to initiate a partition dump by booting a stand-alone dump program on a DPM-enabled system.
Enable Dynamic Partition Manager		SE only	Use the Enable Dynamic Partition Manager task to enable DPM for the system and identify the OSA-Express5S 1000 Base-T adapters for z System Management (OSM).

Table 3. Summary of key HMC/SE tasks and displays for working with DPM-enabled systems, partitions, and adapters (continued)

Task name	Icon	On HMC or SE	Description
Getting Started with Dynamic Partition Manager		HMC	Use the Getting Started with Dynamic Partition Manager task for quick links to the Manage Adapters and New Partition tasks, along with step-by-step instructions for using the New Partition task in basic mode. The Getting Started with Dynamic Partition Manager task also provides an overview of DPM concepts, with links to additional reference information.
Integrated 3270 Console		HMC	Use the Integrated 3270 Console task to log on to a z/VM hypervisor that is running in a DPM partition.
Integrated ASCII Console		HMC	Use the Integrated ASCII Console task to log on to a Linux operating system.
Manage Adapters		HMC	Use the Manage Adapters task to view and customize the adapters and devices of an DPM-enabled system.
Adapter Details		HMC	Use the Adapter Details task to view or modify the adapter settings of the selected adapter.
Create Hipersockets Adapter		HMC	Use the Create Hipersockets Adapter task to create an adapter and switch for HyperSockets, which provide high-speed communications between partitions within a single system, without the need for any physical cabling or external networking connections.
Delete Hipersockets Adapter		HMC	Use the Delete Hipersockets Adapter task to delete only one selected HyperSockets adapter.
Export WWPNS		HMC	Use the Export WWPNS task to export the world wide port names (WWPNs) of the host bus adapters for one or more partitions.
Reassign Channel Path IDs		HMC	Use the Reassign Channel Path IDs task to change the channel path IDs that are assigned to DPM adapters.
Reassign Devices		HMC	Use the Reassign Devices task to change the adapter, port, or switch for one or more devices.
Manage Processor Sharing		HMC	Use the Manage Processor Sharing task to set processor weights, weight capping, and absolute capping for partitions with shared processors. You can also use this task to define one or more groups of partitions to set absolute capping limits.
Monitor tab under the Systems Management node			Use the Monitor tab to view the overall system monitoring data for the various components that make up a DPM-enabled system.
Monitor System Events		HMC	Use the Monitor System Events task to create and manage event monitors. An <i>event monitor</i> listens for events from managed objects, such as partitions, adapters, and other system resources or states. When an event is received, the monitor tests it with user defined time and text filters. If the event passes the tests, the monitor enables email to be sent to interested users.

Table 3. Summary of key HMC/SE tasks and displays for working with DPM-enabled systems, partitions, and adapters (continued)

Task name	Icon	On HMC or SE	Description
New Partition		HMC	Use the New Partition task to create a new partition on an DPM system. The New Partition task offers two modes through which you can create a partition: basic and advanced. For a comparison of the two modes and the implications of switching between them, see “Selecting which New Partition task mode to use” on page 35.
Operating System Messages		HMC	Use the Operating System Messages task as an operating system console while the operating system or hypervisor in a partition is being initialized. To do so, the operating system or hypervisor must support console integration, which is an HMC facility.
Partition Details		Both (view-only mode on SE)	Use the Partition Details task to view or modify an existing definition for a specific partition on an DPM system. Note that you cannot change the partition type through the Partition Details task.
Start		Both	Use the Start task to start a single DPM-enabled system, or to start one or more partitions on a DPM-enabled system.
Stop		Both	Use the Stop task to stop a single DPM-enabled system, or to stop one or more partitions on a DPM-enabled system.
System Details		Both (view-only mode on SE)	Use the System Details task to view and manage properties of the selected DPM-enabled system.

Part 4. Topics for system planners

Topics in this part provide the prerequisites for enabling DPM on a mainframe system, information about supported functions, and the engineering changes (ECs) or machine change levels (MCLs) for upgrading to the latest DPM version. This part also includes migration instructions and information about I/O adapter configuration. These topics are appropriate for experienced system planners and other administrators who are familiar with mainframe systems.

Topics covered in this part are:

- Chapter 10, “Prerequisites for using Dynamic Partition Manager,” on page 71
- Chapter 11, “Migrating DPM configuration data to a new system,” on page 79
- Chapter 13, “Adapter configuration,” on page 93

Chapter 10. Prerequisites for using Dynamic Partition Manager

This topic lists the IBM Z and IBM LinuxONE (LinuxONE) products that can run in IBM Dynamic Partition Manager (DPM) mode, lists prerequisites, and provides additional configuration details.

The following IBM mainframe systems can be configured to run in either standard Processor Resource/Systems Manager (PR/SM) mode or IBM Dynamic Partition Manager (DPM) mode. DPM uses PR/SM functions but presents a simplified user interface for creating partitions and managing system resources.

- IBM z14 (z14) (machine type 3906 or 3907)
- IBM z13 (z13) or IBM z13s (z13s)
- IBM LinuxONE Emperor II (Emperor II), or IBM LinuxONE Rockhopper II (Rockhopper II)
- IBM LinuxONE Emperor (Emperor) or IBM LinuxONE Rockhopper (Rockhopper) machine type 2965

The latest version of DPM is available starting with the Hardware Management Console / Support Element (HMC/SE) Version 2.14.0. For additional details, see “DPM versioning.”

For information about using prior versions of DPM on other systems, see *IBM Dynamic Partition Manager Guide*, SB10-7168.

Required feature codes

To run one of these systems in DPM mode, your company must order the system with the features listed in Table 4.

Table 4. Required feature codes for IBM Dynamic Partition Manager (DPM)

Feature code	Description
0016	Hardware for IBM Dynamic Partition Manager (DPM) feature
0417	OSA-Express 1000BASE-T Ethernet features. Two features are required for primary and backup connectivity.

DPM versioning

DPM is available through specific engineering changes (EC) or machine change levels (MCL) for the mainframe systems that support it. Table 5 on page 72 lists the DPM versions and the EC MCLs that must be applied to use any new functions or updates. The machine types indicate the systems on which you can apply the EC MCLs.

Important:

- Before applying an MCL, upgrade the HMC to the latest HMC/SE version and level.
- Note that each DPM version requires the application of both an HMC MCL and an SE MCL. The suggested practice is to apply MCLs to the HMC first, then to apply the SE MCL. Until the SE MCL is applied, new functions or updates are not available even if you have applied the HMC MCL.

Table 5. Summary of DPM releases and their associated EC MCLs

DPM release	Machine type	HMC/SE version and required MCLs	Description
Release 3.2 (R3.2)	3906 and 3907	HMC/SE Version 2.14.1	Enhancements to storage management through the Configure Storage task and its subtasks. Connect System to Storage Search and cloning capabilities, as well as automatic configuration functions, have been added or improved for specific storage elements, such as adapter ports and switches. Request Storage or Create Storage Group To make storage requests even easier to complete, administrators can create a master copy, or template, for a FICON or FCP storage group. Templates can reflect typical usage patterns (production, staging, test); standardize storage for specific user groups; or document company requirements or restrictions that might be in place for storage use. Users can select an available template and, with minimal changes, quickly submit a request for a new storage group. Configure FICON Connections - Administrators can dynamically reconfigure the logical control unit (LCU) paths that are mapped for a storage group, while the partitions to which the storage group is attached are active and running workloads. - Search and cloning capabilities, as described for Connect System to Storage. Storage Overview and Storage Group details - Administrators can use the new FICON Connection Report to view the logical paths that are in use for a FICON storage group. Administrators can also download the report contents in spreadsheet format. - The behavior of the background check for the FCP Connection Report has changed. Previously, DPM started a background check of the connections for an FCP storage group as soon as the user opened the Connection Report page. Now administrators can use the the Update report icon (🔄) to request a new report. - Administrators can use the new HISTORY tab on the Storage Group details page to view a list of actions that users have taken for a specific storage group. The ACTION column entry briefly describes the activity and preserves additional details that are available through a link. When the storage group is attached to one or more partitions, some changes that are made to the storage group can require corresponding changes on the operating system that the partition hosts. In this case, use the link in the ACTION column entry to display a list of required updates.

Table 5. Summary of DPM releases and their associated EC MCLs (continued)

DPM release	Machine type	HMC/SE version and required MCLs	Description
Release 3.1 (R3.1)	3906 and 3907	HMC/SE Version 2.14.0 with MCLs in the following bundles: • H33 or later • S55 and S55a, or later	Release of DPM for the z14 (machine type 3907) and Rockhopper II systems. Enhancements include the following items, which also can be applied to the z14 (machine type 3906) and Emperor II systems. • A new HMC task, Configure Storage, which provides a simplified, visual, and automated user interface for configuring and managing storage resources, without requiring extensive knowledge of the mainframe. Through this new task: – Administrators can connect a single system to storage devices through the HMC, rather than using either the Hardware Configuration Definition (HCD) or the Input/Output Configuration Program (IOCP) to define storage devices in an input/output configuration data set (IOCDs) for the system. – Administrators also use this task to request and fulfill storage resources for use by partitions on that system. – For automation purposes, the functions that are available through the Configure Storage task are also available through the HMC Web Services Application Programming Interface (API). For more information, see Chapter 17, “Managing storage on a DPM-enabled system,” on page 113. • Support for Fibre Channel connections (FICON) cards that operate in FICON native mode (FC or FICON). FICON native (FC or FICON) mode provides access to extended count key data (ECKD) devices, and tape devices, through point-to-point (direct) connections, or single- or multiple-channel switches. ECKD devices are more commonly known as direct-access storage devices (DASD). DPM supports access to ECKD devices, but does not support FICON channel-to-channel (CTC) or FICON tape devices. Also, DPM supports the use of parallel access volumes, but only through the optional HyperPAV feature on the IBM System Storage DS8000 series. • The Import DPM Configuration application programming interface (API), which automates most of the changes required to migrate the DPM configuration from an existing system to a new system that replaces it. This migration support also requires the application of specific MCLs on the existing system, as well as the new system. For more information, see “Migration instructions for a system replacement” on page 82.
Release 3.0 (R3.0)	3906	HMC/SE Version 2.14.0 with bundle S52 or later	Support for using Dynamic Host Configuration Protocol (DHCP) option 93 to tailor the processing related to a network boot loader request from an IBM Z system.
Release 3.0 (R3.0)	3906	HMC/SE Version 2.14.0 with bundle H13 or later	Support for using the Integrated ASCII Console through a new DPM API: Get ASCII Console Websocket URI. This API enables an API client to communicate with the operating system on an active partition. For more information, see <i>Hardware Management Console Web Services API</i> , SC27-2636.

Table 5. Summary of DPM releases and their associated EC MCLs (continued)

DPM release	Machine type	HMC/SE version and required MCLs	Description
Release 3.0 (R3.0)	3906	HMC/SE Version 2.14.0	Release of DPM for the z14 and Emperor II systems. Enhancements include the following items. • Support for the auto-configuration of Fibre Channel connections (FICON) in Fibre Channel Protocol (FCP) mode, IBM HiperSockets, and Open Systems Adapter-Express (OSA-Express) devices. This support not only simplifies the installation of the Linux operating system in a partition, but also enables dynamic updates to those virtual resources. To benefit from auto-configuration and dynamic updates, you need to use a version of the Linux operating system that supports these functions. See the Red Hat, SUSE, or Ubuntu product information page to determine which RHEL, SLES, or Ubuntu Server version provides this support. • For Secure Service Container partitions, several enhancements related to the configuration and use of network interface cards (NICs): – Specification of adapter port 1 (as well as port 0) – Specification of a default IPv6 gateway servers for NICs with a static IPv6 type address For more information, see the online help for the Network section of the New Partition or Partition Details tasks. • Support for additional file transfer protocols for booting and installing an operating system in a partition: FTP Secure (FTPS) and Secure File Transfer Protocol (SFTP). See Chapter 15, “Securing FTP operations,” on page 105. • Support for optional network traffic controls: – For partitions with a type of Linux or z/VM, support for the use of VLAN IDs. – For partitions of any type, support for unique, user-specified or firmware-generated MAC addresses. To display these MAC addresses on the operating system, you need to use a version of the Linux operating system that supports the specification or generation of MAC addresses through firmware. See the Red Hat, SUSE, or Ubuntu product information page to determine which RHEL, SLES, or Ubuntu Server version provides this support. (On an operating system without this support, the operating system's network device driver overwrites any MAC addresses with addresses that it generates.) For more information, see Chapter 16, “Using VLAN IDs to secure network communications,” on page 107. • Support for migrating existing DPM configuration data from earlier systems. See Chapter 11, “Migrating DPM configuration data to a new system,” on page 79.

Table 5. Summary of DPM releases and their associated EC MCLs (continued)

DPM release	Machine type	HMC/SE version and required MCLs	Description
Release 2.1 (R2.1)	2964 or 2965	HMC / SE Version 2.13.1 with the following MCLs: - HMC-SYSTEM P08462.261 and SE-SYSTEM P00339.304 or later - To use the support for IBM Secure Service Container partitions, also install SE-FWPART P08442.024 or later	- Support for displaying the number of processor threads through the Partition Details task. - Support for assigning a partition type of z/VM, for partitions in which the z/VM operating system will run. If you have already installed and run z/VM on DPM R2.0 or DPM R1, see the z/VM migration information in <i>IBM Dynamic Partition Manager Guide</i>, SB10-7168. - Support for creating, starting, and managing Secure Service Container partitions. A Secure Service Container partition contains its own embedded operating system, security mechanisms, and other features that are specifically designed for simplifying the installation of appliances, and for securely hosting them. For information about working with Secure Service Container partitions and the appliances they support, see <i>Secure Service Container User's Guide</i>, SC28-6971, which is available on http://www.ibm.com/servers/resourcelink.
Release 2.0 (R2.0)	2964 or 2965	HMC / SE Version 2.13.1 with the following MCLs: - SE-SYSTEM P00339.243 or later - SE-I390ML P08414.044 or later - SE-ENABLE2 P08434.001 or later - SE-MCS P08456.019 or later	- Support for z/VM as a hypervisor, and for the use of the Integrated 3270 Console task through which users can interact with the z/VM hypervisor. For additional details, see "Requirements for the hypervisor or operating system" on page 76. - Support for universally unique identifiers (UUIDs), which enables users to define their own device numbers for Peripheral Component Interconnect Express (PCIe) devices, when using a Linux distribution that also supports the use of UUIDs. For related information, see the migration considerations topic in <i>IBM Dynamic Partition Manager Guide</i>, SB10-7168.
Release 1 (R1)	2964 or 2965	HMC / SE Version 2.13.1	Initial release of DPM for the z13, z13s, Emperor, and Rockhopper (machine type 2965) systems.

About network adapters

Several types of network adapters enable communication through different networking transport protocols. For a DPM-enabled system, these network adapters include the following:

- Open Systems Adapter-Express (OSA-Express) features, which provide direct, industry-standard Ethernet LAN connectivity through various operational modes and protocols. OSA adapters can provide connectivity between partitions on the same system, as well as connectivity to external LANs.
- HiperSockets, which provide high-speed communications between partitions within a single system, without the need for any physical cabling or external networking connections.
- Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) Express features. These features provide similar benefits for system-to-system communication as the HiperSockets technology provides for communication between partitions on the same system.

The supported OSA and RoCE adapters vary, depending on the system. For a list of the network adapters that are supported on a specific system, see the appropriate system overview on the IBM Redbooks web site at <http://www.redbooks.ibm.com/>.

OSA and RoCE adapters are defined through the use of the Hardware Configuration Definition (HCD) or the Input/Output Configuration Program (IOCP) to define storage devices. To create HiperSockets on a

DPM-enabled system, use the **Create HiperSockets Adapter** task, which is available through the **Actions** list on the **Adapters** tab of the **Manage Adapters** task.

About storage adapters

IBM Z and IBM LinuxONE (LinuxONE) systems use Fibre Channel connections (FICON) to provide high-speed data transfer to storage devices. Through cables, FICON Express features (adapter cards) connect a DPM-enabled system to the devices in the storage area network (SAN). The supported FICON Express adapter cards vary, depending on the system configuration; for example, the z14 supports the FICON Express16S+, FICON Express16S, and FICON Express8S adapter cards. For a list of the storage adapters that are supported on a specific system, see the appropriate system overview on the IBM Redbooks web site at <http://www.redbooks.ibm.com/>.

FICON adapter cards operate in different modes, which determine the type of storage devices that you can access. Typically, storage administrators configure the mode in which each FICON adapter card operates.

- Fibre Channel Protocol (FCP) mode provides access to Small Computer System Interface (SCSI) devices, through single- or multiple-channel switches. Note that, because DPM runs all FCP adapters in N Port Identifier Virtualization (NPIV) mode, the SAN must support NPIV. DPM senses but does not display FCP tape devices, so administrators cannot manage them through DPM tasks on the HMC; however, operating systems that are hosted by partitions on a DPM-enabled system can access those tape devices. Support for FCP mode is available with all DPM releases.
- FICON native (FC or FICON) mode provides access to extended count key data (ECKD) devices, and tape devices, through point-to-point (direct) connections, or single- or multiple-channel switches. ECKD devices are more commonly known as direct-access storage devices (DASD). DPM supports access to ECKD devices, but does not support FICON channel-to-channel (CTC) or FICON tape devices. Also, DPM supports the use of parallel access volumes, but only through the optional HyperPAV feature on the IBM System Storage DS8000 series. Support for FICON mode is available only with DPM R3.1 or later.

Starting with DPM R3.1, the Hardware Management Console (HMC) **Configure Storage** task replaces the use of either the Hardware Configuration Definition (HCD) or the Input/Output Configuration Program (IOCP) to define storage devices in an input/output configuration data set (IOCDs) for the system.

Requirements for the hypervisor or operating system

Partitions on a DPM-enabled system support the following operating systems and hypervisors:

- The following minimum distribution levels of Linux on Z:
 - Red Hat Enterprise Linux (RHEL) 6 or later.
 - SUSE Linux Enterprise Server (SLES) 11 or later.
 - Ubuntu Server (KVM or LPAR DPM): 16.04 LTS or later.

For recommended Linux on Z distribution levels on IBM Z and LinuxONE servers, see the IBM tested platforms at: <http://www.ibm.com/systems/z/os/linux/resources/testedplatforms.html>

For information about installing and running a Linux distribution on an IBM Z or LinuxONE server, see the Linux on IBM Systems topics in IBM Knowledge Center, at: https://www.ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz_r_lib.html

- z/VM V6.4 or later, with the following limitations:
 - On a DPM-enabled system, z/VM can host only Linux on Z guests.
 - Because DPM does not support FICON channel-to-channel (CTC) connections, administrators cannot configure and use z/VM Single System Image (SSI) and Live Guest Relocation (LGR).
 - Because DPM exclusively provides dynamic input/output (I/O), administrators cannot use z/VM to run Hardware Configuration Definition (HCD) and Hardware Configuration Manager (HCM).

- To boot and install z/VM in a partition, administrators must use one of the following HMC **New Partition** or **Partition Details** boot options and z/VM installation methods.
 - Select the **Hardware Management Console removable media** boot option for the partition, and then use the **DVDPRIME** command with the **dvd** option to install z/VM. Note that, when you use this option, installation can take up to several hours to complete.
 - Select one of the FTP boot options for the partition, and then use the **DVDPRIME** command with the **server** option to install z/VM through an FTP server.

For more information about partition boot options and the z/VM installation process, see Chapter 5, “Operating systems and hypervisors,” on page 23. For information about z/VM, see the product documentation. The z/VM library is available at <http://www.vm.ibm.com/library/>

Where to find hardware planning and corequisite software information

For the most recent hardware planning and corequisite software information, go to IBM Resource Link: <http://www.ibm.com/servers/resourceLink>

- For hardware updates, click **Tools** on the navigation panel. Then click **Machine information** under **Servers**, and enter your enterprise number, customer number, or machine serial number for the host system (CPC). You must register with IBM to search machine information.
- For software updates, click **Fixes** on the navigation panel. Then click **Preventative Service Planning buckets (PSP)** under **Preventive actions**, and check the PSP bucket for the appropriate system::
 - For a z14, the 3906DEVICE or 3907DEVICE PSP bucket
 - For an Emperor II, the 3906DEVICE PSP bucket
 - For a Rockhopper II, the 3907DEVICE PSP bucket
 - For a z13 or Emperor, the 2964DEVICE PSP bucket
 - For a z13s or Rockhopper (machine type 2965), the 2965DEVICE PSP bucket

Chapter 11. Migrating DPM configuration data to a new system

This topic describes the migration considerations or actions that you need to take when you migrate your IBM Dynamic Partition Manager (DPM) environment from an existing system to a new system.

If you configured partitions and adapters on an existing IBM Z or IBM LinuxONE (LinuxONE) system with DPM enabled, you can migrate the configuration information for those resources to a new DPM-enabled system.

The tasks that you need to complete vary, depending on whether your company is acquiring the new system through a miscellaneous equipment specification (MES) frame upgrade, or through a complete system replacement.

MES frame upgrade

When you acquire a new system through an MES frame upgrade, you are getting a new system that reuses the serial number, I/O worldwide port number (WWPN), and actual hardware parts (such as I/O cages, adapters, and so on) from an existing system.

Through the MES frame upgrade, most reconfiguration is automated through the system installation and start-up process itself, in part by DPM. However, you do need to complete several post-installation checks, including the modification of some partition definitions. For instructions, see “Migration instructions for an MES frame upgrade.”

System replacement

In this case, you are getting a completely new system to replace an existing system. This new system has its own serial number and, in most cases, the I/O WWPN is also different from the WWPN of the prior system.

When you replace an existing system with a new system, you need to write a script that uses the Hardware Management Console (HMC) Web Services API to export the DPM configuration from the existing system, and to import the configuration to the new system. You can use any programming language or tool that you currently use to issue the HMC Web Services API. Note that only specific migration paths are supported, and specific machine change levels (MCLs) must be applied to the new system. The partitions on a newly migrated system will have access to storage area network (SAN) resources through storage groups, so specific zoning and masking requirements must be met before starting the migration process. For instructions, see “Migration instructions for a system replacement” on page 82.

Migration instructions for an MES frame upgrade

When you acquire a new system through a miscellaneous equipment specification (MES) frame upgrade, your IBM service representative copies configuration data from the existing system to the new system, and starts the Support Element (SE) on the new system. During the start process, DPM automatically configures new adapters and reconfigures existing adapters and partitions to fit the new system. However, because some information is not available on the new system, DPM cannot completely reconfigure all objects that are carried forward from the existing system. Follow this set of instructions to complete the reconfiguration of your DPM on the new system.

Before you begin

- Before working with your service representative on the MES upgrade, make sure that you apply the latest available firmware for DPM on the existing system. For more information about DPM firmware updates, see the following topics:
 - “DPM versioning” on page 71

- “Where to find hardware planning and corequisite software information” on page 77
- Only the following migration paths are supported for an MES frame upgrade.
 - From a z13 to a z14 (machine type 3906)
 - From an Emperor to an Emperor II
 - From a z13s to a z14 (machine type 3907)
 - From a Rockhopper (machine type 2965) to a Rockhopper II

About this task

As part of the MES frame upgrade process, your service representative saves upgrade data that resides on the SE of the existing system; that data is restored on the new hardware when the SE for the new system is first started. After the data is restored, DPM automatically updates its configuration data by detecting new adapters, updating existing adapter properties, updating partition definitions, and so on. For example, if you reuse a specific adapter on the new system but install it in a different I/O cage, DPM updates the adapter properties to reflect the new location.

The updates that DPM makes are listed in the hardware messages for the new system. This message includes several lists:

- A list of partitions that were defined having reserved resources on the prior system. Because DPM cannot guarantee that the same resources are available on the new system, DPM deletes the reserved property from the partition definition.
- A list of partitions that were defined with the boot option of ISO image. ISO images are not part of the configuration data that can be copied over to a new system, so DPM changes the partition definition to specify a boot option of None.
- A list of the physical channel IDs (PCHIDs) that DPM converted for adapters that moved to a new location. This list maps the prior PCHID value and the new PCHID value for moved or replaced adapters.

Use this message to help you complete the post-installation steps in this procedure. Figure 37 shows a sample hardware message for the automated DPM process.




Figure 37. Sample hardware message that shows the results of a successful DPM conversion of migrated configuration information

Procedure

1. If you are using new OSA-Express 1000BASE-T Ethernet adapters (required feature 0417) on the new system, configure these new system management (OSM) adapters. On the SE for the new system, using the default SERVICE user ID or a user ID that is assigned to the SERVICE task role, complete the following steps.
 - a. Open the **System Details** task.
 - b. Select the **Management Networks** tab.
 - c. Choose a new adapter from both of the OSM drop-down lists.
 - d. Click **Apply** to save your changes.
2. On an HMC through which you can manage the new system, open the Hardware Messages to view the updates that DPM has made.
 - a. In the **System Management** view, select the new system.
 - b. From the pop-up menu next to the system name, select **Daily**, then select **Hardware Messages** to display hardware messages for this system.
 - c. If necessary, scroll through the list of messages to the date when the system was first started, or look for the message text: The migration of the partitions configuration to a new version was successful.
 - d. Select the message and click **Details** to display the results of the automated process. Keep this tab open for reference as you complete the remaining steps in this procedure.
3. Review and, if necessary, modify the migrated adapters. For any adapter with a status value of Exceptions, evaluate the exceptions and determine what action might be required.
 - a. From the **Systems** tab in the **System Management** view, open the pop-up menu next to the system name, select **Configuration**, then select **Manage Adapters**.
 - b. On the **Adapters** tab, sort the **Status** column to show all adapters for which the value Exceptions is displayed, and review those adapters to determine what action to take.
 If the Number of Partitions column indicates that one or more partitions have devices that are backed by an adapter that was removed (that is, they are not installed on the new system), you need to reassign those devices to a different adapter. Depending on the adapter type, these devices can be network interface cards (NICs), host bus adapters (HBAs), or virtual functions (VFs). To reassign devices, complete the following steps.
 - 1) Click on the **Devices** tab and sort the table by adapter name, and select the devices that are assigned to an adapter that was removed.
 - 2) Select **Reassign Devices** from the **Actions** list.
 - 3) On the Reassign Devices window, select the new backing adapter for the devices, and click **Reassign**.
 - 4) Repeat these steps as necessary, until all devices are assigned to an adapter that is currently installed on the new system.

When the SE for the new system is rebooted, the adapters that were removed are deleted from the DPM configuration.

4. Review and modify the migrated partitions that are listed in the hardware message.
 - a. From the **System Management** view, click the **Partitions** tab to display the partitions for the new system.
 - b. If necessary, sort or scroll the list to find the partition names that are listed in the hardware message.
 - c. Select one of the listed partitions, open the pop-up menu next to the partition name, and select **Partition Details**.
 - For a partition that was reserved on the prior system, complete the following steps.

- 1) Select or scroll to the **Processors** section, and adjust the processor type, mode, and number of virtual processors as necessary. Review the Processors bar chart to determine how many processors are available on this system, and how many are already in use or reserved for other partitions.
- 2) Select or scroll to the **Memory** section, and adjust the initial and maximum amounts of memory that are assigned to the partition. Review the Installed Memory bar chart to determine how much memory is available on this system, and how much is already in use or reserved for other partitions.
- 3) Select or scroll to the **General** section, and select the **Reserve resources...** check box.
- For a partition that had a boot option of ISO image, complete the following steps.
 - 1) Select or scroll to the **Boot** section, and select **ISO image** from the **Boot from** list.
 - 2) Upload an ISO file that is located on your workstation file system.
- d. When you finish updating the partition definition, click **Apply** to save your changes.
- e. Repeat these steps, as necessary, for each migrated partition that is listed in the hardware message.

Results

The DPM configuration on the new system is complete.

What to do next

You can start the partitions on the new system.

Migration instructions for a system replacement

When you replace an existing system with a new system, you need to write a script that uses the Hardware Management Console (HMC) Web Services API to export the DPM configuration from the existing system, and to import the configuration to the new system. You can use any programming language or tool that you currently use to issue the HMC Web Services API. Note that only specific migration paths are supported, and specific machine change levels (MCLs) must be applied to the new system. The partitions on a newly migrated system will have access to storage area network (SAN) resources through storage groups, so specific zoning and masking requirements must be met before starting the migration process.

Before you begin

- Review Table 6, which lists the machine types and DPM releases that are supported for migrating a DPM configuration from an existing (source) system to a new (target) system.
 - Both systems must have DPM enabled.
 - The target system must have DPM R3.1 installed. DPM R3.1 requires the installation of an enablement MCL that your service representative can supply.

Table 6. Supported migration paths for exporting and importing a DPM configuration

Source system	Target system
One of the following systems, with DPM R2.1 or earlier: • z13 (machine type 2964) or z13s (machine type 2965) • Emperor (machine type 2964) or Rockhopper (machine type 2965)	One of the following systems, with DPM R3.1: • z14 (machine type 3906 or 3907) • Emperor II (machine type 3906) or Rockhopper II (machine type 3907)
One of the following systems, with DPM R3.0: • z14 (machine type 3906 or 3907) • Emperor II (machine type 3906) or Rockhopper II (machine type 3907)	One of the following systems, with DPM R3.1: • z14 (machine type 3906 or 3907) • Emperor II (machine type 3906) or Rockhopper II (machine type 3907)

- During the migration process, the host bus adapters (HBAs) that a partition uses are converted into a dedicated FCP storage group. A storage group is a logical group of storage volumes that share certain attributes; a storage group provides access to storage resources starting with DPM R3.1. To ensure that this conversion is successful, make sure that the storage administrator has complied with the following SAN zoning and masking requirements:
 - All world wide port names (WWPNs) that are used by partitions must be added to the zoning of all switches.
 - All WWPNs that are used by partitions must be added to the host-mapping of the storage subsystems that provide the logical unit numbers (LUNs) for the storage group.
- Make sure that you have enabled the use of the HMC Web Services API on both the source and the target system. You can enable the use of the API through the HMC **Customize API Settings** task. For information about the authorization requirements and syntax of specific APIs, see the appropriate edition of *Hardware Management Console Web Services API*, which is available on IBM Resource Link.
- To perform this procedure, choose a user ID that has authorization to access all of the DPM configuration data that you want to export from the existing system. If you perform the steps in this procedure with a user ID that has access to only specific DPM partitions or adapters, only the configuration data for those resources is exported.
 - The suggested choice is the default SERVICE user ID. API access is not enabled for the SERVICE user ID by default, so you must authorize API access for the SERVICE user ID through the HMC **User Management** task.
 - If you choose a user ID other than the default SERVICE user ID, make sure that user ID has the following permissions.
 - API access
 - Object permission to both the source and target systems
 - Task permission to **Import Dynamic Partition Manager Configuration**.
- If you have controlled user access to DPM resources, and want to transfer those authorizations, use the **Save/Restore Customizable Console Data** task to save the data from the source system and restore it on the target system.
- Download the following scripts:
 - **zhmcclient**, which is an Open Source client library that interacts with the HMC Web Services API. For more information and installation details, see <https://github.com/zhmcclient>.
 - **exportDpmResourcesToFile.py**, which is a Python script that exports the DPM configuration of a system to a configuration file. For more information, see the accompanying PDF on the download page in Resource Link: <https://www.ibm.com/servers/resource link/lib03020.nsf/pages/zHmcWebServicesExportDpmConfig?OpenDocument>

About this task

Through this procedure, you are exporting DPM configuration data that includes the objects and properties of the DPM-enabled source system, its adapters, and the auto-start partition list and capacity groups, if any. The configuration data also includes the objects and properties of partitions and their virtual adapters: network interface cards (NICs), host bus adapters (HBAs), virtual functions, and cryptos.

When this configuration data is imported to a target system with DPM R3.1, most of the data is preserved but adapter IDs are changed and HBAs become part of storage groups. You need to review these changes and modify them, if necessary. Figure 38 on page 84 illustrates how partitions access storage resources both before and after migration.



Before migration**After migration**

Figure 38. Partition access to storage before and after migration

1. On the source system, an administrator defines storage resources for Partition A by creating four HBAs, each with its own backing FCP-mode adapter. When the partition definition is saved, DPM generates WWPNs that the administrator exports and gives to the storage administrator to complete zoning and LUN masking tasks for the storage subsystem configuration.
2. The backing adapters are configured on the source system and are physically connected to a storage subsystem through a switch. In this example, adapter 0124 is connected to Controller 1 through Switch A in Fabric 0. Adapter 0118 is also connected to Controller 1 through Switch A in Fabric 0. Similarly, adapters 01CC and 01D8 are connected to Controller 2 through Switch B in Fabric 1. Partition A uses the WWPNs to access specific volumes on each storage subsystem. This configuration gives Partition A access to eight storage volumes.
3. During the migration process, DPM creates a dedicated (not shared) FCP storage group, SG01, and attaches it to Partition A. The HBAs and WWPNs that were defined for Partition A on the source system become part of the infrastructure for the storage group. Through the storage group, Partition A on the target system has access to the same eight volumes.
4. Note that the adapter IDs that were in use on the source system are copied to the target system. You can change the adapter IDs by providing an adapter mapping for the migration process (as described in step 4 on page 85 in the following procedure).
5. Note that, on a DPM R3.1 system, switches must be connected to all storage controllers that provide LUNs for a storage group. Because storage group SG01 uses volumes on both Controller 1 (LUNs 1-4) and Controller 2 (LUNs 5-8), Switch A must be connected to both Controller 1 and Controller 2, and Switch B also must be connected to both controllers.

- For more information about the results of the migration process and actions that you might need to take, see Results and What to do next.
- For more information about DPM R3.1 storage management, see Chapter 17, “Managing storage on a DPM-enabled system,” on page 113.

Procedure

1. For your script, install the latest version of the **zhmcclient** package.
2. Using the central processor complex (CPC) object identifier for the source system, run the **exportDpmResourcesToFile.py** script to retrieve the existing DPM configuration.
The script generates a configuration file that you use as input for the **Import DPM Configuration** API. The generated file name is *cpcname.dpm.resources*, where *cpcname* is the name of the source system.
3. On the HMC through which you are managing the target system, check the system status. To successfully import the DPM configuration, the target system must be in one of the following states: Active, Service Required, Degraded, or Exceptions.
4. Use the **Import DPM Configuration** API to import the DPM configuration to the new system.

The following information provides the syntax and high-level instructions for using the API. For the best results, see the description of the **Import DPM Configuration** API in *Hardware Management Console Web Services API*, SC27-2636, which is available on IBM Resource Link.

The syntax for the **Import DPM Configuration** API is:

POST /api/cpcs/{cpc-id}/operations/import-dpm-config

- Specify the CPC object identifier for the target system.
- Optionally provide an adapter mapping, using the following format as an example.

```
{ "adapter-mapping":
  [
    { "old-adapter-id" : "100", "new-adapter-id" : "124" },
    { "old-adapter-id" : "110", "new-adapter-id" : "134" }
  ]
}
```

If you do not provide a mapping, DPM uses a one-to-one mapping of adapters on the source system to adapters on the target system.

If the request is successful, the API returns HTTP status code 204. Otherwise, use the HTTP status code and reason codes for troubleshooting.

5. Restart the target system.

Results

The DPM configuration has been imported to the target system.

- For each partition on the source system that has HBAs assigned for access to storage, DPM creates a dedicated FCP storage group and reuses the HBAs and WWPNs for that storage group. The new storage group is attached to the same partition on the target system.
- If the partition definition specified an HBA, port, and logical unit number (LUN) for the **Storage (SAN)** boot option, the new storage group contains a storage volume that is mapped to the same port and LUN. If the partition does not use a boot volume, the storage group that DPM creates does not contain any storage volumes.
- For each new storage group, DPM sets the Connectivity attribute to the number of distinct adapters that were in use by the HBAs for the partition on the source system. For example, the Connectivity attribute for storage group SG01 in Figure 38 on page 84 is set to 4 because of the four adapters that were selected to back the HBAs.
- DPM starts automatic discovery of the LUNs that are configured for the WWPNs assigned to each storage group.

- If no LUNs are discovered (that is, no zoning and masking has been done), DPM sets the fulfillment state of the storage group to Checking migration. For any storage groups in this fulfillment state, DPM attempts LUN discovery every 10 minutes.
- If a discovery attempt results in the same set of LUNs that are sensed (detected) for all WWPNS over the required number of adapters (as defined by the Connectivity attribute for the storage group), DPM automatically accepts the volumes and sets the fulfillment state of the storage group to Complete. For any storage groups in this fulfillment state, DPM attempts LUN discovery every 24 hours.
- If a discovery attempt results in LUNs that are sensed for all WWPNS, but not all WWPNS are configured for the same set of LUNs over the required number of adapters (as defined by the Connectivity attribute for the storage group), DPM sets the fulfillment state of the storage group to Pending with mismatches. For any storage groups in this fulfillment state, DPM attempts LUN discovery every 10 minutes.

What to do next

- Go to **Storage Overview** in the **Configure Storage** task to check storage groups and adapter assignments. If necessary, you can modify the storage group names and other details, or reconfigure system adapters through **Storage Cards**.

The Storage Overview page includes the Storage Groups table, which contains one row for each storage group. To view more details about a specific storage group, click anywhere in the table row for that storage group to open the Storage Group details page.

Check the fulfillment state to determine what actions you might need to take.

- If the fulfillment state does not change from Checking migration to either Complete or Pending with mismatches, have the storage administrator fix the configuration in the storage subsystem. Then open the **Connection Report** on the Storage Details page for the storage group, and select the **Update report** icon (🔄) so that DPM rechecks the storage group connections and changes the fulfillment state.
- For a fulfillment state of Pending with mismatches, go to the **Volumes** tab on the Storage Details page for the storage group. All mismatched volumes are displayed at the top of the Volumes table, and are marked with a warning icon (⚠️). Volumes are considered mismatched when one or more of the following conditions are true.
 - The volumes are zoned and masked equally for all of the WWPNS. For these volumes, system administrators have the option of either deleting or keeping the mismatched volumes in the storage group. They can select one or more volumes to keep or delete.
 - The volumes are masked and zoned only for a subset of WWPNS. For these volumes, storage administrators must correct the zoning and masking configurations for the WWPNS.

For more information about fulfillment states and possible corrective actions, see “All of the details: the Storage Group details page” on page 126 (this information is also available in the online help for the Storage Group details page).

- If any partitions have multiple HBAs that are backed by the same adapter, consider modifying the storage group to make sure that the required number of HBAs are retained if the storage group is detached and then reattached to the same or another partition. For example, suppose that your source system has one partition with five HBAs that are backed by only two adapters.
 1. During the migration process, DPM creates the dedicated FCP storage group, retains all five HBAs that are defined for the partition, and sets the Connectivity attribute of the storage group to 2, which is the number of distinct adapters that were in use by the five HBAs.
 2. After the migration, an administrator detaches the storage group from the partition.
 3. If an administrator then reattaches the storage group to the same partition, or attaches it to a different partition, DPM uses only the current Connectivity attribute setting to determine how many HBAs to create when the storage group is reattached. In this case, DPM creates only two HBAs, not five.

If the partition needs additional HBAs, you need to modify the storage group attributes to either change the Connectivity attribute setting, or specify additional connections. To do so, go to **Storage Overview**, select the storage group, and select **Modify Storage Group**.

- If you have imported configuration data for any Secure Service Container partitions, use the **Partition Details** task to check the following information.
 - Go to the **General** section to reset the default master user ID, the password, or both.
 - Go to the **Boot** section to check the install location for the appliance that the partition hosts. If necessary, reinstall the appliance.

Important:

- After you have successfully migrated to DPM R3.1, use the **Backup Critical Data** task on the Support Element.
- When you detach an FCP storage group from a partition, DPM does not preserve the HBAs, device numbers, or backing adapters that were in use for that storage group. Consequently, if you reattach the storage group to the same partition, the device numbers and the backing adapters of the HBAs are not guaranteed to be the same as they were before the detachment. The same condition is true if you attach the FCP storage group to a new partition: the device numbers and the backing adapters of the HBAs are not guaranteed to be the same as they were when the storage group was attached to a different partition.

If the operating system image for the partition resides on a boot volume in the storage group that you reattach, the operating system might not start when the partition is restarted. (Operating systems that are started from a storage volume usually have a preconfigured device number and a path to the volume.) To avoid this situation, administrators must review the device numbers and the backing adapters of the HBAs, after reattaching the storage group and before restarting the partition. If necessary, the administrators must change the device number to match the preconfigured device number for the operating system, and make sure that the preconfigured device number is assigned to the backing adapter that was assigned to the HBA before the detachment. For more information, see “Instructions: Detaching and reattaching an FCP storage group” on page 61.

Chapter 12. Upgrading from DPM R3.0 to R3.1

When you upgrade your system from DPM R3.0 to R3.1, partitions on the upgraded system will have access to storage area network (SAN) resources through storage groups, rather than through manually created host bus adapters (HBAs). A storage group is a logical group of storage volumes that share certain attributes, such as type (FCP or FICON) and shareability. During the upgrade process, which is concurrent and does not affect any active partitions with running workloads, DPM converts the HBAs that are defined for a partition into a dedicated FCP storage group. To ensure that this conversion is successful, administrators need to prepare for the upgrade, check the results of the conversion, and make sure that the newly created storage groups are ready for use.

Before you begin

- Make sure that DPM is enabled on your system.
- Make sure that the storage administrator has complied with the following SAN zoning and masking requirements:
 - All world wide port names (WWPNs) that are used by partitions must be added to the zoning of all switches.
 - All WWPNs that are used by partitions must be added to the host-mapping of the storage subsystems that provide the logical unit numbers (LUNs) for the storage group.
- If your installation is using the HMC Web Services Application Programming Interfaces (APIs) to automate DPM tasks, you need to update your API scripts. For more information about new or changed APIs for DPM R3.1, see the appropriate edition of *Hardware Management Console Web Services API*, which is available on IBM Resource Link.
- The suggested practice is to work with your IBM service representative, who installs and activates the required MCLs. DPM R3.1 requires the installation of an enablement MCL that your service representative can supply.

About this task

DPM R3.1 introduces a simplified, visual, and automated user interface for configuring and managing storage resources, without requiring extensive knowledge of the mainframe. The **Configure Storage** task replaces the use of either the Hardware Configuration Definition (HCD) or the Input/Output Configuration Program (IOCP) to define storage devices in an input/output configuration data set (IOCDS) for the system. It also introduces the storage group, through which partitions access the storage resources that are connected to the system.

When you apply the MCLs to upgrade from DPM R3.0 to R3.1, DPM preserves the configuration data that includes the objects and properties of the DPM-enabled system, its adapters, and the auto-start partition list and capacity groups, if any. The configuration data also includes the objects and properties of partitions and their virtual adapters: network interface cards (NICs), host bus adapters (HBAs), virtual functions, and cryptos. HBAs, however, become part of storage groups, as shown in Figure 39 on page 90.



Before upgrade**After upgrade to DPM R3.1**

Figure 39. Partition access to storage before and after upgrading to DPM R3.1

1. On a DPM R3.0 system, an administrator defines storage resources for Partition A by creating four HBAs, each with its own backing FCP-mode adapter. When the partition definition is saved, DPM generates WWPNs that the administrator exports and gives to the storage administrator to complete zoning and LUN masking tasks for the storage subsystem configuration.
2. The backing adapters are configured on the system and are physically connected to a storage subsystem through a switch. In this example, adapter 0124 is connected to Controller 1 through Switch A in Fabric 0. Adapter 0118 is also connected to Controller 1 through Switch A in Fabric 0. Similarly, adapters 01CC and 01D8 are connected to Controller 2 through Switch B in Fabric 1. Partition A uses the WWPNs to access specific volumes on each storage subsystem. This configuration gives Partition A access to eight storage volumes.
3. During the upgrade process, DPM creates a dedicated (not shared) FCP storage group, SG01, and attaches it to Partition A. The HBAs and WWPNs that were defined for Partition A become part of the infrastructure for the storage group. Through the storage group, Partition A has access to the same eight volumes.
4. Note that, on a DPM R3.1 system, switches must be connected to all storage controllers that provide LUNs for a storage group. Because storage group SG01 uses volumes on both Controller 1 (LUNs 1-4) and Controller 2 (LUNs 5-8), Switch A must be connected to both Controller 1 and Controller 2, and Switch B also must be connected to both controllers.

For more information about DPM R3.1 storage management, see Chapter 17, “Managing storage on a DPM-enabled system,” on page 113.

Procedure

1. On the Hardware Management Console (HMC) through which you are managing the system, check that the system is in an operating state.
2. Have your service representative install and activate the required MCLs.
 - a. Install and activate the required MCLs on the HMC.
 - b. Install and activate the enablement MCL on the Support Element (SE).
 - c. Restart the SE.
 - d. Restart the HMC.

Results

The system has been successfully upgraded to DPM R3.1.

- For each partition that has HBAs assigned for access to storage, DPM creates a dedicated FCP storage group and reuses the HBAs and WWPNNs for that storage group. The new storage group is attached to the same partition.
- If the partition definition specified an HBA, port, and logical unit number (LUN) for the **Storage (SAN)** boot option, the new storage group contains a storage volume that is mapped to the same port and LUN. If the partition does not use a boot volume, the storage group that DPM creates does not contain any storage volumes.
- For each new storage group, DPM sets the Connectivity attribute to the number of distinct adapters that were in use by the HBAs for the partition. For example, the Connectivity attribute for storage group SG01 in Figure 39 on page 90 is set to 4 because of the four adapters that were selected to back the HBAs.
- DPM starts automatic discovery of the LUNs that are configured for the WWPNNs assigned to each storage group.
 - If no LUNs are discovered (that is, no zoning and masking has been done), DPM sets the fulfillment state of the storage group to Checking migration. For any storage groups in this fulfillment state, DPM attempts LUN discovery every 10 minutes.
 - If a discovery attempt results in the same set of LUNs that are sensed (detected) for all WWPNNs over the required number of adapters (as defined by the Connectivity attribute for the storage group), DPM automatically accepts the volumes and sets the fulfillment state of the storage group to Complete. For any storage groups in this fulfillment state, DPM attempts LUN discovery every 24 hours.
 - If a discovery attempt results in LUNs that are sensed for all WWPNNs, but not all WWPNNs are configured for the same set of LUNs over the required number of adapters (as defined by the Connectivity attribute for the storage group), DPM sets the fulfillment state of the storage group to Pending with mismatches. For any storage groups in this fulfillment state, DPM attempts LUN discovery every 10 minutes.

What to do next

- Go to **Storage Overview** in the **Configure Storage** task to check storage groups and adapter assignments. If necessary, you can modify the storage group names and other details, or reconfigure system adapters through **Storage Cards**.

The Storage Overview page includes the Storage Groups table, which contains one row for each storage group. To view more details about a specific storage group, click anywhere in the table row for that storage group to open the Storage Group details page.

Check the fulfillment state to determine what actions you might need to take.

- If the fulfillment state does not change from Checking migration to either Complete or Pending with mismatches, have the storage administrator fix the configuration in the storage subsystem. Then open the **Connection Report** on the Storage Details page for the storage group, and select the **Update report** icon (⌂) so that DPM rechecks the storage group connections and changes the fulfillment state.

- For a fulfillment state of Pending with mismatches, go to the **Volumes** tab on the Storage Details page for the storage group. All mismatched volumes are displayed at the top of the Volumes table, and are marked with a warning icon (⚠). Volumes are considered mismatched when one or more of the following conditions are true.
 - The volumes are zoned and masked equally for all of the WWPNS. For these volumes, system administrators have the option of either deleting or keeping the mismatched volumes in the storage group. They can select one or more volumes to keep or delete.
 - The volumes are masked and zoned only for a subset of WWPNS. For these volumes, storage administrators must correct the zoning and masking configurations for the WWPNS.

For more information about fulfillment states and possible corrective actions, see “All of the details: the Storage Group details page” on page 126 (this information is also available in the online help for the Storage Group details page).

- If any partitions have multiple HBAs that are backed by the same adapter, consider modifying the storage group to make sure that the required number of HBAs are retained if the storage group is detached and then reattached to the same or another partition. For example, suppose that your system has one partition with five HBAs that are backed by only two adapters.
 1. During the upgrade process, DPM creates the dedicated FCP storage group, retains all five HBAs that are defined for the partition, and sets the Connectivity attribute of the storage group to 2, which is the number of distinct adapters that were in use by the five HBAs.
 2. After the upgrade, an administrator detaches the storage group from the partition.
 3. If an administrator then reattaches the storage group to the same partition, or attaches it to a different partition, DPM uses only the current Connectivity attribute setting to determine how many HBAs to create when the storage group is reattached. In this case, DPM creates only two HBAs, not five.

If the partition needs additional HBAs, you need to modify the storage group attributes to either change the Connectivity attribute setting, or specify additional connections. To do so, go to **Storage Overview**, select the storage group, and select **Modify Storage Group**.

- If, at a later time, you need to deactivate the enablement MCL, you first need to detach all FICON storage groups from the partitions that are using them, and then delete the FICON storage groups from the system. After you deactivate the MCL, partitions can access FCP but not FICON storage devices. Deactivating the enablement MCL is a concurrent process that does not affect partitions with running workloads.

Important:

- After you have successfully upgraded to DPM R3.1, use the **Backup Critical Data** task on the Support Element.
- When you detach an FCP storage group from a partition, DPM does not preserve the HBAs, device numbers, or backing adapters that were in use for that storage group. Consequently, if you reattach the storage group to the same partition, the device numbers and the backing adapters of the HBAs are not guaranteed to be the same as they were before the detachment. The same condition is true if you attach the FCP storage group to a new partition: the device numbers and the backing adapters of the HBAs are not guaranteed to be the same as they were when the storage group was attached to a different partition.

If the operating system image for the partition resides on a boot volume in the storage group that you reattach, the operating system might not start when the partition is restarted. (Operating systems that are started from a storage volume usually have a preconfigured device number and a path to the volume.) To avoid this situation, administrators must review the device numbers and the backing adapters of the HBAs, after reattaching the storage group and before restarting the partition. If necessary, the administrators must change the device number to match the preconfigured device number for the operating system, and make sure that the preconfigured device number is assigned to the backing adapter that was assigned to the HBA before the detachment. For more information, see “Instructions: Detaching and reattaching an FCP storage group” on page 61.

Chapter 13. Adapter configuration

DPM automatically discovers any adapters that are installed in the processor frame, and assigns names to them, using a default naming convention. Before any administrators create and start partitions that use these adapters, use the **Manage Adapters** task to review details about the adapters, and make any adjustments that might be necessary.

Also, if HiperSockets are required on the system, you need to define them. See “Creating a HiperSockets adapter” on page 95 for an overview of the procedure.

For a system that has the DPM R3.1 storage management feature applied, you need to use the **Configure Storage** task to configure storage adapter cards as either FCP or FICON, and to specify which adapters are used to connect the system to storage devices. For more details about storage management with DPM R3.1 or later, see Chapter 17, “Managing storage on a DPM-enabled system,” on page 113.

Default naming convention

For most installed adapters, DPM assigns a default name in the form *adapter_type adapter_ID partial_location*. For cryptographic adapters, DPM uses a more specific adapter type. Table 7 provides examples of the default names that DPM assigns to installed adapters. In the table, PCHID refers to the physical channel path identifier, and Location includes jack and slot information.

Table 7. Sample default names for installed adapters

Type	Crypto type	PCHID	Card type	Location	Default name
Crypto	CCA coprocessor	01B8	Crypto Express4S	Z15B-LG18	CCA 01B8 Z15B-18
Crypto	EP11 coprocessor	01BC	Crypto Express4S	Z15B-LG19	EP11 01BC Z15B-19
Crypto	Accelerator	01C0	Crypto Express4S	Z15B-LG20	Accel 01C0 Z15B-20
FCP	—	0171	FICON Express 8	Z22B-D211-J.01	FCP 0171 Z22B-11
FCP	—	0161	FCP Express 32S	Z22B-D230-J.01	FCP 0161 Z22B-30
FICON	—	013C	FICON Express 16S	Z22B-D119-J.01	FICON 013C Z22B-19
OSD	—	018C	OSA-Express5S	Z15B-D104-J.01	OSD 018C Z15B-04
OSD	—	0144	OSA-Express7S 25GbE	Z22B-D121-J.01	OSD 0144 Z22B-21
RoCE	—	0184	10GbE RoCE Express	Z15B-LG02J.01-LG02J.02	RoCE 0184 Z15B-02
RoCE	—	014C	25 GbE RoCE Express	Z22B-LG23J.01-LG23J.02	RoCE 014C Z22B-23
zEDC	—	0188	zEDC Express	Z15B-LG03	zEDC 0188 Z15B-03

To help Linux administrators select the appropriate adapters for their partitions to use, consider either renaming the adapters with more descriptive names, or add a description to each adapter. If your company has an established naming convention for adapters, you can change the adapter names to follow that convention. To change an adapter name or add a description, open the **Manage Adapters** task. On the **Adapters** tab, select an adapter and right-click on its name to display the task menu, and click **Adapter Details**, as shown in Figure 40 on page 94.

Manage Adapters - S202B

Adapters | Devices | Cryptos | Related Tasks

Filter: Network Storage Accelerator Crypto

Search

Name	ID	Type	Status	State	Card Type	Location	Device Allocation	Number of Partitions	Description
Accel 0128 A13B-13	0128	Crypto	Active	Online	Crypto Express5s	A13B-LG13	0.0%	8	
CCA 0138 A13B-10	0138	Crypto	Not active	Standby	Crypto Express5s	A13B-LG18	0.0%	0	
CCA 0144 A13B-2	0144	Crypto	Not active	Standby	Crypto Express5s	A13B-LG21	0.0%	0	
EP11 0104 A13B-02	0104	Crypto	Active	Online	Crypto Express5s	A13B-LG02	0.0%	7	

Adapter Details

- Delete HiperSockets Adapter
- Create HiperSockets Adapter
- Reassign Channel Path IDs

Figure 40. Sample screen of the Manage Adapters task

The Adapter Details window opens. The details vary, depending on the adapter type, but all have Name and Description text boxes for you to use, if you want to modify the name or add a description. Figure 41 on page 95 shows the adapter details for a cryptographic adapter. After you modify the name or add a description, click **Apply** to save your changes.

Adapter Details - CCA 0138 A13B-18

General

Assigned Domains

Related Tasks

[Monitor System](#)

— ▾ General

* Name:	CCA 0138 A13B-18
Description:	
Object ID:	2BE5929C-79B8-11E6-A7F9-00106F0D5A70
System:	S202B
Status:	Not active
Detailed status:	Stopped
State:	Standby
Usage domain allocation: ?	0%
Adapter ID:	0138
Adapter type:	Crypto
Card type:	Crypto Express5S
Location:	A13B-LG18
LED:	
* Crypto number:	2
* Crypto type:	CCA Coprocessor ▾
UDX-Loaded:	Unknown

TKE access control can not be queried or modified when the adapter state is not Online

Figure 41. Sample screen of the Adapter Details window, with Name and Description fields highlighted

Creating a HiperSockets adapter

To create a HiperSockets adapter, open the **Manage Adapters** task, and click the down arrow to display the Actions list, as shown in Figure 42. Click **Create HiperSockets Adapter** in the list.

Manage Adapters - S223B

Adapters Devices Cryptos

Filter: Network Storage Accelerator Crypto

Actions Search

Name	Status	State	Card Type	Location	Device Allocation	Number of Partitions	Description
0131 A13B-17	Not active	Reserved	FICON Express16s	A13B-D216-J.01	0%	0	
0134 A13B-17	Not active	Reserved	FICON Express16s	A13B-D117-J.01	0%	0	
0135 A13B-17	Not active	Reserved	FICON Express16s	A13B-D217-J.01	0%	0	

Figure 42. Sample screen of the Manage Adapters task, with the expanded Actions list

The New HiperSockets Adapter window opens.

New HiperSockets Adapter - S223B

HiperSockets adapters have a switch that provides high-speed communications between partitions on a system using in-memory TCP/IP connections between and among partitions running several different operating systems, including Linux.

Name: HS10

Description:

MTU/Frame size (KB): 8/16, 16/24, 32/40, **56/64**, 56/64

OK Cancel Help

Figure 43. Sample screen of the New HiperSockets Adapter window

Enter a name, description, and MFS setting, and click **OK** to create the new adapter. A HiperSockets adapter supports several different maximum frame size (MFS) settings to accommodate different bandwidth requirements. The MFS setting determines the size of the largest packet that TCP/IP can transmit; on activation, TCP/IP adjusts its maximum transmission unit (MTU) according to the value of the MFS setting.

By default, this field is set to the smallest size: an MTU of 8 and frame size of 16 kilobytes (KB), which is displayed as 8/16. Select a value based on your knowledge of the bandwidth requirements for partition-to-partition communication. In most workload environments, the default value provides the most efficient use of system resources; however, for workloads that require increased bandwidth, for tasks such as large-file transfers and file backup, select a higher value to improve performance.

To display additional information about the new HiperSockets adapter, open the Adapter Details window.

Adapter Details - HS10

General

Network Interface Cards

General

* Name:

HS10

Description:

Object ID:

A56C8508-F06F-11E5-9051-9ABE94228451

System:

S223B

Status:

Not active

Detailed status:

Not active

State:

Standby

NIC allocation: ?

0%

Adapter ID:

07C0

Adapter type:

Internal queued direct communications (HiperSockets)

Channel path ID (CHPID): ?

48

Card type:

HiperSockets

MTU/Frame size (KB): ?

56/64

Figure 44. Sample screen of the new HiperSockets adapter in the Adapter Details window

Part 5. Topics for network administrators

Topics in this part provide information about network-related concepts and tasks that are specific to working with DPM-enabled systems.

Topics covered in this part are:

- Chapter 14, “Customizing the network boot server environment,” on page 101
- Chapter 15, “Securing FTP operations,” on page 105
- Chapter 16, “Using VLAN IDs to secure network communications,” on page 107

Chapter 14. Customizing the network boot server environment

Many companies use a preboot execution environment (PXE) to boot, install, and deploy hypervisors and operating systems. To use your company's network boot server, you might need to make some configuration changes to match the DPM implementation for the network boot server environment. Use the information in this procedure to customize your existing configuration files to work with the DPM NetworkBoot function.

Before you begin

- A network boot server environment usually consists of a Dynamic Host Configuration Protocol (DHCP) server, a Trivial File Transfer Protocol (TFTP) server and, optionally, a Domain Name System (DNS) server. These network boot server environments are based on the Intel Preboot eXecution Environment specification and specific Internet Engineering Task Force (IETF) requests for comments (RFCs) related to networking.

Before you complete the steps in this procedure, your installation must set up a DHCP server and a TFTP server, according to the IETF RFCs. These servers can reside on the same or on different systems within the network. Note that the configuration examples in this procedure work for Internet Systems Consortium (ISC) DHCP servers; the configuration statements for other DHCP servers are different.

- When communicating with network boot servers, DPM sends the universally unique identifier (UUID) of the partition in DHCP option 61, so the DHCP server must be configured to use that option and UUID in class or pool definitions.

You need to know the UUID of the partition, and the UUIDs of any other partitions for which you want to use the network boot server to install the hypervisor or operating system. To determine the UUID, open the **Partition Details** task for a specific partition, and find the value listed for Object ID in the General section.

To indicate that a request comes from an IBM Z system, DPM also sends an identifier in DHCP option 93. This **client-arch** option has a fixed value of 32. You can use this option to tailor the processing related to a request from an IBM Z system.

Procedure

- Create a configuration file for the DHCP server (for example: `/etc/dhcp/dhcpd.conf`), for IPv4 addresses.
 - Specify the following options.


```
allow booting;
allow bootp;
```
 - In the subnet specification for your network boot server environment, add one of the following statements, substituting the IPv4 address of your TFTP server for the variable `ipv4_addr`.


```
option tftp-server-name <ipv4_addr>;
      OR
next-server <ipv4_addr>
```
 - Complete one of the following options.

Option 1

- Define a class with a match for the option **dhcp-client-identifier**. This class can be used for program selection.


```
class "NetworkBoot" {
    match if option dhcp-client-identifier = "df7f2566-05aa-11e6-a31d-9abe94227999";
}
```
- Define a pool that includes both the dynamic address range and the boot configuration to load for all members of the class defined in the previous step. For example:

```
pool {
    allow members of "NetworkBoot";
    option bootfile-name      "NetworkBoot.PCI_D22H_004";
    filename                  "NetworkBoot.PCI_D22H_004";
    range 192.168.1.10 192.168.1.30;
}
```

- To specify the boot configuration, you can use either **option bootfile-name** (as shown in the previous example), or the **filename** keyword.
- For the boot configuration, specify the file name, optionally preceded by a relative path.
- The boot configuration has to be located under the “file-root” directory of the TFTP server, as described in step 2.

Option 2

Specify the dynamic address range and the boot configuration to load globally for your subnet. Use the same statements as described for the pool definition in Option 1.

- d. Optional: Create a configuration file for the DHCP server (for example: /etc/dhcp/dhcpd6.conf), for IPv6 addresses. The DHCP server configuration for IPv6 is similar to that for IPv4, only with different keywords. The key difference is the use of a URL to specify the boot configuration; for example:

```
option dhcp6.bootfile-url "tftp://[[fe80::400:faff:fe0a:802]/NetworkBoot.RHEL-7.0";
```

2. Verify the structure of the file-root directory of the TFTP server.

The file-root directory of the TFTP server is the top-level directory, where the TFTP server loads and stores files. It is specified by the *-s file root* option to **tftpd**. This directory contains the files associated with the program components to be loaded, as specified in the boot configuration. The supported program components are:

kernel A required component; the kernel of the program to load.

initrd A required component; a RAM disk for the program to load.

append

An optional component; additional parameters to the command line of the program to load.

To work with the DPM NetworkBoot function, the files associated with the kernel and initrd program components must reside under the file-root directory of the TFTP server. The file names can be preceded by a relative path.

The following examples illustrate boot configuration files. (Note that the examples are formatted to fit on this page, so line breaks might not match the formatting required for a boot configuration file.)

Example: Boot configuration file for an automated KVM server installation

```
~ # cat /tftpboot/kvm.DPM
PROMPT 1
DEFAULT linux
TIMEOUT 60
label linux
    kernel=s390x/kernel.kvm.s390x
    initrd=s390x/initrd.kvm.s390x
    append=ro systemd.show_status=0 ramdisk_size=40000 cio_ignore=all,!condev,
        !0.0.eb00-0.0.eb02,!0.0.7000,!0.0.7100
rd.zfcp=0.0.7000,0x500507680210d2ce,0x0000000000000000 rd.zfcp=0.0.7100,0x500507680220d2ce,
    0x0000000000000000
rd.znet=qeth,0.0.eb00,0.0.eb01,0.0.eb02,layer2=1,portno=1,portname=DUMMY
ip=172.16.47.231::172.16.47.1:255.255.255.0:zkvm231:enccw0.0.eb00:none nameserver=172.16.47.221
    searchdomain=z0plex.com vnc
vncpassword=pw4demo RUNKS=1 inst.repo=http://172.16.47.221/kvmibm
    inst.auto=http://172.16.47.221/kvmibm/auto/zkvm231.ks rd_NO_LUKS
rd_NO_LVM rd_NO_MD rd_NO_DM LANG=en_US.UTF-8
```

Example: Boot configuration file for an Ubuntu network installation system

```

~ # cat /tftpboot/ubuntu.DPM
PROMPT 1
DEFAULT linux
TIMEOUT 60
label linux
    kernel=s390x/kernel.ubuntu.s390x
    initrd=s390x/initrd.ubuntu.s390x
    append=ro locale=C DEBCONF_DEBUG=5 s390-netdevice/choose_networktype=qeth
        s390-netdevice/qeth/choose=0.0.eb00-0.0.eb01-0.0.eb02 s390-netdevice/qeth/port=1
        s390-netdevice/qeth/layer2=true netcfg/use_autoconfig=false
        netcfg/disable_dhcp=true
netcfg/get_ipaddress=172.16.47.251 netcfg/get_netmask=255.255.255.0
netcfg/get_gateway=172.16.47.1 netcfg/get_nameservers=194.25.0.60
netcfg/get_hostname=zsvm251 netcfg/get_domain=z0plex.com network-console/password=lin390
    network-console/password-again=lin390
preseed/url=http://172.16.47.221/auto/ubuntu.preseed ro systemd.show_status=0
ramdisk_size=40000

```

For these examples of boot configuration files, the file structure under the file-root directory of the TFTP server must look as follows:

```

~ # ls -al /tftpboot/s390x/
total 83524
drwxrwxr-x 2 atftp root      4096 15. Sep 12:41 .
drwxrwxr-x 5 atftp users    4096 21. Jul 19:12 ..
-rwxr-xr-x 1 atftp root    26113392 22. Feb 2016 initrd.kvm.s390x
-rwxr-xr-x 1 atftp root    30881788 21. Apr 21:02 initrd.sles121.s390x
-rwxr-xr-x 1 atftp root    10605354 21. Apr 14:46 initrd.ubuntu.s390x
-rwxr-xr-x 1 atftp root    3859968 25. Jan 2016 kernel.kvm.s390x
-rwxr-xr-x 1 atftp root    10245888 21. Apr 21:02 kernel.sles121.s390x
-rwxr-xr-x 1 atftp root    3687400 21. Apr 14:46 kernel.ubuntu.s390x
~ #

```

3. Through the **Partition Details** task on the HMC, complete following steps to specify the network boot server as the method of booting the operating system or hypervisor for the partition.
 - a. In the Network section, define a network interface card (NIC) to connect the partition to the network on which the network boot server resides.
 - b. In the Boot section:
 - 1) Select the **Network server (PXE)** option from the “Boot from” menu list.
 - 2) Select the NIC for the adapter that connects the partition to the network on which the network boot server resides.
 - c. Click **OK** to create the partition or to save your changes.
4. Edit the configuration file for the DHCP server to contain the Object ID value of each partition that is to be a client of the network boot server.

The following sample shows the option `dhcp-client-identifier` specified in a partial DHCP server configuration file.

```

pool {
    option bootfile-name      "NetworkBoot.DPM";
    filename                  "NetworkBoot.DPM";
    range 192.168.9.210 192.168.9.214;
}
host zsvm251 {
    option dhcp-client-identifier "df7f2566-05aa-11e6-a31d-9abe94227999";
    fixed-address 192.168.9.219;
    option bootfile-name      "kvm.DPM";
    filename                  "kvm.DPM";
}

```

Results

The network boot server environment is ready for use, and you can successfully start the partitions that use the **Network server (PXE)** boot option to load and initialize the operating system or hypervisor.

Chapter 15. Securing FTP operations


On a DPM-enabled system, administrators can use their choice of protocol for file transfers: standard FTP, FTP Secure (FTPS), and Secure File Transfer Protocol (SFTP). Standard FTP requires an FTP server to be on the same network as your DPM-enabled system. In contrast, with FTPS or SFTP, you can keep your system on an isolated network for maximum security.

FTPS and SFTP work through a proxy feature on the Hardware Management Console (HMC) through which you manage your DPM-enabled system. To use these secure protocols, you must first prepare at least one of the HMCs through which you manage the Support Element of the DPM-enabled system. For instructions, see “Setting up FTPS” or “Setting up SFTP” on page 106.


Setting up FTPS

FTPS uses the Secure Socket Layer (SSL) protocol to secure data through certificates that authenticate the FTP servers. To use the FTPS protocol, complete the following steps to import an FTPS server certificate.


1. Open the **Certificate Management** task, click the arrow for the **Advanced** list, and select **Manage Trusted Signing Certificate** from the list.



2. On the Manage Trusted Signing Certificate page, click the arrow for the **Import** list, and select **From Remote Server**.



3. On the Import Remote Certificate page, provide an IP/Host address and a valid port number, then click **OK**.



Import Remote Certificate

Enter the IP or hostname and port number.

IP/Host

Port


OK Cancel Help

4. Confirm the request. When processing completes, view the resulting message to determine whether the operation was successful.
5. Repeat, as necessary, on other managing HMCs. When you acquire the appropriate SSL certificates, you can use the FTPS or SFTP selections that are available for all the tasks that support FTP or removable media as options to import or export files.

Setting up SFTP

SFTP uses the Secure Shell (SSH) protocol to secure data through SSH keys that authenticate the FTP servers. To use the SFTP protocol, complete the following steps to import SSH server keys.

1. Open the **Manage SSH Keys** task.



Manage SSH Keys

Known Host Keys

--- Select Action ---

Select	IP Address	Key Type	Key Fingerprint
☐	9.60.14.37	RSA	47:9d:4a:e4:39:6b:c3:81:82:bb:bb:1a:f3:a6:a0:a3
☐	9.60.14.37	RSA	47:9d:4a:e4:39:6b:c3:81:82:bb:bb:1a:f3:a6:a0:a3
☐	progers.endicott.ibm.com	RSA	51:bd:51:be:af:83:c1:57:d4:84:45:1d:74:4e:25:cd

Total: 3 Selected: 0

Delete

Add Host Key

Address: Add

OK Help

2. In the Address field, provide the SFTP server ID or host name, and click **Add** to add the SSH key.
3. When processing completes, view the resulting message to determine whether the operation was successful.
4. Repeat, as necessary, on other managing HMCs. When you acquire the appropriate SSH keys, you can use the FTPS or SFTP selections that are available for all the tasks that support FTP or removable media as options to import or export files.

Chapter 16. Using VLAN IDs to secure network communications

When you use the **New Partition** task in basic mode to create a network interface card (NIC) for a new partition, that new partition can receive all traffic that flows across the physical adapter to which the NIC is assigned. If you want to limit that traffic to a particular virtual local area network (VLAN), you can specify a VLAN identifier (ID) for each NIC through the **New Partition** task in advanced mode, or through the **Partition Details** task. This *VLAN enforcement* capability can help you avoid data security issues.

DPM provides VLAN enforcement only for partitions with a type of **Linux** or **z/VM**, and only when you select an OSA-Express or HiperSockets adapter for the NIC. By specifying a VLAN ID, you are configuring the network adapter to send only packets that are tagged with the same ID, or untagged packets, to the partition and the operating system or hypervisor that it hosts. The operating system or hypervisor must tag and untag the packets.

For example, consider the configuration in Figure 45, which shows eight partitions, each hosting a guest operating system. Each of these partitions has a NIC through which it accesses one of two adapters that are connected with a crossover cable.




Figure 45. Sample VLAN configuration in partitions and guest operating systems

- All but two of the guest operating systems have a configuration file that contains a specific VLAN ID:
 - The operating systems on partitions A and E are configured to use VLAN 20.
 - The operating systems on partitions B, C, F, and G are configured to use VLAN 32.
 - The operating systems on partitions D and H do not have a VLAN ID configured.
- Each of the partitions has a NIC, but only six of the partitions have a NIC that is configured with a specific VLAN ID.
 - Partitions A and E have NICs that are configured with VLAN ID 20, which matches the ID value that is configured on their guest operating systems. In this case, VLAN enforcement is in effect, and the guests on those partitions are able to both send and receive network traffic on VLAN 20.
 - Partitions B and F have NICs that are configured with VLAN ID 25, but their guest operating systems are using a different VLAN ID value of 32. In this case, VLAN enforcement is in effect. However, because the VLAN ID values do not match, the guests on these partitions cannot receive

any network traffic on VLAN 32. To correct this communication problem, the VLAN ID values in both the partition definitions and the operating system configuration files must match.

- Partitions C and G have NICs that are not configured with any VLAN ID, so VLAN enforcement is not in effect. In this case, the guest operating systems can receive network traffic on any VLAN, even though they are configured to send and receive network traffic on VLAN ID 32.
- Partitions D and H have a NIC that is configured with VLAN ID 25, but their guest operating systems do not have any configured VLAN ID value. In this case, VLAN enforcement is in effect and the guests on these partitions cannot receive any network traffic. To correct this communication problem, VLAN ID 25 must be configured on both of the operating systems.

In addition to configuring a VLAN ID, you also can specify a unique media access control (MAC) address to identify a NIC. You must use either the **New Partition** task (advanced mode) or the **Partition Details** task to specify a MAC address.

- You can specify a MAC address for any type of partition, but only when you select an OSA-Express or HiperSockets adapter for the NIC. If you do not specify a value, DPM automatically generates a unique MAC address for the NIC.
- To display these MAC addresses on the operating system, you need to use a version of the Linux operating system that supports the specification or generation of MAC addresses through firmware. See the Red Hat, SUSE, or Ubuntu product information page to determine which RHEL, SLES, or Ubuntu Server version provides this support. (On an operating system without this support, the operating system's network device driver overwrites any MAC addresses with addresses that it generates.)

To configure VLAN enforcement for an existing partition, complete the following steps. You can follow the same steps for a new partition as well, but you must use the advanced mode of the **New Partition** task. Changing a NIC for an existing, active partition is considered a disruptive task, so you might want to stop the partition first.

1. Select the existing partition and open the **Partition Details** task.
2. In the **General** section, make sure that the partition type is either **Linux** or **z/VM**.
3. In the **Network** section, either select an existing NIC or create a new one. On the NIC Details or New Network Interface Card page, complete the following steps:
 - a. In the Adapter Ports and Switches table, verify that the selected adapter port or switch is either an OSA-Express or HiperSockets adapter.
 - b. Enter a value in the **VLAN ID** field. The valid range of VLAN IDs is 1 - 4094.
 - c. Verify that the **VLAN Type** field displays VLAN Enforcement.
 - d. Optionally, specify a MAC address that is both locally administered and unicast, which consists of six groups of two lower-case hexadecimal digits, separated by colons; for example: 02:ff:12:34:56:78. DPM checks the validity and uniqueness of the value that you supply, and issues a message if it finds an error.
 - e. Click **OK**.
 - f. In the **Network** section, verify your changes in the NICs table.
4. Click **OK** to save your changes. (If the partition is active, the Confirm Disruptive Action dialog opens; review the action and supply any confirmation values and your password, as necessary. Then click **Save** to continue.)

When your changes are successfully saved, the Validation window lists the device number and the VLAN ID for each new NIC that you defined.

5. To complete the setup for VLAN enforcement, you must specify the same VLAN ID in the network configuration files for the operating system or hypervisor.

For example, the following network configuration script, `ifcfg-Vnic-0SD`, in the `/etc/sysconfig/network-scripts` folder on a KVM hypervisor, contains the information that is required to start the

VLAN on the guest operating system. Note the device number and the VLAN ID that are highlighted; these values must match the device number and VLAN ID specified for the NIC in the **Partition Details** task.

```
NAME=Vnic-OSD
IPADDR=192.168.1.201
PHYSDEV=enccw0.0.5010
VLAN=yes
DEFROUTE=no
MTU=1500
NETMASK=255.255.255.0
BOOTPROTO=none
DEVICE=Vnic-OSD
TYPE=vlan
ONBOOT=no
VLAN_ID=700
IPV6INIT=no
```


Part 6. Topics for storage administrators

Topics in this part provide information about storage-related concepts and tasks that are specific to working with DPM-enabled systems.

Topics covered in this part are:

- Chapter 17, “Managing storage on a DPM-enabled system,” on page 113

Chapter 17. Managing storage on a DPM-enabled system

Starting with DPM R3.1, the **Configure Storage** task provides the controls through which you can configure and manage the storage resources for a DPM-enabled system. Through this task, system and storage administrators collaborate to connect a system to devices in the storage area network (SAN) through a simplified, visual, and automated process that does not require extensive knowledge of the mainframe. Administrators also use the **Configure Storage** task to request and fulfill storage resources for use by partitions on a DPM-enabled system.

In effect, the **Configure Storage** task replaces the use of either the Hardware Configuration Definition (HCD) or the Input/Output Configuration Program (IOCP) to define storage devices in an input/output configuration data set (IOCDS) for the system. Although multiple systems can share the same storage resources, note that the scope of this task is *one single system only*.

The **Configure Storage** task is available on the HMC only for a system that has the DPM R3.1 storage management feature applied. You can open it either from the context menu next to the system name, or from the Configuration task group. For automation purposes, the functions that are available through the **Configure Storage** task are also available through the HMC Web Services Application Programming Interface (API).

For more information about configuring and managing storage resources for a DPM R3.1 system, see the following topics:

- “A simplified and improved approach to storage management”
- “Setting up storage resources for the first time” on page 118
- “Managing storage resources” on page 125

A simplified and improved approach to storage management

To learn how the **Configure Storage** task simplifies and improves storage configuration and management, see the following topics.

- “A simplified, visual process for connecting storage to one system”
- “Integrated requests and notifications” on page 115
- “Access to FICON and FCP storage through storage groups” on page 116
- “Automatic detection and assignment of redundant resources” on page 117
- “Device number allocation for FICON storage groups” on page 118

A simplified, visual process for connecting storage to one system

To accomplish the initial setup of storage resources for a DPM R3.1 system, system and storage administrators use the **CONNECT TO STORAGE** wizard of the **Configure Storage** task. The initial setup includes defining the protocol of the adapter cards that are installed in the system I/O drawers, and building a visual copy of the FICON connections of the SAN to which this system is attached.

To make adapter card configuration easier to accomplish, the Configure Storage Cards page of the **CONNECT TO STORAGE** wizard displays a visual abstraction of each physical I/O drawer in the system. DPM automatically detects the installed adapters, which you can define as either FICON or FCP devices. Figure 46 on page 114 shows the visual display of one drawer on the Configure Storage Cards page.




Figure 46. Sample frame and card display on the Configure Storage Cards page

For a comparison, look at Figure 47, which is a photograph of a physical I/O drawer in a system. Note how the adapter card slots in the I/O drawer in Figure 46 mirror the size and location of actual slots in the photograph.



Figure 47. Photograph of a physical adapter card in a system I/O drawer

An I/O drawer can contain different types of adapter cards, so DPM highlights only the slots containing an adapter card that can be used for connections to storage devices. In Figure 46, the adapter cards that can be configured contain selectable labels. You can manually configure cards by clicking the FCP or FICON label. Or, for an automated approach, you can use other controls on the page to provide a total number of cards of each type for DPM to automatically select and configure for optimal redundancy and availability. This automated approach is available only when the system is not yet physically attached to SAN hardware through cables.

After configuring storage adapter cards, the next step is to configure all of the FICON storage devices in the SAN that are or will be connected to the DPM R3.1 system. As it does for adapter cards in an I/O drawer, the **Configure Storage** task also provides a simplified, visual approach to building a replica of FICON storage devices in the SAN. Figure 48 on page 115 shows a sample Configure FICON Connections page of the **CONNECT TO STORAGE** wizard, which provides a basic visual layout of the storage configuration, along with hover help to guide you through the process of defining a replica of the storage hardware that is or will be connected to the system. This configuration can contain at most two physical sites where storage devices are located. The primary site is always where the DPM-enabled system is physically located.

Highlighted controls on this page make it easy to add more storage subsystems, fabrics, and switches to match the physical hardware configuration of the SAN. Links within the boxes open windows through which you define the connections between specific devices. After you define these connections, lines in the display indicate which devices are connected, and the number of connections.




Figure 48. Sample display of two sites with physical storage hardware configured for the primary site

If the DPM-enabled system is not yet physically attached to SAN hardware through cables, DPM provides several automated options for this configuration process; for example, storage administrators have the option of having DPM select port connections. However, if the system is already cabled, you need to supply information that reflects the physical connections that are already in use.

As you progress through the **CONNECT TO STORAGE** wizard, DPM generates the virtual configuration that is required to connect the system to the physical SAN hardware; this virtual configuration is equivalent to the contents of an Input/Output Configuration Program (IOCP) file for the system. Completing the Connect to Storage wizard produces the following results:

- Fully configured adapter cards on the DPM-enabled system.
- Fully enabled FICON connections that link the system to physical elements in the SAN.

Integrated requests and notifications

System administrators and storage administrators work together to connect a system to storage, and to create and fulfill storage groups for partitions to use. The **Configure Storage** task provides integrated controls to help you collaborate with co-workers. For example, if you do not have authorization or need help to complete the initial storage setup for a system, you can click a button to invite a storage administrator to complete the configuration.

To facilitate the collaboration between system administrators and storage administrators, DPM automatically generates detailed requests for system administrators to send, and provides inline notification of results when the storage administrator has fulfilled the request. For integrated requests and notifications, storage administrators (recipients) must have an email address associated with their user IDs, and Simple Mail Transfer Protocol (SMTP) settings must be defined. Users who send email through the **Configure Storage** task do not require an assigned email address because DPM can generate one based on the user name, but the suggested practice is to assign email addresses for senders as well, so recipients know which person sent the email.

To take advantage of the integrated requests and notifications, your installation must define email addresses for users, and set up the Simple Mail Transfer Protocol (SMTP).

- For information about setting up email addresses, see “Authorizing users of the Configure Storage task” on page 118.
- For information about setting up SMTP, see the online help for the **Monitor System Events** task.

Access to FICON and FCP storage through storage groups

Administrators use the **Configure Storage** task to configure Fibre Connection (FICON) extended count key data (ECKD) direct-access storage devices (DASD), and Fibre Channel Protocol (FCP) Small Computer System Interface (SCSI) storage devices. Earlier releases of DPM supported only FCP SCSI storage, while DPM R3.1 introduces support for FICON ECKD DASD, with the following exceptions.

- DPM supports the use of parallel access volumes, but only through the optional HyperPAV feature on the IBM System Storage DS8000 series.
- DPM does not support the use of FICON channel-to-channel (CTC) or FICON tape devices.

After the initial storage configuration is complete and the system is connected to physical storage hardware by cables, system administrators can enable partitions to access and use those resources by creating storage groups. A *storage group* is a logical group of storage volumes that share certain attributes. One attribute is the type, which can be either FICON or FCP. Another example of an attribute is shareability, which indicates whether the storage group is defined as either dedicated for use by only one partition, or shared by multiple partitions.

To create a storage group, open the **Configure Storage** task and select one of the following options: **REQUEST STORAGE** or **CREATE STORAGE GROUP**. (The option that you can select depends on the authorization of your user ID.) Using this task does not require extensive mainframe knowledge; however, a storage administrator must fulfill any requests for storage through tools for managing storage subsystems.

You can define storage groups with or without the use of a template. By creating templates, administrators can predefine storage group attributes and volumes for one or more FCP or FICON storage groups. Depending on the template contents, creating a storage group can be as quick and easy as providing a unique name for the storage group, and sending a fulfillment request to a storage administrator.

When you submit a request for one or more storage groups, DPM automatically generates the world wide port names (WWPNs) that are allocated to virtual storage resources when the storage group is attached to a partition. Also, DPM automatically generates a request that you can send as email to one or more storage administrators to fulfill through the storage subsystem management interface.

To access storage, attach one or more storage groups to a partition, through either the **New Partition** task or **Partition Details** task. Figure 49 on page 117 shows several storage groups with key attributes, and information about the physical devices that an administrator has configured to fulfill them.

DPM-enabled system




Figure 49. Sample storage groups that are attached to partitions on the system

In Figure 49, the ovals in the partitions indicate which storage groups are attached to each partition.

- Partition A can access storage group SG01, which consists of a set of FICON base volumes that are dedicated for use by only one partition. Partition A accesses these volumes through FICON-mode adapter card 01 and switch A.
- Partitions A and B share storage group SG02, which consists of FCP disks that both partitions can access through FCP-mode adapter card 22 and switch B.
- Similarly, partitions B and C share storage group SG03, which consists of FCP disks that both partitions can access through FCP-mode adapter card 31 and cascaded switches B and C.

Storage group templates

To make storage requests even easier to complete, administrators can create a master copy, or template, for a FICON or FCP storage group. Templates can reflect typical usage patterns (production, staging, test); standardize storage for specific user groups; or document company requirements or restrictions that might be in place for storage use. With templates in place, users can select an available template and, with minimal changes, quickly submit a request for a new storage group. They can modify the template contents.

To create and manage storage templates, you can use the default SYSPROG, STORAGEADMIN, or SERVICE user IDs, or any user IDs that an access administrator has authorized to this task through customization controls in the **User Management** task. When you open the **Configure Storage** task, the options displayed on the Configure Storage page depend on the authorization of your user ID. To work with templates, select **REQUEST STORAGE**, **CREATE STORAGE GROUP**, or **MANAGE TEMPLATES**.

Automatic detection and assignment of redundant resources

DPM automatically discovers any adapters that are installed in the processor frame, and assigns names to them, using a default naming convention. When you use the **Configure Storage** task to configure the installed storage adapter cards, DPM highlights only the slots containing an adapter card that you can use for connections to storage devices.

- When you enter the number of FICON or FCP cards to be configured, DPM automatically selects a combination of the unconfigured cards to satisfy the request. This combination is *redundant*; that is, the configured cards of each type are spread across domains and drawers to ensure availability, in case of a card or drawer failure. You can override these selections by selecting the appropriate protocol label in each card: FICON or FCP. Selecting NONE resets the adapter card to the unconfigured state.

Depending on the type of adapter card and its cabling, you might not be able to change the protocol through these labels. For these adapter card types, the protocol labels are disabled.

- If the DPM-enabled system is not yet physically attached to SAN hardware through cables, DPM provides several automated options for configuring the FICON connections to storage devices in the SAN. For example, when you are defining the adapter ports that connect a system to switches or storage subsystems, you can select the Auto-connect icon (🔌) to have DPM automatically configure redundant ports.

Device number allocation for FICON storage groups

When you create a request for a dedicated or shared FICON storage group, you have the option of assigning device numbers to the requested volumes. If you do not assign device numbers, DPM automatically assigns them when the storage group is attached to a partition. At that time, DPM looks for device numbers that are already in use by existing partitions, if any, and avoids assigning any conflicting numbers.

Storage Overview, details, and connection reports

The **Configure Storage** task includes a storage overview through which you can view information about all of the storage groups on a DPM-enabled system. The Storage Overview page includes the Storage Groups table, which contains one row for each storage group. To view more details about a specific storage group, click anywhere in the table row for that storage group to open the Storage Group details page. The fulfillment state indicates whether the storage group is available for use.

Depending on the permissions that are associated with your user ID, you can modify or duplicate a storage group, delete one or more storage groups, or view a report with details about the connections between the system and physical storage hardware (subsystem, switches, and so on) that is in use for a specific storage group. The connection report for a selected FCP or FICON storage group can be useful for diagnosing configuration problems.

Setting up storage resources for the first time

When your company orders an IBM system with the required DPM features, IBM service representatives install the system and enable DPM, so it is ready for use when the system is powered on. Although an input/output configuration data set (IOCDs) is still required for configuring network, accelerator, and crypto devices, administrators use the **Configure Storage** task to initially configure the devices in the storage area network (SAN) that are or will be connected to the DPM R3.1 system.

The following topics provide more information about planning for and using the **Configure Storage** task to set up storage resources for a DPM-enabled system.

- “Authorizing users of the Configure Storage task”
- “Connecting a system to storage” on page 119
- “Requesting storage resources for partitions to use” on page 123

Authorizing users of the Configure Storage task

Use this procedure to create a new user for a system administrator who can also perform functions that are usually restricted to storage administrators. This new user has access permissions to all functions that are available through the **Configure Storage** task. When you create a new user, you can assign an email address so the new user can, if necessary, participate in the request, invitation, and notification functions that are integrated into the **Configure Storage** task.

Before you begin

Log in to the HMC with the default ACSADMIN user ID, or a user ID defined through the **User Management** task with equivalent permissions.

About this task

To open the **Configure Storage** task, administrators can use the default SYSPROG, STORAGEADMIN, or SERVICE user IDs, or any user IDs that an access administrator has authorized to this task through customization controls in the **User Management** task.

Some functions in the **Configure Storage** task require the use of the default STORAGEADMIN user ID, or a user ID defined through the **User Management** task with equivalent permissions. For access to all functions in the **Configure Storage** task, users must have the same permissions as both the default SYSPROG and STORAGEADMIN user IDs. Also, access permission to all storage adapters is required to configure storage cards, and access permission to all FICON adapters is required to configure FICON connections.

Procedure

1. In the HMC navigation pane, select **HMC Management** or **Tasks Index** and select the link to open the **User Management** task. The User Management dashboard is displayed.
2. Select the New User action icon (). The New User wizard opens to the Welcome page.
3. Select **Next** to open the Name page, and complete the following steps.
 - a. For Create Option, select **New based on**, and select SYSPROG from the list.
By default, selecting SYSPROG gives the new user permission to the **Configure Storage** task, as well as permission to the object types FCP Storage Group, FICON Storage Group, and FICON Adapter.
 - b. For User Details, provide a name and optional description. The suggested practice is to also specify an email address in the Email Address field, so the new user can participate in the request, invitation, and notification functions that are integrated into the **Configure Storage** task.
 - c. Select **Next** to go to the next page.
4. On the Authentication page, select an authentication type and provide any additional required details. Then select **Next** to go to the next page.
5. On the Roles page, select the following storage administrator roles.
 - Storage Administrator Objects
 - Storage Administrator Tasks
 Then select **Next** to go to the next page.
6. On the Summary page, review the information for the new user, and select **Finish**.

Results

The new user has the task and role permissions that are required to complete all actions that are available through the **Configure Storage** task. The new user also can send or receive requests or invitations, and receive notifications about work that has been completed through the **Configure Storage** task.

Connecting a system to storage

Review this topic to become familiar with the steps in the Connect to Storage wizard, through which administrators set up the initial storage configuration for a DPM R3.1 system. When you use the Connect to Storage wizard, it provides hover help to guide you through the process of defining a replica of the storage hardware that is or will be connected to the system. If you need more guidance, detailed instructions are available through the online help for the **Configure Storage** task.

Before you start the initial storage configuration

Before you open the **Configure Storage** task on the HMC, review the following information.

- DPM automatically discovers any adapters that are installed in the processor frame, and assigns names to them, using a default naming convention. Consider using the **Manage Adapters** task to review the installed storage card adapters, and make any adjustments that might be necessary.
- For integrated requests and notifications, storage administrators (recipients) must have an email address associated with their user IDs, and Simple Mail Transfer Protocol (SMTP) settings must be defined.
 - For information about setting up email addresses, see “Authorizing users of the Configure Storage task” on page 118.
 - For information about setting up SMTP, see the online help for the **Monitor System Events** task.
- To open the **Configure Storage** task, you can use the default SYSPROG, STORAGEADMIN, or SERVICE user IDs, or any user IDs that an access administrator has authorized to this task through customization controls in the **User Management** task. Also, access permission to all storage adapters is required to configure storage cards, and access permission to all FICON adapters is required to configure FICON connections. For instructions to authorize users, see “Authorizing users of the Configure Storage task” on page 118.

An overview of the initial storage configuration process

When you open the **Configure Storage** task, DPM detects whether the storage configuration has been started; if not, the Connect to Storage wizard opens automatically.

This configuration process does require administrators to know high-level information about the physical elements of the SAN, such as the names of storage subsystems, the types of devices and communication protocols, intended use, and so on. This information is usually available through a system plan for the company's physical IT site. Starting with the system plan, system and storage administrators use the **Connect to Storage** task to define the initial storage configuration, which consists of the following activities.

Step 1. Define storage adapter cards as FCP or FICON mode

The first task is defining the protocol of the adapter cards that are installed in the system I/O drawers. DPM automatically detects the installed adapters, which you can define as either FICON or FCP devices. Based on the number of cards that you specify, DPM automatically selects a combination of the unconfigured cards, but you can override these selections, if necessary.

Figure 50 shows a partial Configure Storage Cards page that contains the visual representation of one I/O drawer with FCP, FICON, and unconfigured storage adapter cards. Blank slots are either empty or contain network, accelerator, or crypto adapter cards.




Figure 50. Sample frame and card display on the Configure Storage Cards page

Step 2. Build a visual copy of the SAN hardware configuration

The second task is building a visual copy of the FICON connections of the SAN to which this system is or will be attached. This configuration can contain at most two physical sites where storage devices are located. The primary site is always where the DPM-enabled system is physically located. Figure 51 shows a sample Configure FICON Connections page of the **CONNECT TO STORAGE** wizard, which provides a basic visual layout of the storage configuration, along with hover help to guide you through the process of defining a replica of the storage hardware that is or will be connected to the system.




Figure 51. Sample display of two sites with physical storage hardware configured for the primary site

System administrators can define physical elements, such as switches, fabrics, and storage subsystems, but storage administrators are required to define port connections and logical control units (LCUs). If you invite a co-worker to help finish the configuration, DPM automatically generates an invitation that you can send to one or more storage administrators.

If the DPM-enabled system is not yet physically attached to SAN hardware through cables, DPM provides several automated options for this configuration process; for example, storage administrators have the option of having DPM select port connections. However, if the system is already cabled, you need to supply information that reflects the physical connections that are already in use.

Planning considerations for your SAN configuration

Use the following information for planning purposes. Note that these details are also available in the online help for the **Configure Storage** task.

Suggested practice for configuring storage adapter cards

On the Configure Storage Cards page, you configure adapter cards as either FICON or FCP mode by either specifying a number of each mode type, or configuring each card manually. The suggested practice is to specify numbers of each type, so that DPM automatically selects a combination of the unconfigured cards to satisfy your request. This combination is *redundant*; that is, the configured cards of each type are spread across domains and drawers to ensure

availability, in case of a card or drawer failure. You can override these selections by selecting the appropriate protocol label in each card: FICON or FCP. Selecting NONE resets the adapter card to the unconfigured state. Depending on the type of adapter card and its cabling, you might not be able to change the protocol through these labels. For these adapter card types, the protocol labels are disabled.

How many adapters you specify depends on your system plan; for example, if this system is expected to host test rather than production workloads, you might specify fewer adapter cards. Redundancy might be less important for a test workload, as well.

Conventions and maximums for configuring FICON connections

On the Configure FICON Connections page, you need to define the storage hardware that is or will be connected to this DPM-enabled system. The hardware that you define on this page must match the planned or actual configuration of storage subsystems and switches in the SAN.

- As you provide names for physical storage elements, note that supported characters are alphanumeric, blanks, periods, underscores, dashes, or at symbols (@). Names cannot start or end with blank characters.
- Add physical storage elements to the primary site and, optionally, to a secondary site. To create the secondary site, you can clone the primary site any time after you add the first storage subsystem or switch to the primary site. DPM duplicates the primary site subsystems, fabrics, and switches, along with adapter ports and LCUs, but you must provide unique names and IDs for the storage subsystems and switches.
- Your SAN configuration can consist of point-to-point (direct) connections or switch connections, but not both.
- Fabric names must be unique among fabrics, and cannot exceed 32 characters. Switch IDs must be unique within a fabric, and consist of hexadecimal values in the range 01 - EF. You can create a maximum of 256 fabrics per site, and a maximum of 239 switches per fabric.

The suggested practice is to define all fabrics as high integrity fabrics. A high integrity fabric is configured to use 2-byte link addresses, rather than 1-byte link addresses.

- If you define a fabric as high integrity, you must ensure that the physical switches are configured to use 2-byte link addresses. A high integrity fabric is required when you configure cascaded switches.
- Note that, after you save the FICON configuration, you cannot modify the high integrity fabric setting when LCUs are in use for switches in the fabric. If you need to modify the high integrity fabric setting, you must first remove all of the paths from the LCUs that are using that fabric, and then modify the high integrity fabric setting. After modifying the setting, you can restore the paths by adding them back to the LCUs.
- Storage subsystem names must be unique among storage subsystems, and cannot exceed 64 characters. You can create a maximum of 256 storage subsystems per site.
- For availability, connect each storage subsystem or switch to at least two adapter ports, each of which resides in a different frame and domain of the system.
- For availability, the suggested practice is to connect each storage subsystem to at least one switch per fabric, and at least two ports on each switch. You can define a maximum of 64 connections to switches on each subsystem. When you select adapter ports to connect the system to storage subsystems or switches, you can have DPM automatically select and balance connections for availability.
- LCU numbers must be unique within the storage subsystem; valid values are in the range 00 - FF. You can specify a maximum of 256 LCUs per storage subsystem. You can specify up to a combined total of 256 base and alias volumes for one LCU. When you add LCUs, DPM automatically selects paths to maximize redundancy and to reduce common points of failure.

Results of using the Connect to Storage wizard

Using the Connect to Storage wizard produces the following results:

- Fully configured adapter cards on the DPM-enabled system.
- Fully enabled FICON connections that link the system to physical elements in the SAN.

What to do next

- Use the **Request Storage** task to define storage resources, known as storage groups, for partitions to use. For more information, see “Requesting storage resources for partitions to use.”
- If you need to modify this initial storage configuration at a later time, open the **Configure Storage** task and, depending on what you need to change, select **STORAGE CARDS** or **FICON CONNECTIONS**. Before making any modifications, review the planning considerations in “Modifying the storage configuration of the system” on page 133.

Requesting storage resources for partitions to use

Review this topic to become familiar with the steps in the Request Storage task, through which administrators request storage resources for use by partitions on a DPM R3.1 system. You can access the Request Storage task by selecting **REQUEST STORAGE** or **CREATE STORAGE GROUP** in the **Configure Storage** task. (The option that you can select depends on the authorization of your user ID.) Using this task does not require extensive mainframe knowledge; however, a storage administrator must fulfill any requests for storage through tools for managing storage subsystems. If you need more guidance when you use this task, see the detailed instructions that are available through the online help for the **Configure Storage** task.

Before you create the request for a storage group

- Before you can successfully request storage, physical storage hardware (subsystems, switches, and so on) must be connected by cables, and storage cards must be configured.
- For integrated requests and notifications, storage administrators (recipients) must have an email address associated with their user IDs, and Simple Mail Transfer Protocol (SMTP) settings must be defined.
 - For information about setting up email addresses, see “Authorizing users of the Configure Storage task” on page 118.
 - For information about setting up SMTP, see the online help for the **Monitor System Events** task.
- To open the **Configure Storage** task, you can use the default SYSPROG, STORAGEADMIN, or SERVICE user IDs, or any user IDs that an access administrator has authorized to this task through customization controls in the **User Management** task. For instructions to authorize users, see “Authorizing users of the Configure Storage task” on page 118.

An overview of the Request Storage process

To request storage for partitions to use, you define one or more storage groups. A *storage group* is a logical group of storage volumes that share certain attributes. Storage groups can be shared by multiple partitions, and multiple storage groups can be attached to one partition.

You can define storage groups with or without the use of a template. By creating templates, administrators can predefine storage group attributes and volumes for one or more FCP or FICON storage groups. Templates can reflect typical usage patterns (production, staging, test); standardize storage for specific user groups; or document company requirements or restrictions that might be in place for storage use. You can select an available template and, with minimal changes, quickly submit a request for a new storage group.

The following steps describe the process of creating a storage group with or without the use of a template. Note that templates are based on the storage group type: FCP or FICON. An administrator can create one or more templates for each type, but you can select only a template that matches the current

system configuration. For example, if the system does not have any storage adapter cards that are configured as FICON, you cannot select any template for a FICON storage group.

Step 1. Specify the attributes for the storage group

On the Specify Storage Attributes page, define the attributes for the storage group.

- **Type:** FICON or FCP. This setting represents the type of storage devices that the storage group can use, and also controls the other attributes or default settings that are displayed on this page. The adapter card configuration for the system determines which type you can select; for example, if the system does not have any adapter cards that are configured as FICON, you cannot select that type. Also, if you are using a template for this storage request, you cannot change the storage group type.
- **Shareability:** Shared or Dedicated. If you select **Dedicated**, then only one partition is able to use this storage group. If you select **Shared**, specify the number of partitions that can share this storage group. The maximum number of partitions is set automatically to the system limit.
- **Connectivity:** The the number of paths to be available for use by each operating system with access to this storage group. The number of paths that you can define varies, depending on the storage group type. For FCP, the limit is the total number of adapters that are configured as FCP on the system; for FICON, the limit is the number of adapters that are configured as FICON on the system, up to a maximum of eight. The number that you select affects overall bandwidth, performance, and redundancy.
- **Optimized for 2nd level virtualization:** For a dedicated FCP storage group only, select **Optimized for 2nd level virtualization** when you want to enable the direct assignment of host bus adapters (HBAs) so an operating system or its guests can access the storage group. DPM distributes additional HBAs as equally as possible, taking into account both fabrics and adapters that will be assigned to this storage group, as indicated through the Connectivity attribute setting.

Step 2. Specify the size and type of each volume to be added to the storage group

On the Add Storage Volumes page, specify the size and type of each volume to be added to the storage group. If you are using a template for this storage request, the table on this page might contain predefined volumes that you can modify. Otherwise, this page initially contains a table heading with controls for defining a volume. The information you provide for this step varies, depending on the type of storage group: FCP or FICON.

1. For FCP, provide the capacity in gibibytes (GiB); for FICON, select a predefined model or custom size, and DPM automatically fills in the appropriate values for GiBs and cylinders.
2. For both FCP and FICON storage group volumes, specify the type as either **Data** or **Boot**.
3. Optionally, for FICON volumes only, you can assign device numbers. The suggested practice is to have DPM automatically assign device numbers to avoid conflicts; in this case, DPM assigns device numbers only when the storage group is first attached to a partition.
4. Optionally, you can provide a description of the FCP or FICON volume, and duplicate it by specifying the number of copies that you want.

Step 3. Name the storage group and, optionally, specify duplicates

On the Name and Duplicate page, specify the name of your new storage group. Optionally, provide a description and, if you want to easily duplicate the storage group, enter the number of duplicates and modify the names as necessary. For the name of the storage group, specify a value that is 1 - 64 characters in length. Supported characters are alphanumeric, blanks, periods, underscores, dashes, or at symbols (@). Names cannot start or end with blank characters. The name must uniquely identify the storage group from all other storage groups that are defined for this system.

If you request any duplicates, DPM uses the storage group name that you provided, and appends an underscore and sequential number to give each duplicate a unique name. You can edit the name of any duplicate, but all names must be unique.

Step 4. Send the storage group request to a storage administrator for fulfillment

DPM automatically generates a storage request for each storage group that you have requested. If your installation has configured SMTP, you can fill in the name of one or more storage administrators and submit the request. Otherwise, you can download or copy the generated request and send it to one or more storage administrators.

Results of using the Request Storage task

When you submit your request for one or more storage groups, DPM automatically generates the world wide port names (WWPNs) that are allocated to virtual storage resources when the storage group is attached to a partition.

After your request is received, the storage administrator uses storage management tools to select and configure the physical storage volumes that fulfill your request.

- For an FCP storage group, DPM periodically checks for the requested volumes, and updates their fulfillment status on the **Storage Overview** tab of the **Configure Storage** task. When the storage administrator completes the configuration through tools for managing storage subsystems, DPM changes the storage group status to Complete.
- For a FICON storage group, the storage administrator not only completes the configuration through tools for managing storage subsystems, but also maps the base volumes (which were created on the storage subsystem to fulfill your request) to volumes that you requested for the storage group. DPM changes the storage group status to Complete only after the volumes have been mapped through the **Map Volumes** action in the **Configure Storage** task. The **Map Volumes** action is accessible only through the Storage Details page for a FICON group, and is accessible only to storage administrators who are using the default STORAGEADMIN user ID, or a user ID with equivalent permissions.

What to do next

- You can create another storage group or select **GO TO STORAGE OVERVIEW** to view the status of the requested storage groups. If you are viewing a page in the **Configure Storage** task when the fulfillment status for the storage group changes to Complete, you receive an online notification.
- If you have defined boot and data volumes that are the same size, and your storage administrator fulfills your storage request with a preinstalled boot volume, you need to make sure that the correct volume is identified as the boot volume. Using the preinstalled boot volume ID that you receive from your storage administrator, go to **Storage Overview**, select the storage group to open the Storage Group details page, look up the volume ID on the **VOLUMES** tab, and make sure that the type for the preinstalled volume is Boot. If it is not, select the **Modify** icon to change the volume type.
- Use the **Partition Details** task to attach one or more storage groups to an existing partition, or attach them when you create a new partition through the **New Partition** task. You can attach storage groups that are not fulfilled yet; however, DPM issues a warning if you try to start a partition or apply changes to an existing partition before the storage group is fulfilled. Although the partition can be started, the operating system and applications that run on it might not function properly because some storage is not available until the fulfillment status for the storage group changes to Complete.

Managing storage resources

After the initial storage configuration for the system is complete, and storage groups are available for partitions to use, you can use the **Configure Storage** task to modify the system configuration, or to view and modify storage groups. You can also use either the **New Partition** task or **Partition Details** task to attach storage groups to a partition. For overviews of specific storage management tasks, see the following topics.

- “Viewing and managing storage groups” on page 126
- “Attaching storage resources to a partition” on page 130
- “Modifying the storage configuration of the system” on page 133

Viewing and managing storage groups

Review this topic to become familiar with the content of the Storage Overview and Storage Details pages, through which administrators can view and modify storage groups on a DPM R3.1 system. More detailed information is available through the online help for the **Configure Storage** task.

At-a-glance view of storage groups: the Storage Overview page

Use the Storage Overview to view information about all of the storage groups on a DPM-enabled system. The Storage Overview page includes the Storage Groups table, which contains one row for each storage group. To view more details about a specific storage group, click anywhere in the table row for that storage group to open the Storage Group details page. The fulfillment state indicates whether the storage group is available for use. You can access the Storage Overview page by selecting **STORAGE OVERVIEW** in the **Configure Storage** task.

Depending on the permissions that are associated with your user ID, you can modify or duplicate a storage group, delete one or more storage groups, or view a report with details about the connections between the system and physical storage hardware (subsystem, switches, and so on) that is in use for a specific storage group. The default STORAGEADMIN user ID (or users with equivalent permissions) can complete a storage request by mapping volumes for a FICON storage group through the Storage Group details page.

The Storage Groups table contains the following information about each storage group.

NAME

Specifies the user-defined name of the storage group.

TYPE Specifies the type of storage group: FICON or FCP.

PARTITIONS

Specifies the number of partitions to which the storage group is attached.

SHAREABLE

Specifies whether the storage group can be shared among partitions, or whether it is dedicated to only one partition.

TOTAL CAPACITY

Specifies the total amount of storage in gibibytes (GiBs) that is assigned to the storage group.

DESCRIPTION

Specifies the user-provided description, if any, of this storage group. The description can be up to 200 characters in length.

FULFILLMENT STATE

Identifies the current state of the storage group. DPM runs a background check of storage resources for FCP storage groups and, if necessary, changes the fulfillment state. These checks are more frequent (every 10 minutes) for fulfillment states other than Complete (every 24 hours). Users can manually start a background check by selecting the **Connection Report** icon, and selecting the **Update report** icon (🔄).

ACTIONS

In any table row, select the ellipsis (...) to display a selectable list of actions that you can take for the storage group. The listed actions vary, depending on the type of storage group, and on the permissions that are associated with your user ID.

All of the details: the Storage Group details page

Use the Storage Group details page to view or modify information about a specific FCP or FICON storage group on a DPM-enabled system. The Storage Group details page consists of a summary, a set of action icons, and tabbed sections that you can select to change the lower portion of the page display. The

actions that are displayed depend on the fulfillment state of the storage group, the type of storage group, and the authorization of the user who is accessing the Storage Group details page. DPM runs a background check of storage resources for FCP storage groups and, if necessary, changes the fulfillment state. These checks are more frequent (every 10 minutes) for fulfillment states other than Complete (every 24 hours). Users can manually start a background check by selecting the **Connection Report** icon, and selecting the **Update report** icon (🔄).

For specific fulfillment states, various summary fields, tabbed sections, and table entries have a pending, incomplete, or warning icon to alert you to details that might need your attention or action. The following list describes the fulfillment state values.

Checking migration

An existing DPM configuration was either upgraded to DPM R3.1, or was migrated from one system to a different system that has the DPM R3.1 storage management feature applied. During the migration or upgrade process, DPM automatically creates a dedicated FCP storage group for each partition, based on all of the host bus adapters (HBAs) assigned to the partition. This fulfillment state indicates that DPM is checking the logical and physical elements that support an FCP storage group it created during a system migration or firmware upgrade process.

In some cases, this check detects a storage group that cannot be fulfilled because logical unit numbers (LUNs) are not visible. For such cases, the fulfillment state does not change from Checking migration. The storage administrator must fix the configuration in the storage subsystem, and an administrator must open the **Connection Report** and select the **Update report** icon (🔄), so that DPM can recheck the storage group and change the fulfillment state.

For more details about the migration process, see the appropriate topic:

- “Migration instructions for a system replacement” on page 82
- Chapter 12, “Upgrading from DPM R3.0 to R3.1,” on page 89

Complete

The FCP or FICON storage group is ready for use.

Incomplete

One or more volumes or adapters that are used for an FCP or FICON storage group are marked as incomplete. DPM periodically checks the availability of storage volumes or adapters for FCP storage groups, so resources that were functioning properly can become incomplete.

The **VOLUMES** or **ADAPTERS** tab, and specific table entries for the tab display, are marked with the incomplete icon (❗). To diagnose the problem, select the **Connection Report** action icon (🔗).

Pending

A system administrator has sent a request to create or modify an FCP or FICON storage group, but the storage administrator has not finished fulfilling that request through tools for managing storage subsystems. When a creation, modification, or deletion request is in progress, the affected summary fields, section tabs, and table entries are marked with the pending icon (🕒).

Pending with mismatches

A system administrator has sent a request to create or modify an FCP storage group, and the storage administrator has fulfilled that request, but with an amount of storage that does not exactly match the original request. In this case, the **VOLUMES** tab has a pending icon (🕒) next to it, and the table rows for mismatched volumes are marked with a warning icon (⚠️).

The tabbed sections provide more details related to the storage group. The tabbed sections vary, depending on the type of storage group. Note that the pending icons (🕒 or ⚠️), incomplete icon (❗), and warning icon (⚠️) indicate details that might require your attention or action.

The following list describes the tabbed sections on the Storage Details page.

VOLUMES

The **VOLUMES** tab display varies, depending on the type of storage group. For an FCP storage group, the tab displays a single table that lists all volumes associated with the storage group. For a FICON storage group, the tab displays two tables: Base Volumes and Alias Volumes. Depending on the volume type, the table display includes the following information: storage subsystem, volume ID or UUID, capacity, type, device ID, and a user-supplied description, if any.

The **VOLUMES** tab display also varies depending on the current fulfillment state of the storage group. Table 8 lists the fulfillment states, describes the table display, and provides possible actions you might take to resolve any issues. Note that, in some cases, more than one type of pending request might be in effect.

Table 8. Affect of fulfillment status on the Volume tab display

Fulfillment status	Description	Volume tab display	Possible action
Checking migration	This fulfillment state indicates that DPM is checking the logical and physical elements that support an FCP storage group it created during a system migration or firmware upgrade process.	The display contains an empty Volumes table.	When DPM completes the check, it changes the fulfillment state to Complete or Pending with mismatches. In some cases, this check detects a storage group that cannot be fulfilled because logical unit numbers (LUNs) are not visible. For such cases, the fulfillment state does not change from Checking migration. The storage administrator must fix the configuration in the storage subsystem, and an administrator must open the Connection Report and select the Update report icon (🔄), so that DPM can recheck the storage group and change the fulfillment state.
Complete	DPM has successfully detected all of the logical and physical elements that support the volumes in this storage group.	The display contains a table entry with complete information for all volumes in this storage group.	None.
Incomplete	One or more volumes are incomplete. Volumes can be marked as incomplete under the following conditions. - When DPM can no longer detect (or sense) them in the FICON configuration. - When DPM found a problem when checking the results of a system migration or firmware upgrade process to the DPM R3.1 storage management feature.	An incomplete icon (🔴) is displayed in the table entry of each incomplete volume.	To diagnose the problem, select the Connection Report action icon (🔍).
Pending (creation request)	When a creation request is in progress, some information in the volume table is not available yet.	A pending icon (🟡) marks incomplete table rows until DPM provides the information.	In the case of a FICON storage group, a storage administrator must select the Map Volumes action icon (📍) and complete that task before DPM can complete the table.

Table 8. Affect of fulfillment status on the Volume tab display (continued)

Fulfillment status	Description	Volume tab display	Possible action
Pending (modification request)	A modification request is in progress for one or more volumes.	A pending icon (🕒) marks the table entries for volumes to be modified. The CAPACITY column indicates not only the current volume size, but also the pending change.	None. When the modification request is satisfied, the pending icons are removed, and the CAPACITY column values are updated to show only the modified size.
Pending (deletion request)	A deletion request is in progress for one or more volumes.	A pending icon (🕒) marks the table entries for volumes to be deleted, and the values in those table entries are crossed out.	None. When the deletion request is satisfied, the table entries are removed from the display.
Pending with mismatches	The volume sizes do not match the original storage request for an FCP storage group, or a migration process resulted in mismatched volumes. The following conditions describe the various conditions under which DPM assigns this fulfillment state to a storage group. 1. DPM detected more volumes for this storage group than the number that was originally requested. 2. DPM detected the correct number of volumes, but the volume sizes are either larger or smaller than originally requested. 3. DPM detected a volume but that volume is not accessible through all of the worldwide port numbers (WWPNs) that are available for use with this storage group.	All of the mismatched volumes are displayed at the top of the table, enclosed in a box. A warning icon (⚠️) is displayed in the CAPACITY column of each mismatched volume. Both the actual and requested sizes are shown in the CAPACITY column, along with a message that explains the mismatch.	The first two conditions in the Description column can be resolved by selecting either REQUEST DELETION to remove the volumes or KEEP IN GROUP. Use the check boxes to select the mismatched volumes that you want to keep or delete. Note that you can modify the type or description of these mismatched volumes, and these changes are saved when you select KEEP IN GROUP. The third condition can be resolved only by a storage administrator, through the storage management subsystem; checking the connection report can help identify the errors that need to be corrected.

PARTITIONS

The **PARTITIONS** tab lists the partitions to which the storage group is attached. If necessary, you can select one or more partitions in the table and detach the storage group from them by selecting **DETACH STORAGE GROUP**. This action can be disruptive when the partitions are in Active state, are in Paused state, or are using a volume in the storage group as a boot volume. In such cases, a warning, error, or informational message is displayed and you are prompted to confirm the detachment.

The Partitions table on this tab includes the following information: partition name, current status, and a user-supplied description, if any.

- For an FCP storage group only, each row in the Partitions table can be expanded to show details about the host bus adapters (HBAs) that the partition is using to access storage. If necessary, you can select **Change Adapter** to review the assigned adapters and remove or replace them with other adapters that are available for use by a partition. You can change an adapter only when an HBA with a backing adapter is available.
- You also have the option of changing the device number for HBAs. If you change the device number, specify a four-character hexadecimal device number in the range 0000 - ffff. For

FICON storage groups, device numbers must be unique within a storage group and across all attached partitions and partition resources. For FCP storage groups, the device number must be unique only within the partition. For example, this volume device ID cannot be the same as the device ID that a partition is using for a network connection (network interface card, or NIC).

ADAPTERS

The **ADAPTERS** tab lists the adapters that are assigned to the storage group. The adapter attributes include the name, ID, type, location, and current allocation.

- If adapters are not yet assigned, the **ADAPTERS** tab has a pending icon (🟡) next to it, and the Adapters table is empty. The total at the foot of the Adapters table lists how many adapters are assigned to the storage group.
- If an existing adapter becomes incomplete, the **ADAPTERS** tab has an incomplete icon (🔴) next to it, and the table entry for that adapter is highlighted with the incomplete icon. If one or more adapters are incomplete, the fulfillment status of the storage group is Incomplete.

WWPN (displayed for FCP storage groups only)

The **WWPN** tab lists the worldwide port numbers (WWPNs) that are available for use. This tab is displayed only for an FCP storage group. The tab display contains two tables: one table that lists each WWPN that is in use by a partition, and one table that lists unused WWPNs. The tables include the following information: the WWPN, its current state and, if the WWPN is in use, the name of the associated partition and host bus adapter (HBA).

HISTORY

The **HISTORY** tab lists the actions that users have taken for this storage group. The most recent action is listed at the top of the History table. Information in the ACTION column not only briefly describes the activity, but also preserves details such as requests that were sent to storage administrators for fulfillment.

- When the storage group is attached to one or more partitions, some changes that are made to the storage group can require corresponding changes on the operating system that the partition hosts. In this case, use the link in the ACTION column entry to display a list of required updates.
- If the storage group is deleted, you can access the history details only for the next 30 days, by using the HMC Web Services API for DPM.

Attaching storage resources to a partition

Review this topic to become familiar with the steps to attach one or more storage groups to a partition on a DPM R3.1 system. To attach a storage group to a partition, use either the **New Partition** task or the **Partition Details** task.

Before you attach a storage group to a partition

- To attach storage groups to an existing partition, use Storage section of the **Partition Details** task. To attach them to a new partition, use the Storage page or section of the **New Partition** task in basic or advanced mode. For a comparison of these modes, see “Selecting which New Partition task mode to use” on page 35.
- To open the **Partition Details** task or the **New Partition** task, use either the default SYSPROG user ID or a user ID that a system administrator authorized to this task through customization controls in the **User Management** task. (You can use other default user IDs to open only the **Partition Details** task; for details, see the online help for the task.)
- If the partition has or will have reserved resources, you can attach only fulfilled storage groups to it. To determine whether the partition has reserved resources, go to the General page of the **Partition Details** task or advanced-mode **New Partition** task, and look at the **Reserve resources** check box. (You cannot reserve partition resources through the basic-mode **New Partition** task.) After the storage groups are attached, no one can modify them.

An overview of the attachment process

The Storage section contains a Storage Groups table with controls that you can use to attach storage groups. The Storage Groups table is displayed only when the DPM R3.1 storage management feature is applied on the system.

Step 1. Open the Attach Storage Groups window

Select the plus icon in the Storage Groups table toolbar to open the Attach Storage Groups window.

Step 2. Review and select one or more available storage groups

On the Attach Storage Groups window, select one or more storage groups listed in the Storage Groups table to attach to this partition. You can select FCP storage groups, FICON storage groups, or both types of storage group.

- The suggested practice is to select storage groups that are in the Complete fulfillment state, but you can select any storage group except for those with a fulfillment state of Incomplete, or those that are already attached to the maximum number of partitions. If you do select groups in states other than Complete, some storage might not be available for use when you start the partition.
- Use the additional information in the Storage Groups table, as necessary, to decide which storage groups to attach. This information includes not only the storage group name, type, and fulfillment state, but also the number of partitions to which the storage group is already attached (if any), and the total capacity.

For FCP storage groups only, you can expand the table entry to show the system-generated host bus adapters (HBAs) and their assigned adapters. You can change the device numbers that DPM automatically assigned to the HBAs when you selected the FCP storage group. The expanded display also includes a link through which you can open the FCP adapter assignment window, and remove or replace the adapters that DPM automatically assigned to the HBAs.

Step 3. Save the partition definition

1. When you have finished selecting storage groups to attach, select **OK** to close the Attach Storage Groups window.
2. Check the entries for the storage groups that you selected, which are now displayed in the Storage Groups table in the Storage section. If necessary, you can use the minus icon in the table toolbar to remove a storage group from the table.
3. When you have finished, review another section or click **OK** to save the partition definition. Note that an operating system or hypervisor image can reside on a storage group volume; if you want to boot the image from a storage group, go to the Boot section of the **Partition Details** task or **New Partition** task. Choose the **Storage device (SAN)** boot option, and select a boot volume in an attached storage group.

Results of the attachment process

DPM saves the partition definition. For FICON storage groups, DPM assigns device numbers, avoiding conflicts with any other device numbers for network, accelerator, and crypto adapters.

What to do next

- If the partition is new or not active, use the **Start** task to start the partition.
- After starting a new partition or modifying the storage resources for an active or stopped partition, ensure that storage groups are visible to the Linux operating system. The action you take depends on whether the partition is new (started for the first time), or existing (active or stopped).

For a new partition

After you start the new partition, you might need to enter Linux commands to make the storage groups available to the operating system that the partition hosts. The actions required

depend on the fulfillment state and type of the storage group, and whether the storage group contained the boot volume for the operating system. Typically, the operating system stores the FCP HBA or FICON volume configuration so it can automatically bring the devices online on the next reboot, so you need to take action only for the initial boot of the operating system.

When attaching a storage group in Complete state

- For an FCP storage group:
 - If the storage group contained the boot volume, the operating system brings online all of the HBAs for this storage group, and all volumes in the storage group are available. No action is required unless you have attached other storage groups.
 - If the storage group does not contain the boot volume, and the operating system is not configured to bring HBAs online automatically, you need to issue the **chccwdev** command to bring online all of the HBAs.
- For a FICON storage group, the operating system brings online only the boot volume. You need to issue the **chccwdev** command to bring online all of the remaining volumes in the storage group that contains the boot volume, as well as the volumes in any other storage groups that you attached.

When attaching an unfulfilled storage group that becomes Complete as the partition is running

- For an FCP storage group:
 - If adapters were assigned to HBAs while the partition is running, you need to use the **chchp** command to activate the channel paths for those new adapters.
 - To access the volumes in the storage group, you need to issue the **chccwdev** command to bring online all of the HBAs.
- For a FICON storage group:
 - If the adapters connecting the storage group to the storage subsystem were assigned while the partition is running, use the **chchp** command to activate the channel paths for those new adapters.
 - All volumes are offline. You need to issue the **chccwdev** command to bring online all of the volumes in the storage group.

For an existing partition

If the partition is running, or when you restart a stopped partition, you might need to enter Linux commands to make any newly attached storage groups available to the operating system that the partition hosts. The actions required depend on the fulfillment state and type of the storage group, and whether the storage group contained the boot volume for the operating system.

When attaching a storage group in Complete state when the partition is stopped

- For an FCP storage group:
 - If the storage group contained the boot volume, the operating system brings online all of the HBAs for this storage group, and all volumes in the storage group are available. No action is required unless you have attached other storage groups.
 - If the storage group does not contain the boot volume, and the operating system is not configured to bring HBAs online automatically, you need to issue the **chccwdev** command to bring online all of the HBAs.
- For a FICON storage group, the operating system brings online only the boot volume. You need to issue the **chccwdev** command to bring online all of the remaining volumes in the storage group that contains the boot volume, as well as the volumes in any other storage groups that you attached.

When attaching a Complete storage group to a running partition, or attaching an unfulfilled storage group that becomes Complete as the partition is running

- For an FCP storage group:
 - If adapters were assigned to HBAs while the partition is running, you need to use the **chchp** command to activate the channel paths for those new adapters.
 - To access the volumes in the storage group, you need to issue the **chccwdev** command to bring online all of the HBAs.
- For a FICON storage group:
 - If the adapters connecting the storage group to the storage subsystem were assigned while the partition is running, use the **chchp** command to activate the channel paths for those new adapters.
 - All volumes are offline. You need to issue the **chccwdev** command to bring online all of the volumes in the storage group.

To find the IDs that you need to use for the Linux commands, use the following tasks.

- HBA device numbers are available in the Host Bus Adapters (HBA) table when you expand the storage group table entry in the Storage section of the **Partition Details** task.
- Channel path IDs for FCP adapters are shown in the Host Bus Adapters (HBA) table when you expand the storage group table entry in the Storage section of the **Partition Details** task.
- Channel path IDs for FICON adapters are shown on the **ADAPTERS** tab of the Storage Group details; open the **Configure Storage** task and select the storage group in the **Storage Overview** to open the Storage Group details page.
- FICON volume device numbers are shown on the **VOLUMES** tab of the Storage Group details page; open the **Configure Storage** task and select the storage group in the **Storage Overview** to open the Storage Group details page.

Modifying the storage configuration of the system

After the initial storage configuration is completed, system and storage administrators can use the **Configure Storage** task to modify that configuration. Depending on what you need to change, open the task and select either **STORAGE CARDS** or **FICON CONNECTIONS**. These links open to pages that are almost identical to those of the Connect to Storage wizard, through which administrators initially configure storage for a DPM R3.1 system.

Use the following information for planning purposes. Note that these details are also available in the online help for the **Configure Storage** task.

Planning considerations for modifying installed storage adapter cards

The process of modifying installed storage adapters is very similar to Step 1 in “An overview of the initial storage configuration process” on page 120.

- To use the **Configure Storage** task to modify storage adapter cards, open the task using the default SYSPROG or SERVICE user IDs, or any user IDs that an access administrator has authorized to the **Configure Storage** task through customization controls in the **User Management** task. Also, you must have access permission to all storage adapters.
- Note that DPM does not allow the reconfiguration of any FCP adapter card that is already in use for a storage group, or any FICON adapter that is used in a path to a logical control unit (LCU). Through controls on the Configure Storage Cards page, you can determine which storage groups are using an adapter that cannot be reconfigured (adapter cards that are in use are shown with disabled label selections for changing the current configuration).
- When you open the Storage Group details page for an FCP storage group, you can review the assigned adapters and remove or replace them with other adapters that are available for use by a partition. However, if an FCP adapter is configured while the storage group is attached to an active partition, DPM cannot detect and list the new adapter as available for use by any

partition. To make sure that you can choose a new adapter from a complete list of available adapters, use the instructions in “Instructions: Changing the adapters that are assigned to an FCP storage group” on page 63.

Planning considerations for modifying FICON connections

The process of modifying installed storage adapters is very similar to Step 2 in “An overview of the initial storage configuration process” on page 120.

- To use the **Configure Storage** task to modify FICON connections, open the task using the default SYSPROG, STORAGEADMIN, or SERVICE user IDs, or any user IDs that an access administrator has authorized to this task through customization controls in the **User Management** task. Also, you must have access permission to all FICON adapters.
- If you need help to complete this configuration, you can use the **Invite** link on the Configure FICON Connections page to notify a co-worker about the remaining configuration tasks. DPM automatically generates an invitation that you can send, to which you can add your own greeting and more details, if necessary. For integrated invitations and notifications, users must have an email address associated with their user IDs, and Simple Mail Transfer Protocol (SMTP) settings must be defined.
 - Email addresses for users are assigned through the **User Management** task.
 - The SMTP server and port settings are defined through the **Monitor System Events** task.
- If you are modifying existing fabrics and switches, the following rules apply.
 - You cannot modify the high integrity fabric setting when LCUs are in use for switches in the fabric. If you need to modify the high integrity fabric setting, you must first remove all of the paths from the LCUs that are using that fabric, and then modify the high integrity fabric setting. After modifying the setting, you can restore the paths by adding them back to the LCUs.
 - You can delete a fabric from a site only when the fabric does not contain any switches.
 - You can delete a switch only when all of its existing connections (to an adapter port or to a storage subsystem) are not configured in any storage group.
 - You can disconnect a switch only when both of the following conditions are true.
 - No storage groups are using the switch port.
 - All LCUs that are using the switch port are not being used to fulfill a storage group.
- You can delete a storage subsystem from a site only when that storage subsystem does not contain any LCUs.
- If you are modifying existing adapter ports, note that you can disconnect an adapter port only when both of the following conditions are true.
 - No storage groups are using the adapter port.
 - All LCUs that are using the adapter port are not being used to fulfill a storage group.
- If you are modifying the paths for an existing LCU, the following rules apply.
 - Although you can add a path to or delete a path from an LCU that is being used by a storage group while it is attached to a partition, you cannot delete the last remaining path from an LCU that is in use.
 - You cannot add a path to an LCU in the primary site if doing so requires the use of a cascaded switch. Only the secondary site can have cascaded switches.

Part 7. Appendixes

Appendix A. DPM task and resource roles

Tasks and resources need to be made available or excluded based on the roles to which they are assigned. You may create your own specific task and resource roles which include specific tasks and resources; however, HMC user management provides default roles for your convenience. Table 10 identifies the DPM tasks along with default task roles. The table identifies the task roles in which a particular task is included. It also documents the resource roles that are required to complete a task.

The Details task has unique behavior with respect to roles. The Details task (view only) is always available for all resources accessible to a user ID. If a user ID has permission for the Details task, through an assigned task role, modifications may be made in the details task. Specific Details task permissions are assigned to default task roles as shown in Table 10.

For example consider user ID SIGMUND. SIGMUND has been given the Defined System Managed Objects resource role but not the System Programmer (SP) task role, which contains the Partition Details task. SIGMUND will still be able to launch details task for a Partition, but the content of the task will be displayed read-only such that SIGMUND cannot modify to the resource.

All tasks that can be launched from the HMC workspace are marked in **bold**. Where there are both administrative and operator roles, such as SP and OP, any permissions given to the operator are also available for the administrator role. See Table 9 for the tasks mapping legend.

Table 9. DPM task roles mapping legend

Legend	Description
AA	Access Administrator Tasks
SP	System Programmer Tasks
OP	Operator Tasks
AOP	Advanced Operator Tasks
SER	Service Representative Tasks
SA	Storage Administrator Tasks
X	Required role to perform a task.
O	At least one of the roles is required to perform a task.
*	Denotes a task that is available through the Support Element (SE) only.

Table 10. DPM task roles mapping

DPM Tasks	Task Roles					
	AA	SP	OP	AOP	SER	SA
Configure Storage		O			O	O ¹
Delete Partition		O			O	
Disable Dynamic Partition Manager*					X	
Dump (Partition)		O	O	O	O	
Enable Dynamic Partition Manager*					X	
Getting Started with Dynamic Partition Manager	O	O	O	O	O	
Manage Processor Sharing		O			O	
Manage Adapters		O	O	O	O	

Table 10. DPM task roles mapping (continued)

DPM Tasks	Task Roles					
- Adapter Details		O			O	
- Create HiperSockets Adapter		O			O	
- Delete HiperSockets Adapter		O			O	
- Reassign Channel Path IDs		O			O	
- Reassign Devices		O			O	
- Export WWPNS		O	O	O	O	
New Partition (basic mode)		O			O	
New Partition (advanced mode)		O			O	
- Controls		O			O	
Partition Details		O			O	
- Controls		O			O	
Start (start a single DPM system)		O			O	
Start (start one or more DPM partitions)		O	O	O	O	
Stop (stop a single DPM system)		O			O	
Stop (stop one or more DPM partitions)		O	O	O	O	
System Details		O			O	
- Configure System Management (OSM) Adapters*					X	

Note:

1. During the initial storage setup for a DPM-enabled system, the storage administrator does not have access to the **Configure Storage** task until the system administrator configures storage cards for the system.

Appendix B. Notices

This information was developed for products and services offered in the US.

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

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.

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 is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains 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 enterprise is entirely coincidental.

Trademarks

IBM, the IBM logo, and ibm.com® are 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 “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

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

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Class A Notices

The following Class A statements apply to this IBM product. The statement for other IBM products intended for use with this product will appear in their accompanying manuals.

Federal Communications Commission (FCC) Statement

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Industry Canada Compliance Statement

This Class A digital apparatus complies with Canadian ICES-003.

Avis de conformité à la réglementation d'Industrie Canada

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

European Community Compliance Statement

This product is in conformity with the protection requirements of EU Council Directive 2014/30/EU on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to European Standard EN 55032. The limits for Class A equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

European Community contact:
IBM Deutschland GmbH
Technical Regulations, Department M372
IBM-Allee 1, 71139 Ehningen, Germany
Tele: +49 (0) 800 225 5423 or +49 (0) 180 331 3233
email: halloibm@de.ibm.com

Warning: This is a Class A product. In a domestic environment, this product may cause radio interference, in which case the user may be required to take adequate measures.

VCCI Statement - Japan

この装置は、クラス A 情報技術装置です。この装置を家庭環境で使用する
と電波妨害を引き起こすことがあります。この場合には使用者が適切な対策
を講ずるよう要求されることがあります。 VCCI-A

The following is a summary of the Japanese VCCI statement above:

This is a Class A product based on the standard of the VCCI Council. If this equipment is used in a domestic environment, radio interference may occur, in which case the user may be required to take corrective actions.

Japan JIS C 61000-3-2 Compliance

(一社) 電子情報技術産業協会 高調波電流抑制対策実施
要領に基づく定格入力電力値 : Knowledge Centerの各製品の
仕様ページ参照

For products less than or equal to 20 A per phase, the following statement applies:

高調波電流規格 JIS C 61000-3-2 適合品

For products greater than 20 A, single-phase, the following statements apply:

高調波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
--

回路分類：6（単相、P F C回路付）

換算係数：0

For products greater than 20 A per phase, three-phase, the following statements apply:

高調波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
--

回路分類：5（3相、P F C回路付）

換算係数：0

Electromagnetic Interference (EMI) Statement - People's Republic of China

声 明

此为 A 级产品,在生活环境中,该产品可能会造成无线电干扰。在这种情况下,可能需要用户对其干扰采取切实可行的措施。
--

Declaration: This is a Class A product. In a domestic environment, this product may cause radio interference, in which case the user may need to perform practical action.

Electromagnetic Interference (EMI) Statement - Taiwan

警告使用者:

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

The following is a summary of the Taiwan EMI statement above:

Warning: This is a Class A product. In a domestic environment, this product may cause radio interference, in which case the user will be required to take adequate measures.

IBM Taiwan Contact Information:

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

Electromagnetic Interference (EMI) Statement - Korea

이 기기는 업무용(A급)으로 전자파적합등록을 한 기기이오니
판매자 또는 사용자는 이 점을 주의하시기 바라며, 가정외의
지역에서 사용하는 것을 목적으로 합니다.

Germany Compliance Statement

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55032 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

EN 55032 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden:

"Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem "Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)". Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC EG Richtlinie 2014/30/EU) für Geräte der Klasse A

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Regulations, Abteilung M372

IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 (0) 800 225 5423 or +49 (0) 180 331 3233
email: halloibm@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55032 Klasse A.

Electromagnetic Interference (EMI) Statement - Russia

ВНИМАНИЕ! Настоящее изделие относится к классу A.
В жилых помещениях оно может создавать радиопомехи, для
снижения которых необходимы дополнительные меры

Index

A

accessibility xiii
 contact IBM xiii
 adapter
 auto-configuration
 prerequisites 74
 definition 3
 specifying port number 52
 assistive technologies xiii
 attach a storage group 53

B

boot option
 for Linux operating system 24
 for z/VM hypervisor 24
 network server 73, 101
 PXE 73, 101
 Secure Service Container 24
 boot options
 for operating systems 24
 hypervisor 24
 operating system 24

C

commands
 Linux operating system 52
 z/VM hypervisor 52
 concepts
 hypervisors 23
 operating systems 23
 configure
 FCP adapter card
 initial setup 120
 FICON connection
 initial setup 120
 Configure Storage task
 benefits of using 113
 connect system to storage
 hardware 133
 initial setup 120
 planning considerations 120
 prerequisites 120
 task overview 120
 introduction 113, 118, 119, 125
 prerequisites 73
 request storage
 planning considerations 123
 prerequisites 123
 task overview 123
 required authorization
 procedure 119
 view storage group
 map FICON volume 126
 modify storage group 126
 view connection report 126
 connection report
 for storage group 126
 corequisite software 77

create a partition
 comparison of New Partition task
 modes 35
 overview 5
 selecting New Partition task
 mode 35
 step-by-step instructions 37
 switching New Partition task
 modes 36

D

detach a storage group 53
 DPM
 See also Dynamic Partition Manager
 feature code 71
 prerequisites 71
 upgrade 71
 DPM R3.1
 prerequisites 73
 storage management 113
 DPM version
 upgrading to DPM R3.1 89
 DPM-enabled system
 supported types 71
 DVDPRIME command 24
 Dynamic Partition Manager
 See also DPM
 advantages 4
 definition 3
 enabling for a system 3
 summary of tasks 65
 user interface 3

E

echo command 52
 engineering change
 See MCL
 engineering change (EC) 71

F

FCP adapter card
 change adapter assignment 126
 configure
 initial setup 120
 modify 133
 FCP storage 113
 feature code
 for Dynamic Partition Manager 71
 for OSA-Express 1000BASE-T
 Ethernet 71
 FICON adapter card
 configure
 initial setup 120
 modify 133
 FICON connection
 modify 133

FICON connections
 configure
 initial setup 120
 FICON storage 113
 FICON volume
 mapping 126
 file transfer protocol
 See FTP, FTPS, SFTP
 FTP
 compared to FTPS or SFTP 105
 FTP Secure
 See FTPS
 FTPS
 boot option 24
 prerequisites 74
 setting up on HMC 105

H

Hardware Management Console
 list of Dynamic Partition Manager
 tasks 65
 hardware planning 77
 HBA conversion
 to storage group 82, 89
 HMC
 See also Hardware Management
 Console
 EC MCLs for DPM version 71
 how to get started 5
 hypervisor
 partition type 23

I

Integrated 3270 Console task 27
 Integrated ASCII Console
 web service API 73
 Integrated ASCII Console task 27

K

keyboard
 navigation xiii
 KVM on IBM z
 required version 76

L

Linux operating system
 boot options 24
 commands 52
 Linux operating system 22
 lschp 51
 lscpu 50
 lscss 51
 lsmem 50
 lsqeth 51
 concepts 23

Linux operating system (*continued*)

- device drivers for adapters 22, 52
- Integrated ASCII Console task 27
- link to library 27
- lscpu command 50
- lsmem command 50
- sample configuration 11
- specifying port number 52
- verifying partition resources 50

Linux system

- required version 76

lschp command 51

lscss command 51

lsqeth command 51

M

MAC address

- specify for a NIC 107

machine change level (MCL) 71

media access control address 107

memory

- dynamically add 53

MES frame upgrade

- migration instructions 79
- supported migration options 79

migrate

- DPM configuration 74
- DPM environment 73

migration considerations

- from prior system to new system
- supported migration options 79

HBA conversion

- to storage group 82
- system replacement 82

migration options

- MES frame upgrade 79

miscellaneous equipment specification

See MES frame upgrade

modify

- FCP adapter card 133
- FICON connection 133
- partition resources 53

N

navigation

- keyboard xiii

network adapter

- dynamically add 53

network boot server 73, 101

New Partition

- step-by-step instructions 37

New Partition task

- attach storage
- planning considerations 130
- comparison of modes 35
- selecting the mode 35
- switching modes 36

O

operating system

- partition type 23

OSA-Express 1000BASE-T Ethernet

- feature code 71

P

partition

- boot options 24
- configuration settings 12
- definition 3
- how to start 48
- lifecycle 15
- modify definition 53
- planning checklist 31
- properties 12
- sample configuration 11
- type
- supported operating system or hypervisor 23
- verifying resources in Linux 50

partition definition

- creating through New Partition task 12

Partition Details task

- add adapter 53
- add memory 53
- add processors 53
- attach storage
- planning considerations 130
- attach storage group 53
- reattach storage group 53

planning checklist

- for creating a partition 31

port number

- specifying in Linux 52

preboot execution environment (PXE) 101

prerequisites

- DPM-enabled system 71
- KVM on IBM z 76
- Linux system
- required version 76

Preventative Service Planning buckets

See PSP buckets

processors

- dynamically add 53

PSP buckets

- corequisite software 77
- hardware planning 77
- where to find 77

PXE 101

R

reattach a storage group 53

revisions xv

S

SE

- See also Support Element
- EC MCLs for DPM version 71

Secure File Transfer Protocol

See SFTP

Secure Service Container

- boot option 24
- concepts 23
- definition 11
- prerequisites for using 75

SFTP

- boot option 24

SFTP (*continued*)

- prerequisites 74
- setting up on HMC 105

shortcut keys xiii

start a partition

- step-by-step instructions 48

storage

- attach to partition 130
- request for partition use
- with or without template 123

storage group

- attach 53
- attach to partition 130
- attributes 123
- change FCP adapter assignment 126
- created from HBAs during migration 82
- created from HBAs during upgrade 89
- definition 123
- detach 53
- modify 126
- reattach 53
- request 123
- view 126
- with or without template 123

storage groups 113

- template 117, 123

storage management

- starting with DPM R3.1 113
- support for FICON storage 113

storage management feature

- prerequisites 73

STORAGEADMIN default user ID 119

Support Element

- list of Dynamic Partition Manager tasks 65

system

- connect to storage
- initial setup 120
- modify 133
- types that support DPM 71

system replacement

- migration instructions 82
- required MCLs 82

T

template

- for storage group 117, 123

U

upgrade

- DPM 71

upgrading DPM version

- to DPM R3.1 89

UUID

- use in network boot server 101

V

version

- EC MCL list 71

virtual local area network

- specify an identifier 107

VLAN enforcement 107
 prerequisites 74
VLAN identifier (ID) 107

Z

z/VM
 boot options 24
 concepts 23
 DVDPRIME command 24
 installation 24
 Integrated 3270 Console task 27
 link to library 27
 RAMDISK 24
z/VM hypervisor
 commands 52
 sample configuration 11



Printed in USA

SB10-7170-02

