

IBM Z

*Dynamic Partition Manager
(DPM) Guide*

Level 04a



Note:

Before you use this information and the product it supports, read the information in “[Safety](#)” on page xiii, [Appendix E, “Notices,”](#) on page 177 and *IBM Systems Environmental Notices and User Guide*, Z125-5823.

This edition, SB10-7176-04a, applies to IBM Z and IBM LinuxONE servers. This edition replaces SB10-7176-04.

There might be a newer version of this document in a **PDF** file available on **Resource Link**.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 International Business Machines Corporation 2019, 2021.**

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

Contents

| | |
|--|-------------|
| Figures..... | vii |
| Tables..... | xi |
| Safety..... | xiii |
| Safety notices..... | xiii |
| World trade safety information..... | xiii |
| Laser safety information..... | xiii |
| Laser compliance..... | xiii |
| About this publication..... | xv |
| Intended audience..... | xv |
| Prerequisite and related information..... | xv |
| Related HMC and SE console information..... | xvi |
| Related information for Linux on Z..... | xvi |
| Related information for z/VM..... | xvi |
| Related information for IBM Secure Service Container..... | xvi |
| How to use this publication..... | xvi |
| Accessibility..... | xvii |
| Accessibility features..... | xvii |
| Keyboard navigation..... | xvii |
| Consult assistive technologies..... | xviii |
| IBM and accessibility..... | xviii |
| How to send your comments..... | xviii |
| Summary of changes..... | xix |
| 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 or Linux system..... | 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 or Linux system..... | 11 |
| Chapter 4. Adapters: Connections to networks, storage, and more..... | 17 |
| Network adapters..... | 17 |
| Storage adapters..... | 18 |
| Accelerators..... | 23 |
| Crypto adapters..... | 23 |
| Chapter 5. Operating systems and hypervisors..... | 25 |
| Partition types and supported operating systems and hypervisors..... | 25 |
| Installing an operating system or hypervisor in a partition..... | 26 |
| Working with an installed operating system or hypervisor..... | 29 |

| | |
|---|------------|
| Part 3. Basic tasks for Linux administrators..... | 31 |
| Chapter 6. Planning checklist for creating a partition..... | 33 |
| Chapter 7. Creating and starting a new partition..... | 37 |
| Selecting which New Partition task mode to use..... | 37 |
| Creating a new partition..... | 39 |
| Starting a partition and its operating system or hypervisor..... | 49 |
| Chapter 8. Dynamically modifying the resources of an active partition..... | 55 |
| Chapter 9. Summary of tasks for managing systems, adapters, and partitions..... | 71 |
| Part 4. Topics for system planners..... | 75 |
| Chapter 10. Prerequisites for using Dynamic Partition Manager..... | 77 |
| Chapter 11. Migrating DPM configuration data to a new system..... | 87 |
| Migration instructions for an MES frame upgrade..... | 87 |
| Migration instructions for a system replacement..... | 91 |
| Chapter 12. Adapter configuration..... | 97 |
| Chapter 13. Secure Execution for Linux feature..... | 101 |
| Chapter 14. Enabling FCP tape support for machine types 3906 and 3907..... | 103 |
| Part 5. Topics for network administrators..... | 105 |
| Chapter 15. Customizing the network boot server environment..... | 107 |
| Chapter 16. Securing FTP operations..... | 111 |
| Chapter 17. Using VLAN IDs to secure network communications..... | 113 |
| Part 6. Topics for storage administrators..... | 117 |
| Chapter 18. The Configure Storage task: a simplified and improved approach to storage management..... | 119 |
| A simplified, visual process for connecting storage to one system..... | 119 |
| Integrated requests and notifications..... | 122 |
| Access to FICON and FCP disk storage through storage groups..... | 123 |
| Access to FCP tape storage through tape links..... | 124 |
| Access to NVMe storage through storage groups..... | 126 |
| Storage group templates..... | 127 |
| Automatic detection and assignment of FICON and FCP adapters..... | 127 |
| Device number allocation for FICON storage groups..... | 127 |
| Storage Overview and details..... | 128 |
| Chapter 19. Setting up storage resources for the first time..... | 129 |
| Authorizing users of the Configure Storage task..... | 129 |
| Connecting a system to external storage devices in the SAN..... | 130 |
| Requesting FICON or FCP storage for partitions to use..... | 134 |
| Optimize 2nd level virtualization and share the same FCP disks across partitions..... | 136 |
| Requesting NVMe storage for partitions to use..... | 137 |
| Requesting tape storage for partitions to use..... | 139 |

| | |
|---|------------|
| Chapter 20. Managing storage resources after the initial configuration is complete..... | 143 |
| Viewing and managing storage groups and tape links..... | 143 |
| The Storage Group details page..... | 144 |
| The Tape Link details page..... | 148 |
| Attaching storage resources to a partition..... | 152 |
| Modifying the storage configuration of the system..... | 156 |
| Viewing and managing FCP tape libraries..... | 158 |
| Appendix A. DPM task and resource roles..... | 161 |
| Appendix B. Sample System Cabling Plan..... | 165 |
| Appendix C. Sample NVMe adapter plan..... | 167 |
| Appendix D. Configuring and accessing FCP tape drives on a DPM-enabled system..... | 169 |
| Configuring and accessing tape drives through a shared fabric..... | 169 |
| Configuring and accessing tape drives through a dedicated fabric..... | 172 |
| Appendix E. Notices..... | 177 |
| Trademarks..... | 177 |
| Class A Notices..... | 178 |
| Index..... | 183 |

Figures

| | |
|--|----|
| 1. The HMC Welcome page..... | 3 |
| 2. Portion of the main HMC page, with the Systems Management node expanded, and a system and the New Partition task highlighted..... | 5 |
| 3. Welcome page of the New Partition task in basic mode..... | 6 |
| 4. Methods of starting a new partition..... | 7 |
| 5. Portion of the Start window for a new partition named ASL_Test..... | 7 |
| 6. Sample Linux command and display of configured processor resources..... | 8 |
| 7. Partitions configured on a DPM-enabled system..... | 12 |
| 8. Partitions with NICs configured for access to network adapters on a DPM-enabled system..... | 18 |
| 9. Sample physical FICON and FCP storage configuration..... | 20 |
| 10. Sample FICON and FCP storage groups that are attached to partitions on the system..... | 21 |
| 11. Sample FCP tape links that are attached to partitions on the system..... | 22 |
| 12. Example: Linux servers configured in partitions on a DPM-enabled system..... | 25 |
| 13. Welcome page of the New Partition task in basic mode..... | 40 |
| 14. Processors page of the New Partition task in basic mode..... | 41 |
| 15. Memory page of the New Partition task in basic mode..... | 42 |
| 16. Network page of the New Partition task in basic mode..... | 43 |
| 17. Boot page of the New Partition task in basic mode..... | 46 |
| 18. Sample displays resulting from the lscpu command..... | 51 |
| 19. Sample displays resulting from the lsmem command..... | 52 |
| 20. Sample displays resulting from the lsqeth and lscss commands..... | 52 |
| 21. Sample displays resulting from the lschp command..... | 53 |
| 22. Sample displays resulting from the lspci command..... | 53 |

| | |
|---|-----|
| 23. Sample displays resulting from the lstage command..... | 53 |
| 24. Sample display for the lscpu command showing current processor resources..... | 56 |
| 25. Sample screen of the Processors section in the Partition Details window..... | 56 |
| 26. Sample display for the lscpu command showing additional processor resources..... | 57 |
| 27. Sample display for the lscpu command showing configured processor resources..... | 57 |
| 28. Sample display for the lsmem command..... | 58 |
| 29. Sample screen of the Memory section in the Partition Details window..... | 58 |
| 30. Sample display for the chmem command..... | 58 |
| 31. Sample display for the lscss command showing current devices..... | 59 |
| 32. Sample screen of the Network section with the New NIC icon highlighted..... | 59 |
| 33. Sample screen of the New Network Interface Card window..... | 60 |
| 34. Sample screen of the Network section with the newly added NIC..... | 60 |
| 35. Sample display for the znetconf command showing current devices..... | 61 |
| 36. Sample display for the lsqeth command showing details for devices..... | 62 |
| 37. Sample hardware message that shows the results of a successful DPM conversion of migrated configuration information..... | 89 |
| 38. Partition access to storage when migrating from an early DPM version to DPM R3.1 or later..... | 93 |
| 39. Sample screen of the Manage Adapters task..... | 98 |
| 40. Sample screen of the Adapter Details window, with Name and Description fields highlighted..... | 99 |
| 41. Sample screen of the Manage Adapters task, with the expanded Actions list | 99 |
| 42. Sample screen of the New HiperSockets Adapter window | 100 |
| 43. Sample screen of the new HiperSockets adapter in the Adapter Details window..... | 100 |
| 44. Sample VLAN configuration in partitions and guest operating systems..... | 113 |
| 45. Sample frame and card display on the Configure Storage Cards page..... | 120 |
| 46. Photograph of a physical adapter card in a system I/O drawer..... | 120 |
| 47. Overview map and other viewport navigational controls..... | 121 |

| | |
|---|-----|
| 48. Sample display of two sites with physical storage hardware configured for the primary site..... | 122 |
| 49. Sample FICON and FCP storage groups that are attached to partitions on the system..... | 123 |
| 50. A tape link connects one or more partitions to one FCP tape library in the SAN..... | 125 |
| 51. Sample frame and card display on the Configure Storage Cards page..... | 131 |
| 52. Sample display of two sites with physical storage hardware configured for the primary site..... | 132 |
| 53. A sample SAN configuration with a shared fabric for access to FCP tape storage drives and other storage devices..... | 170 |
| 54. A sample SAN configuration with a dedicated fabric for access to FCP tape storage..... | 173 |

Tables

| | |
|--|-----|
| 1. Planning checklist for defining your new partition..... | 33 |
| 2. Comparison of editable partition properties in the basic and advanced New Partition task modes..... | 37 |
| 3. Summary of key HMC/SE tasks and displays for working with DPM-enabled systems, partitions, and adapters..... | 71 |
| 4. Required feature codes for DPM..... | 77 |
| 5. Summary of DPM versions..... | 78 |
| 6. Supported DPM migration paths for MES frame upgrades..... | 88 |
| 7. Supported DPM migration paths for a system replacement..... | 91 |
| 8. Sample default names for installed adapters..... | 97 |
| 9. Effect of fulfillment status on the Volume tab display..... | 146 |
| 10. Interpreting the ENVIRONMENT tab display for different tape link fulfillment states..... | 149 |
| 11. DPM task roles mapping legend..... | 161 |
| 12. DPM task roles mapping..... | 161 |
| 13. Variables to specify in the Declaration Section of the script..... | 175 |

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), RoCE Express, Integrated Coupling Adapter (ICA SR), 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.

Laser Notice: U.S. FDA CDRH NOTICE if low power lasers are utilized, integrated, or offered with end product systems as applicable. Complies with 21 CFR 1040.10 and 1040.11 except for conformance with IEC 60825-1 Ed. 3., as described in Laser Notice No. 56, dated May 8, 2019.

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 z15™ (z15™): machine types 8561 and 8562
- IBM z14 (z14): machine types 3906 and 3907
- IBM z13® (z13®) or IBM z13s® (z13s®)
- IBM LinuxONE III: machine types 8561 and 8562
- 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 4, which is available starting with Hardware Management Console / Support Element (HMC/SE) Version 2.15.0 on the z15 and LinuxONE III.

- For a list of prerequisites and functions introduced with DPM Version 4, see [Chapter 10, “Prerequisites for using Dynamic Partition Manager,”](#) on page 77.
- For information about the prior DPM Version 3 releases, see *IBM Dynamic Partition Manager Guide*, SB10-7170.

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 supported as a 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 or Linux system 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 or Linux system operators, automation engineers, and experienced system administrators who need to learn how to use DPM tasks.

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 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 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 Documentation, at: <https://www.ibm.com/docs/en/linux-on-systems?topic=linuxone-library-overview>

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 in IBM Documentation at <https://www.ibm.com/docs/en/zvm/7.1>

Related information for IBM Secure Service Container

For information about working with Secure Service Container partitions and the appliances they support, see the appropriate edition of the *Secure Service Container User's Guide*, which is available through the Library link on IBM Resource Link at <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 concepts and terms that might be unfamiliar to system administrators who have little or no experience working on a mainframe or Linux system.

Part 3, “Basic tasks for Linux administrators,” on page 31

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 75

Topics in this part provide the prerequisites for enabling DPM on a mainframe or Linux 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 or Linux systems.

Part 5, “Topics for network administrators,” on page 105

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 117

Topics in this part provide information about storage-related concepts and tasks that are specific to working with DPM-enabled systems. These topics apply to supported types of Fibre Connection (FICON) extended count key data (ECKD) direct-access storage devices (DASD), and Fibre Channel Protocol (FCP) Small Computer System Interface (SCSI) disk storage devices and FCP tape storage devices.

Appendixes

Topics in this part include the default task roles for DPM tasks on the HMC or SE; a sample system cabling plan; a sample Non-Volatile Memory Express® (NVMe) adapter plan; instructions for configuring and accessing FCP tape storage drives; and trademark information.

Accessibility

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 and IBM 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.

Accessibility features

The following list includes the major accessibility features in IBM Z and IBM LinuxONE documentation, and on the Hardware Management Console and Support Element console:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Customizable display attributes such as color, contrast, and font size
- Communication of information independent of color
- Interfaces commonly used by screen magnifiers
- Interfaces that are free of flashing lights that could induce seizures due to photo-sensitivity.

Keyboard navigation

This product uses standard Microsoft Windows navigation keys.

Consult assistive technologies

Assistive technology products such as screen readers function with our publications, the Hardware Management Console, and the Support Element console. Consult the product information for the specific assistive technology product that is used to access the EPUB format publication or console.

IBM and accessibility

See <http://www.ibm.com/able> for more information about the commitment that IBM has to accessibility.

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 thick vertical line to the left of the change.

Summary of changes for SB10-7176-04a

This edition contains updates that describe enhancements to Fibre Channel Protocol (FCP) tape support, which was introduced with DPM R4.3. For details, see the top entry of the versioning table in [Chapter 10, “Prerequisites for using Dynamic Partition Manager,” on page 77](#), and the topic [“The Tape Link details page” on page 148](#).

This edition also contains technical updates to the following topics:

- [“Creating, starting, and managing a partition” on page 15](#)
- [“Example: Adding a new network device” on page 59](#)
- [“Instructions: Attaching storage groups or tape links” on page 62](#)
- [“Migration instructions for a system replacement” on page 91](#)

In this edition, all references to IBM Knowledge Center have been changed to IBM Documentation (<https://www.ibm.com/docs/en>).

Summary of changes for SB10-7176-04

This edition contains the following updates for DPM R4.3, which provides system or storage administrators with a simplified interface in the **Configure Storage** task, through which they can configure access to Fibre Channel Protocol (FCP) tape libraries in the storage area network (SAN). This new support replaces the temporary tape access method described in [Appendix D, “Configuring and accessing FCP tape drives on a DPM-enabled system,” on page 169](#).

New information

- A new table entry that describes the new functions available in DPM R4.3 has been added to the versioning table in [Chapter 10, “Prerequisites for using Dynamic Partition Manager,” on page 77](#). This release of DPM R4.3 applies to 3906 and 3907 systems only.
- A new topic, [Chapter 14, “Enabling FCP tape support for machine types 3906 and 3907,” on page 103](#) has been added to [Part 4, “Topics for system planners,” on page 75](#).
- A new topic, [“Requesting tape storage for partitions to use” on page 139](#), has been added to [Chapter 19, “Setting up storage resources for the first time,” on page 129](#).
- A new topic, [“Viewing and managing FCP tape libraries” on page 158](#), has been added to [Chapter 20, “Managing storage resources after the initial configuration is complete,” on page 143](#).

Changed information

Various topics in the following chapters have been updated to describe the support for FCP tape.

- [Chapter 3, “Partitions: Virtual images of a mainframe or Linux system,” on page 11](#)
- [Chapter 4, “Adapters: Connections to networks, storage, and more,” on page 17](#)
- [Chapter 6, “Planning checklist for creating a partition,” on page 33](#)
- [Chapter 7, “Creating and starting a new partition,” on page 37](#)
- [Chapter 8, “Dynamically modifying the resources of an active partition,” on page 55](#)
- [Chapter 18, “The Configure Storage task: a simplified and improved approach to storage management,” on page 119](#)
- [Chapter 20, “Managing storage resources after the initial configuration is complete,” on page 143](#)

This edition also contains updates that describe enhancements to the **System Details** task to support the use of global and host keys for the IBM Secure Execution for Linux feature on 8561 and 8562 systems. For details, see the top entry of the versioning table in [Chapter 10, “Prerequisites for using Dynamic Partition Manager,”](#) on page 77, and the topic [Chapter 13, “Secure Execution for Linux feature,”](#) on page 101.

Summary of changes for SB10-7176-03

This edition contains the following updates for DPM R4.2.

New information

- A new topic, “NVMe storage adapters” on page 22, has been added to briefly describe Non-Volatile Memory Express (NVMe) storage adapters and how system administrators set them up for use.
- A new table entry that describes the new and improved functions available in DPM R4.2 has been added to the versioning table in [Chapter 10, “Prerequisites for using Dynamic Partition Manager,”](#) on page 77.
- A new topic, “About NVMe storage adapters” on page 84, has been added to [Chapter 10, “Prerequisites for using Dynamic Partition Manager,”](#) on page 77 to briefly introduce this newly supported adapter type.
- A new topic, “Requesting NVMe storage for partitions to use” on page 137, has been added to [Chapter 19, “Setting up storage resources for the first time,”](#) on page 129 to describe how to create NVMe storage groups that administrators can attach to partitions.
- A new topic, [Appendix C, “Sample NVMe adapter plan,”](#) on page 167, provides a sample of the exportable CSV file that is available through the **Configure Storage** task.

Changed information

- Various topics in the following chapters have been updated to describe the support for Non-Volatile Memory Express (NVMe) storage adapters.
 - [Chapter 3, “Partitions: Virtual images of a mainframe or Linux system,”](#) on page 11
 - [Chapter 4, “Adapters: Connections to networks, storage, and more,”](#) on page 17
 - [Chapter 5, “Operating systems and hypervisors,”](#) on page 25
 - [Chapter 6, “Planning checklist for creating a partition,”](#) on page 33
 - [Chapter 7, “Creating and starting a new partition,”](#) on page 37
 - [Chapter 8, “Dynamically modifying the resources of an active partition,”](#) on page 55
 - [Chapter 12, “Adapter configuration,”](#) on page 97
 - [Chapter 18, “The Configure Storage task: a simplified and improved approach to storage management,”](#) on page 119
 - [Chapter 20, “Managing storage resources after the initial configuration is complete,”](#) on page 143

Summary of changes for SB10-7176-02a

The minimum bundle number for DPM R4.1 enablement has been updated in [Table 5](#) on page 78.

Summary of changes for SB10-7176-02

This edition contains the following updates for DPM R4.1.

New information

- The following chapters have been updated to provide information about enabling secure execution, which isolates and protects any guests that run on a hypervisor by restricting host access to guest workloads and data.
 - [Chapter 5, “Operating systems and hypervisors,”](#) on page 25

- [“Starting a partition and its operating system or hypervisor” on page 49](#)
- [Chapter 13, “Secure Execution for Linux feature,” on page 101](#)
- [Appendix A, “DPM task and resource roles,” on page 161](#)
- A new table entry for DPM R4.1 has been added to the versioning table in [Chapter 10, “Prerequisites for using Dynamic Partition Manager,” on page 77](#).

Changed information

- Various topics have been updated with the 8562 machine type for the IBM z15 (z15).
- Topics in [Chapter 11, “Migrating DPM configuration data to a new system,” on page 87](#) have minor editorial updates to account for multiple releases of DPM Version 4.
- [“Modifying the storage configuration of the system” on page 156](#) has been updated to clarify the definition of in-use FCP and FICON storage adapters.

Summary of changes for SB10-7176-01

This edition contains the following updates.

New information

"Configuring and accessing FCP tape drives on a DPM-enabled system" provides instructions for using Fibre Channel Protocol (FCP) tape storage drives.

Changed information

[Table 4 on page 77](#) has been updated to clarify the hardware requirements for enabling DPM on a system.

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 or Linux system,” 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 or Linux system

Linux servers and applications have run on mainframes and Linux 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). The HMC is the user interface for managing mainframes and Linux systems, along with their resources.

DPM is a configuration manager that is designed for setting up and managing Linux servers that run on a mainframe or Linux system. On a DPM-enabled system, the runtime environment for your Linux server is called a *partition*. A partition is also the runtime environment for a hypervisor and its guest operating-system images. 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 or Linux 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 and Linux systems, as on other platforms, an *adapter* is a physical device that connects the system to other computers or devices. In contrast to other platforms, adapters on a mainframe or Linux system can be shared between partitions, which reduces the amount of adapters that might be required to handle a specific workload.

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 on page 3](#) shows the **HMC Welcome** page, which provides a visual summary of the number and status of elements that are managed through the HMC: systems, the partitions that run on those systems, and the adapters that are configured for those systems.

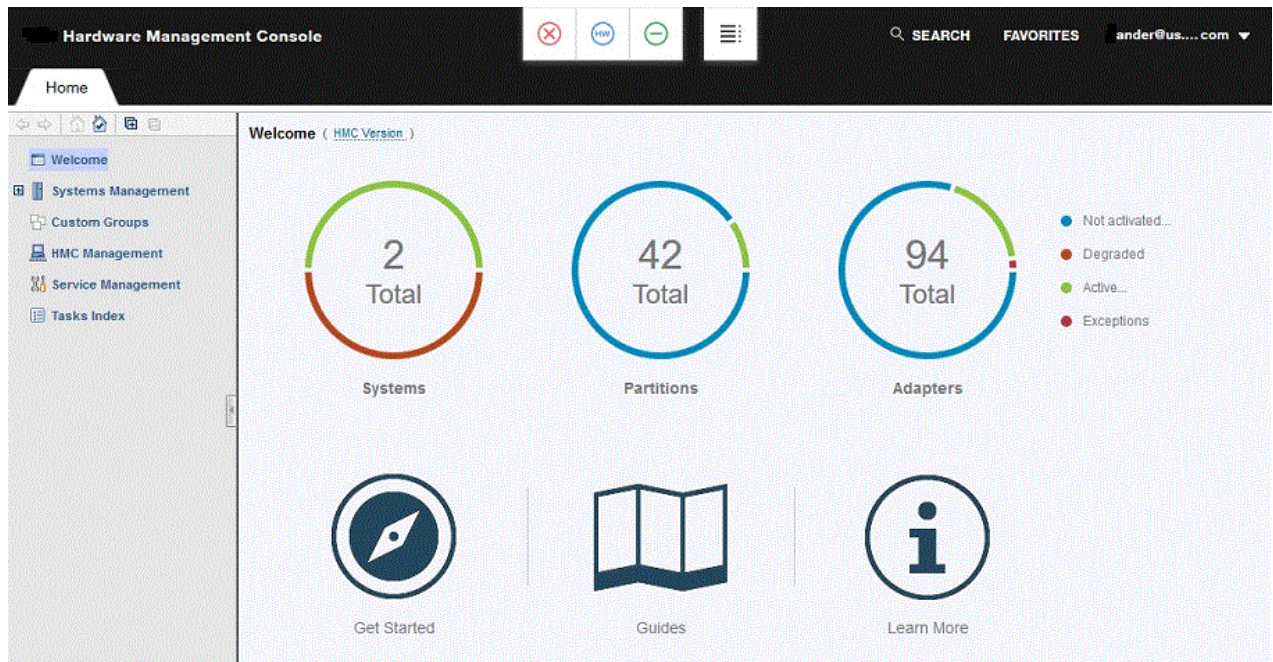


Figure 1. The HMC Welcome page

When your company orders a mainframe or Linux 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 z15 (z15): machine types 8561 and 8562
- IBM z14 (z14): machine types 3906 and 3907
- IBM z13 (z13) or IBM z13s (z13s)
- IBM LinuxONE III: machine types 8561 and 8562
- 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.

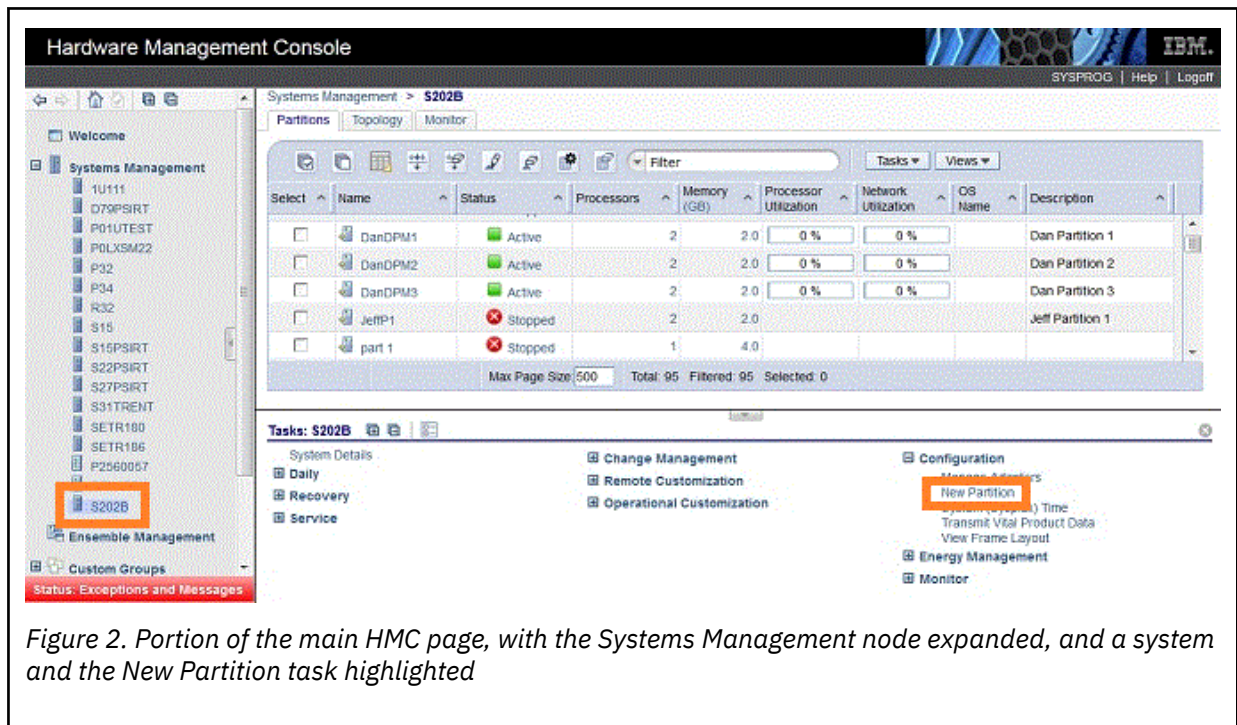


Figure 2. Portion of the main HMC page, with the Systems Management node expanded, and a system and the New Partition task highlighted

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 33](#)
 - [“Creating a new partition” on page 39](#)
- a) Open the **New Partition** task.

This action opens the **New Partition** window, which is shown in [Figure 3 on page 6](#).

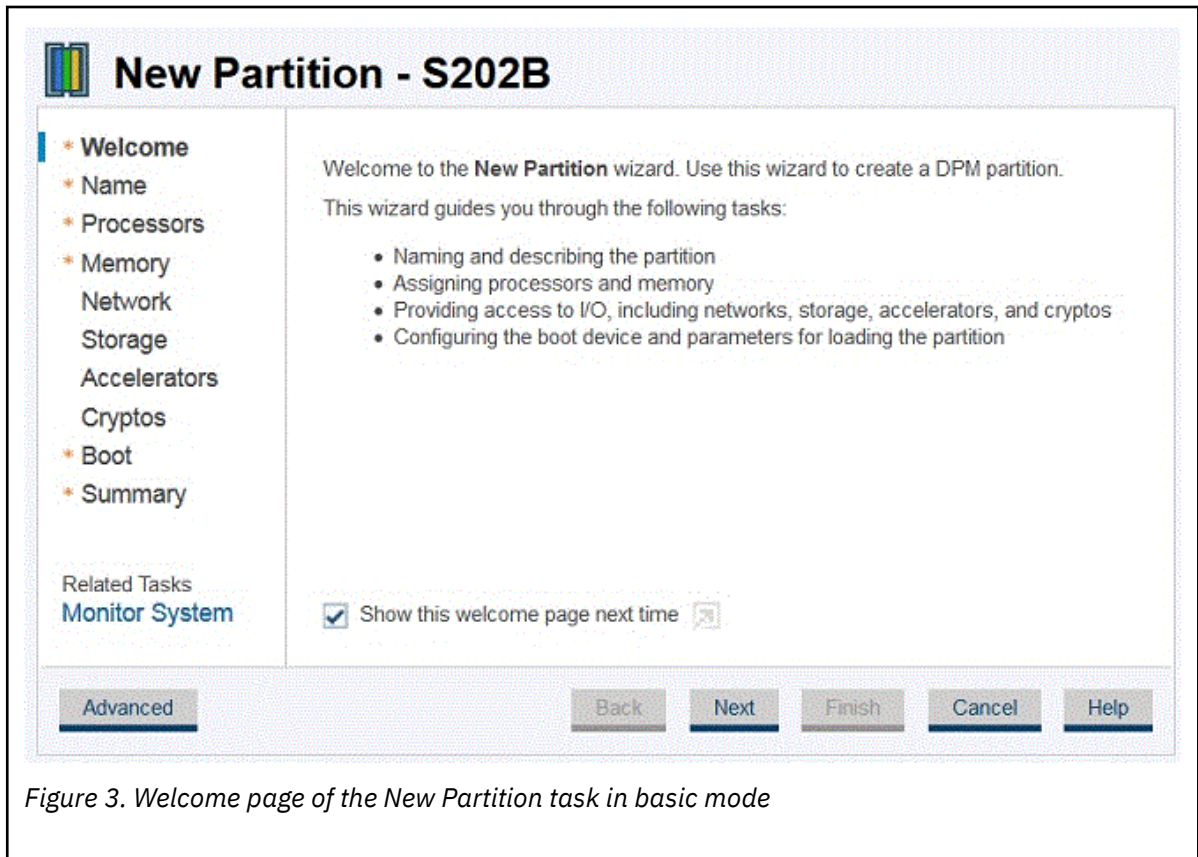


Figure 3. Welcome page of the New Partition task in basic mode

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 37](#).

- b) 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 on page 6](#), 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 39](#).

- c) 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, which displays a message indicating that your partition has been created, and lists more tasks that you can use to work with the new partition.

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 4 on page 7](#).

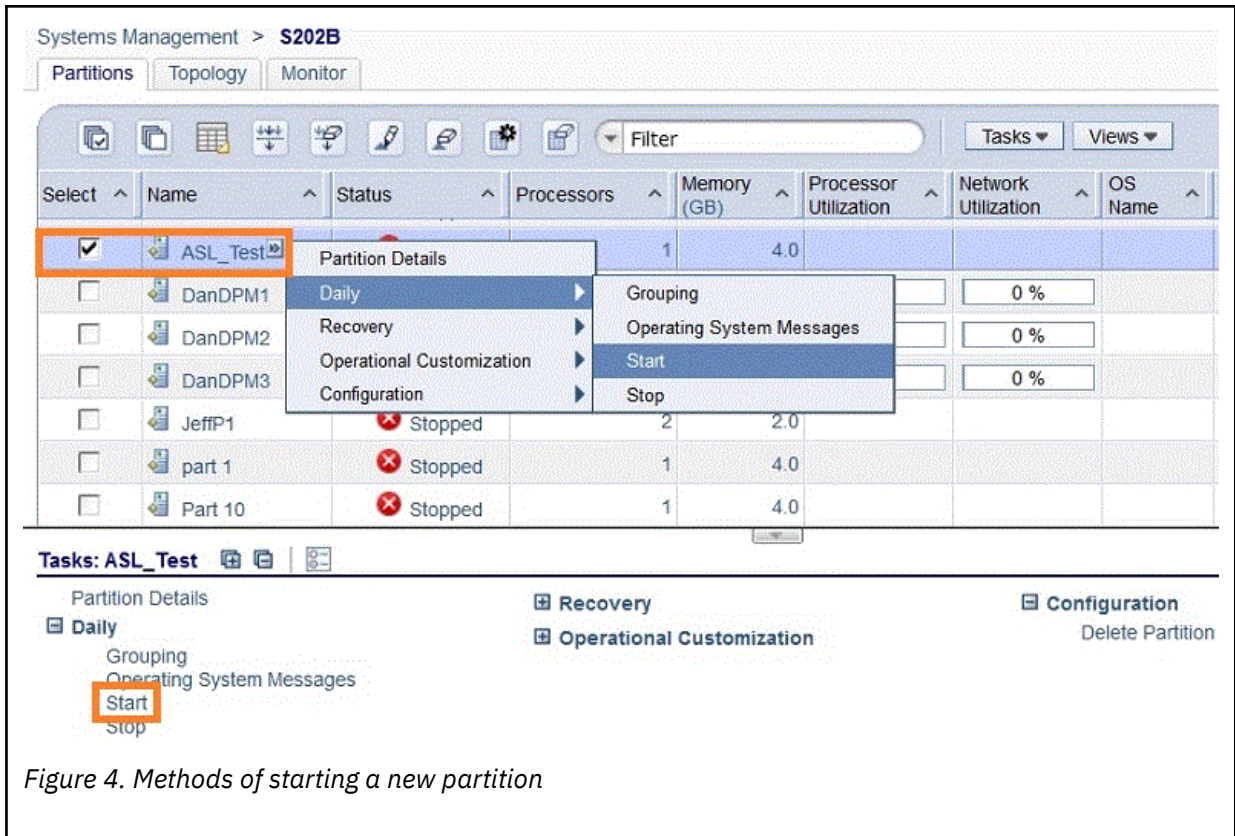


Figure 4. Methods of starting a new partition

- 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 5 on page 7 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.

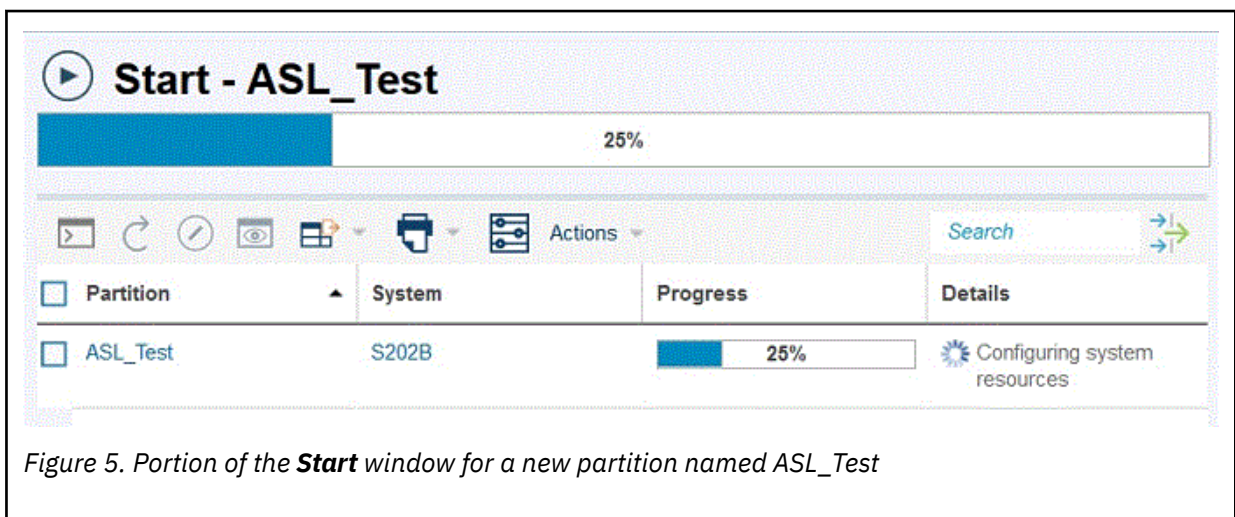


Figure 5. Portion of the **Start** window for a new partition named ASL_Test

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.

5. 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 6 on page 8](#) 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 6. 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 concepts and terms that might be unfamiliar to system administrators who have little or no experience working on a mainframe or Linux system.

Topics covered in this part are:

- [Chapter 3, “Partitions: Virtual images of a mainframe or Linux system,” on page 11](#)
- [Chapter 4, “Adapters: Connections to networks, storage, and more,” on page 17](#)
- [Chapter 5, “Operating systems and hypervisors,” on page 25](#)

Chapter 3. Partitions: Virtual images of a mainframe or Linux system

A partition is a virtual representation of the hardware resources of an IBM Z or LinuxONE system. A partition is the runtime environment for either a hypervisor and its guest operating-system images, each with their own applications; or a single operating system and its applications, which are sometimes called the *workload*.

The system planners at your company order and configure mainframe or Linux 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 in 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 supported as a 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 7 on page 12 illustrates the physical and virtual resources of a mainframe or Linux 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).

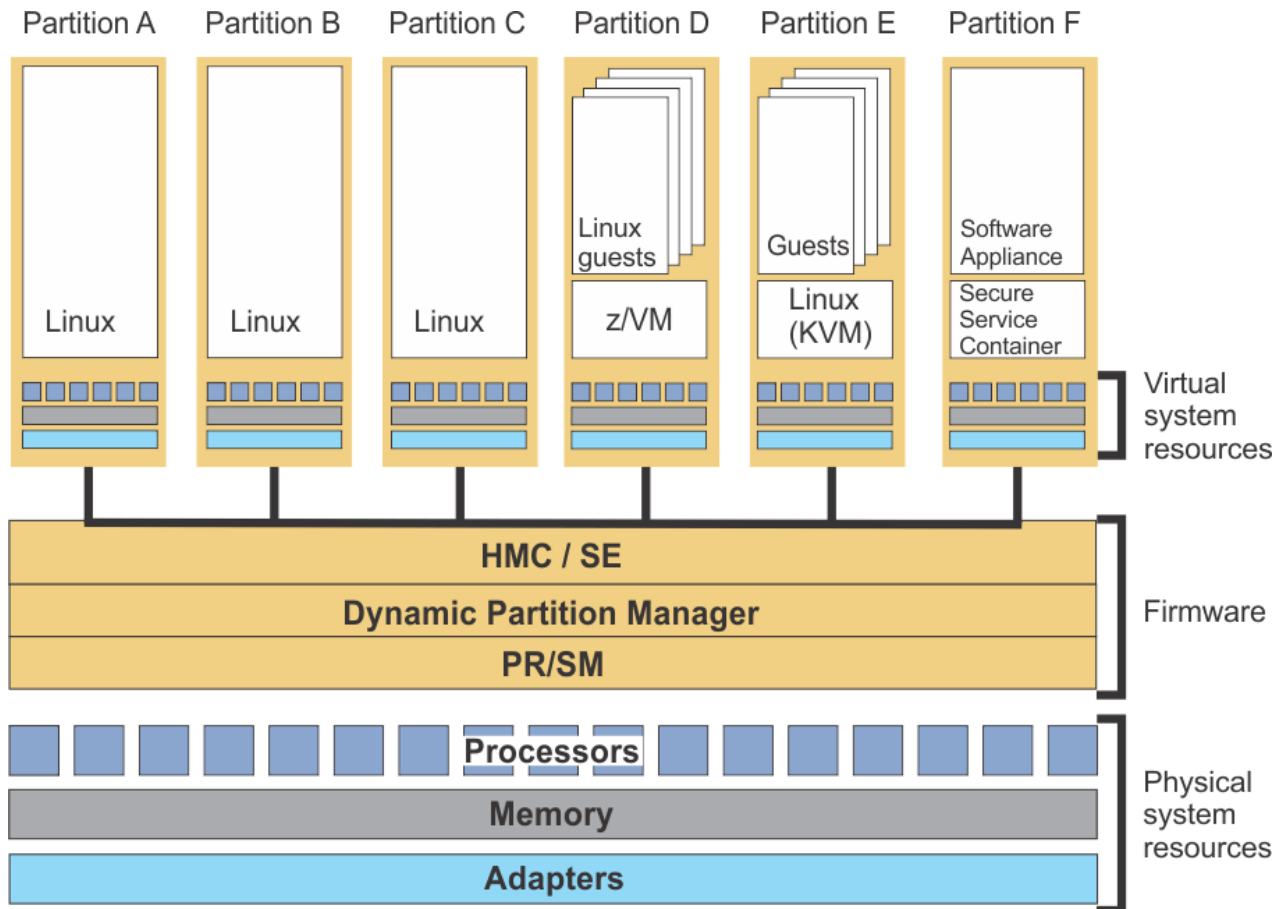


Figure 7. Partitions configured on a DPM-enabled system

In Figure 7 on page 12, 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.
- Partition D hosts one z/VM image and its multiple Linux guests.
- Partition E hosts one Linux hypervisor (for example, Ubuntu KVM) and its multiple guests.
- Partition F is a Secure Service Container partition that hosts a supported software appliance.

Note that DPM does not manage any hypervisor guests, or any appliances that run in a Secure Service Container partition.

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 as a single operating system, or as a hypervisor for multiple guests.

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). In some cases, a system might also support an additional type: Central Processor (CP).

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 in 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 an initial amount of at least 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 HiperSockets, 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

Storage groups or tape links provide a partition with access to internal storage devices, or to external storage area networks (SANs) and devices that are connected to a system. A *storage group* is a logical group of storage volumes that share certain attributes, such as the type or size. A *tape link* defines the attributes of a connection that one or more partitions can use to access one FCP tape library in the SAN.

System administrators create storage groups or tape links for partitions to use. The system administrators work together with storage administrators to correctly configure a storage group or tape link and its associated devices for use.

For partitions to access storage, you attach one or more storage groups or tape links to the partition. Through storage groups and tape links, partitions can access the following types of storage:

- Fibre Connection (FICON) extended count key data (ECKD) direct-access storage devices (DASD), and Fibre Channel Protocol (FCP) Small Computer System Interface (SCSI) disk storage devices, including FCP tape libraries. These devices are physically located in the SAN. FICON and FCP storage groups and FCP tape links can be defined as either dedicated for use by only one partition, or shared by multiple partitions.
- Non-Volatile Memory Express (NVMe) solid state drives, which are installed in a system. Only one partition can use an NVMe storage group at any given time; an NVMe storage group cannot be shared. However, a partition that has attached NVMe storage groups can also have attached FICON and FCP storage groups, and FCP tape links.

Accelerators

An accelerator virtual function provides a partition with access to specific features, such as zEnterprise® Data Compression (zEDC), 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.

Only specific systems support accelerators. To determine whether a specific system supports accelerators, see the appropriate system technical guide on the IBM Redbooks web site at <http://www.redbooks.ibm.com/>

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 in 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.

To check on the status of partitions, select the **Systems Management** node in the HMC navigation pane, and select the **Partitions** tab. The Status column for each partition contains one of the following values.

Active

Indicates that the partition has successfully started and is operating normally.

Communications not active

Indicates a problem with the communication between the Hardware Management Console (HMC) and the Support Element (SE).

Degraded

Indicates that the partition successfully started and is operating, but the availability of physical resources to which it has access is less than required, as stated in the partition definition. This status might be acceptable, for example, for partitions that do not have reserved resources.

Paused

Indicates that, because a user has stopped all processors, the partition is not running its workload. In this case, because the partition was successfully started, its resources are shown as active and are still associated with this partition. A partition can enter this state when, for example, an in-band shutdown is initiated through the operating system; in this case, you must use the **Stop** task to change the partition status to Stopped, which completes the shutdown operation for the partition.

Reservation error

Indicates that the availability of physical resources does not match the reserved resources that are stated in the definition for this partition. The partition cannot start until sufficient resources are available.

Starting

Indicates the transitional phase between Stopped state and Active state, as the result of a Start task issued against this partition.

Status check

Indicates that the current status of the partition is unknown. This condition usually occurs under one of the following circumstances:

- When the SE is starting up; in this case, this partition status is temporary.
- When the SE and the DPM-enabled system to which it is attached cannot communicate.

Stopped

Indicates that the partition has normally ended its operation, and exists only as a partition definition.

Stopping

Indicates the transitional phase between Active state and Stopped state, as the result of a Stop task issued against this partition.

Terminated

Indicates that all of the processors for this partition are in a disabled wait state, or a system check stop occurred. The partition is not running its workload. In this case, because the partition was successfully started, its resources are shown as active and are still associated with this partition.

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. 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 31](#).

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 its 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.

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 Documentation at <https://www.ibm.com/docs/en/linux-on-systems?topic=overview-device-drivers-features-commands>

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 most types of supported adapters, up to specific limits. The following topics provide a description of adapter types, by category.

Network adapters

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

These network adapters are:

- Open Systems Adapter-Express (OSA-Express) adapters, 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. Each HiperSocket adapter represents a single internal network, and only those partitions with access to that adapter participate in that network.
- Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) Express adapters. These adapters provide high speed, low latency data transfer over Ethernet networks.

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 technical guide 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. “[Network adapters](#)” on page 17 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.

DPM-enabled system

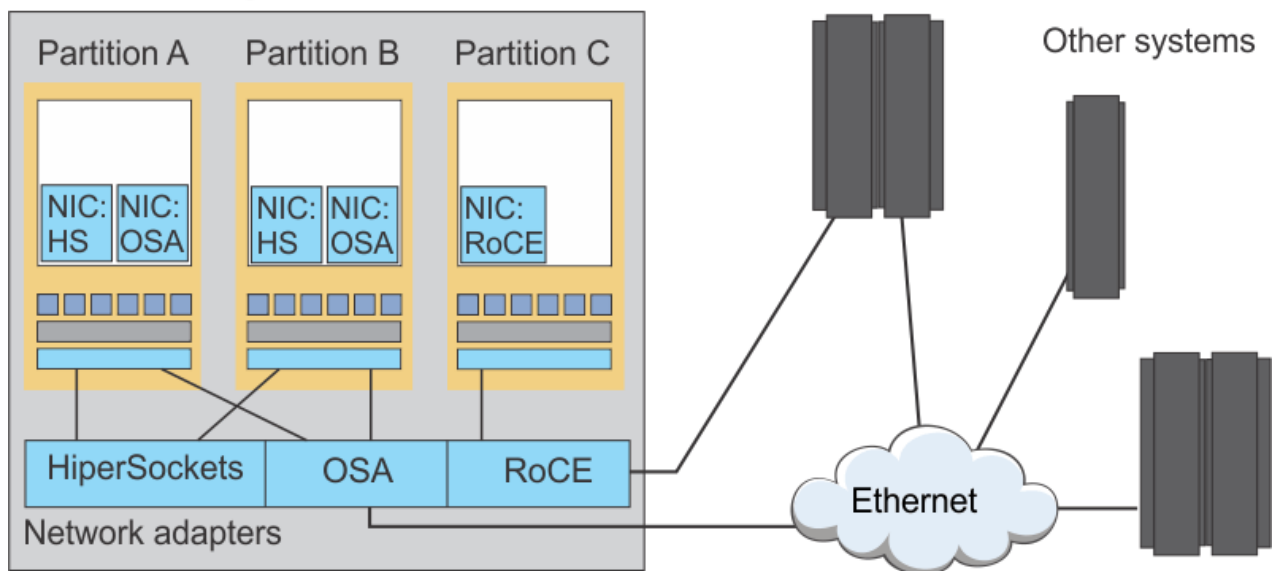


Figure 8. Partitions with NICs configured for access to network adapters on a DPM-enabled system

Storage adapters

Storage adapters provide partitions on a DPM-enabled system with access to external and internal storage devices. The type of available storage adapters varies, depending on the features installed on the system.

- FICON Express adapter cards connect the system to devices in the storage area network (SAN) through Fibre Channel connections. Each FICON adapter card is configured to operate in a specific mode, which determines the type of external storage device that can be connected to the adapter card.
 - Fibre Channel Protocol (FCP) mode provides access to Small Computer System Interface (SCSI) devices, which include tape libraries as well as disk storage.

- FICON native (FC or FICON) mode provides access to extended count key data (ECKD) devices (also known as direct-access storage devices, or DASD).
- Non-Volatile Memory Express (NVMe) storage adapters use the PCI Express (PCIe) protocol to provide high-speed storage within a system. These storage adapters are available on IBM LinuxONE (LinuxONE) systems only, when those systems have IBM Adapter for NVMe1.1 features installed.

DPM automatically detects any storage adapters that are installed in the system. System or storage administrators use the **Configure Storage** task to connect external storage devices to the system, and to define storage groups or tape links that enable partitions to use storage resources. To access storage, you attach one or more storage groups or tape links to a partition, through either the **New Partition** task or **Partition Details** task. Also, you can attach tape links to selected partitions when you use the **Configure Storage** task either to create a tape link or to modify it.

For more information about storage adapter types, see the following topics.

- [“FICON and FCP storage adapters” on page 19](#)
- [“NVMe storage adapters” on page 22](#)

FICON and FCP storage adapters

Fibre Channel connections (FICON) provide high-speed data transfer between systems and storage devices. Fibre Channel networks consist of servers, storage controllers, and other storage 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. Support for FCP mode is available with all DPM versions. In releases prior to R4.3, DPM detects but does not display FCP tape drives, 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 drives. Starting with DPM R4.3, however, administrators can use the **Configure Storage** task to view and manage tape devices, and to configure and provide access to tape storage.
- FICON native (FC or FICON) mode provides access to extended count key data (ECKD) 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 drives. 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 with DPM R3.1 and later DPM versions.

[Figure 9 on page 20](#) shows a sample physical FICON network that can be configured for a system that supports FCP tape and disk storage and FICON DASD.

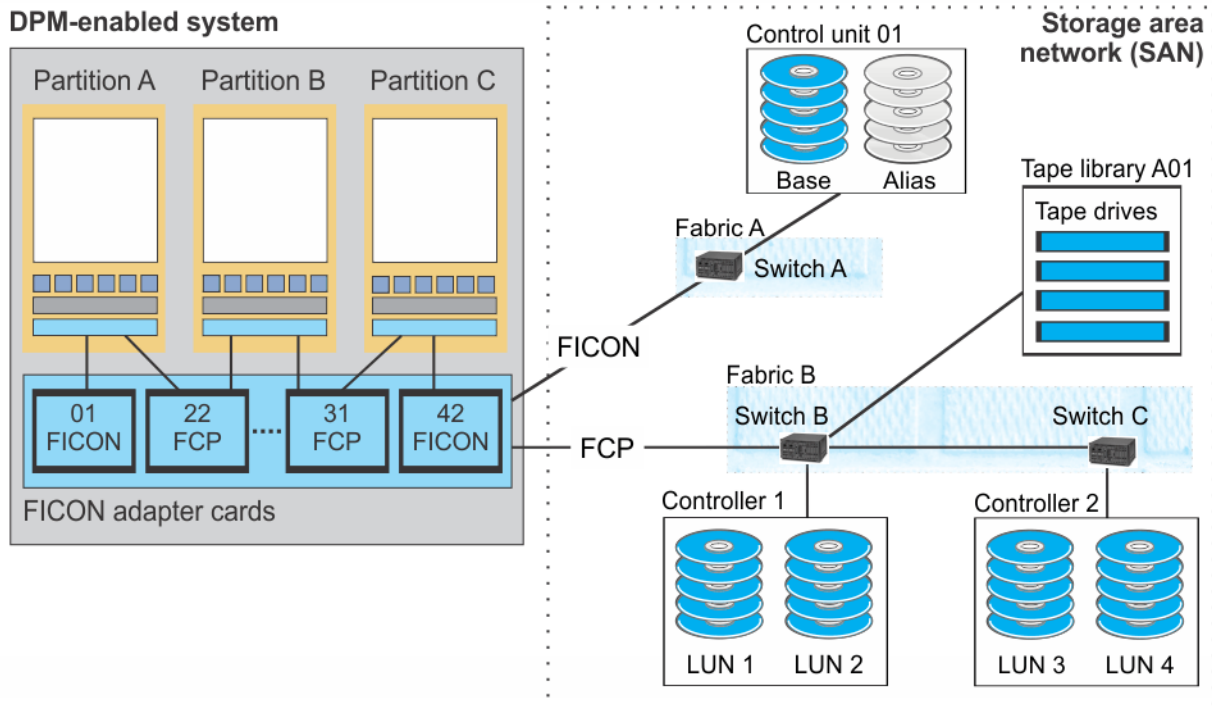


Figure 9. Sample physical FICON and FCP storage configuration

In Figure 9 on page 20:

- 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.
- Adapter cards 22 and 31 also provide access through switch B to tape drives on tape library A01.

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 technical guide on the IBM Redbooks web site at <http://www.redbooks.ibm.com/>.

System or 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 and tape links for partitions to use.

More about FICON and FCP storage groups

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 disk storage, you attach one or more storage groups to a partition, through either the **New Partition** task or **Partition Details** task. 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 in the SAN.

Figure 10 on page 21 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 9 on page 20.

DPM-enabled system

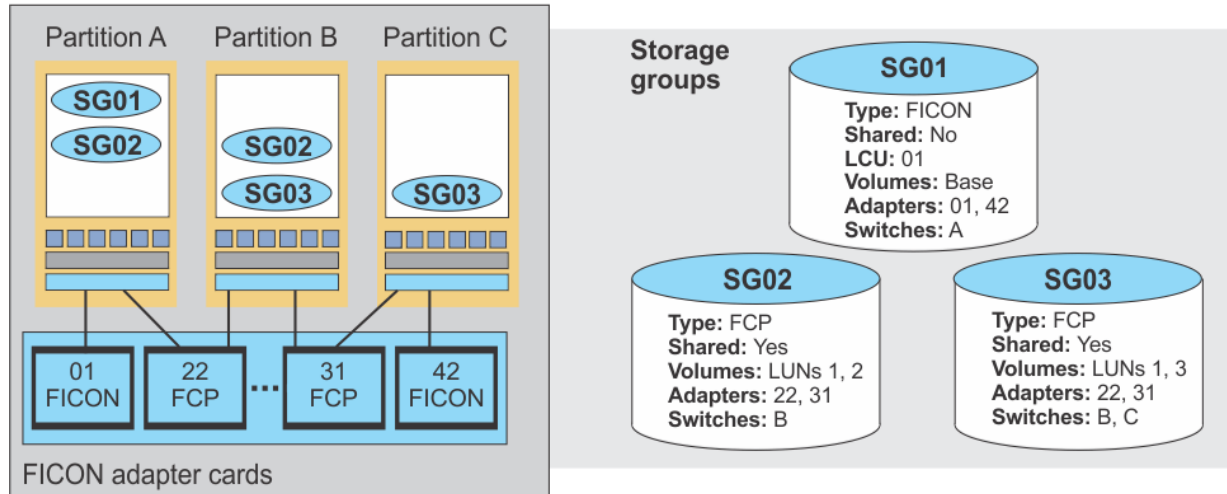


Figure 10. Sample FICON and FCP storage groups that are attached to partitions on the system

In Figure 10 on page 21, 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.

More about FCP tape links

A *tape link* defines the attributes of a connection that one or more partitions can use to access one FCP tape library in the SAN. These attributes include storage resources such as system adapters, world wide port names (WWPNs), and the number of partitions that can share the connection. When an administrator submits a request for a tape link, DPM automatically generates the WWPNs that storage administrators use to fulfill the tape link request. The request goes to one or more storage administrators, who fulfill the request through tools for managing SAN switches.

To access tape storage, you attach one or more tape links to a partition, through the following methods:

- Through the **Configure Storage** task, either when you are creating a new tape link, or modifying an existing tape link. During either process, you can select partitions to which you want to attach the tape link.
- Through the **Storage** section of the **New Partition** task or **Partition Details** task.

When you finish the task or save your changes, DPM generates the HBAs that are required for partitions, and the operating systems that they host, to access the physical tape library in the SAN.

Figure 11 on page 22 shows two tape links with key attributes, and information about the physical devices that an administrator has configured to fulfill them.

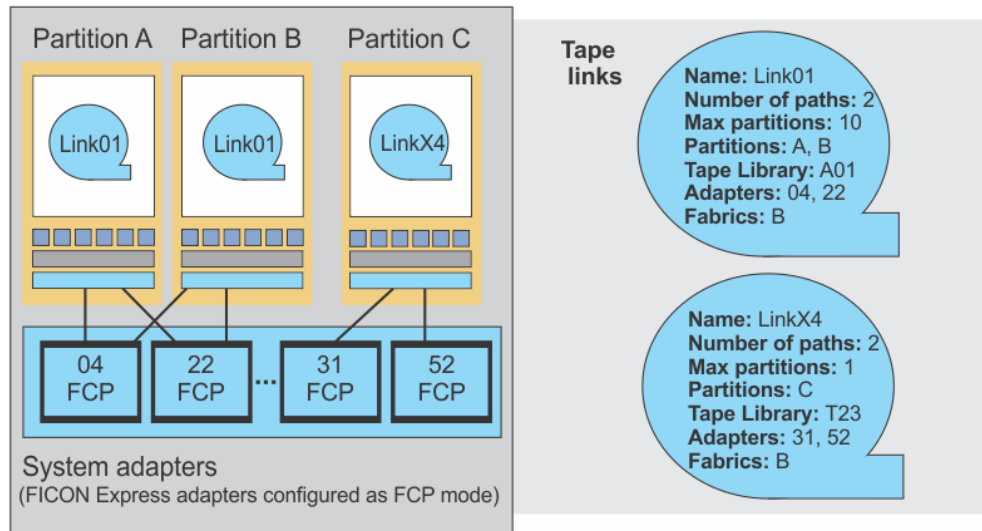
DPM-enabled system

Figure 11. Sample FCP tape links that are attached to partitions on the system

In Figure 11 on page 22, the circles in the partitions indicate which tape links are attached to each partition.

- Partitions A and B share tape link Link01, which provides access to tape drives in tape library A01 through FCP-mode adapter cards 04 and 22 and fabric B.
- Partition C has exclusive access to tape link LinkX4, which provides access to tape drives in tape library T23 through FCP adapters 31 and 52 and fabric B.

NVMe storage adapters

Starting with DPM R4.2, system administrators can enable partitions to use NVMe storage on a system that has IBM Adapter for NVMe1.1 features installed. These features are available on IBM LinuxONE (LinuxONE) systems only.

Non-Volatile Memory Express (NVMe) storage adapters use the PCI Express (PCIe) protocol to provide high-speed storage within a system. Each NVMe adapter consists of two pieces of hardware: an IBM-supplied carrier card installed in a system I/O drawer, and the solid state drive (SSD) that customers purchase. IBM service representatives install the NVMe SSDs in the carrier cards after the system is delivered to the customer site.

DPM automatically discovers any NVMe carrier cards that are installed in the system, and can detect whether an SSD is installed in each carrier card. Through **Connect Storage** and **Storage Cards** in the **Configure Storage** task, system administrators can view NVMe adapters in the visual display of system frames and I/O drawers, and view specific details about each adapter; for example, you can tell whether the carrier card has an SSD installed. They also can use the **Manage Adapters** task to view NVMe adapters, but these adapters are displayed in that task only when an SSD is installed in a carrier card. You cannot use either task to reconfigure NVMe storage adapters; reconfiguration requires properly removing the carrier card and its SSD from the drawer and reinstalling them in a different physical location, as instructed by a service representative.

Through **Request Storage** in the **Configure Storage** task, system administrators can create a storage group containing one or more available NVMe SSDs that are installed in carrier cards. You can define these SSDs as either boot or data volumes.

To enable partitions to use NVMe storage, use either the **New Partition** task or **Partition Details** task to attach the storage group to the partition. You can attach one or more NVMe storage groups, and select an SSD volume in an attached NVMe storage group as the boot volume for an operating system or hypervisor. Note that only one partition can use an NVMe storage group at any given time; an NVMe storage group

cannot be shared. However, a partition that has attached NVMe storage groups can also have attached FICON and FCP storage groups, and FCP tape links.

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.

Only specific systems support accelerators. To determine whether a specific system supports accelerators, see the appropriate system technical guide on the IBM Redbooks web site at <http://www.redbooks.ibm.com/>

Crypto adapters

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. 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 technical guide 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, and the Elliptic Curve Cryptography (ECC) 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, or DES, or ECC 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.

Chapter 5. Operating systems and hypervisors

Partitions on a DPM-enabled system can host a single operating system or hypervisor. 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). Depending on its intended use, the Linux distribution can function as a single operating system image, or as a hypervisor that hosts multiple guests (for example, Ubuntu KVM).

z/VM

In this type of partition, you can install and run z/VM 6.4 or later. z/VM is supported as a 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 12 on page 25 shows a sample configuration of Linux servers and possible workloads.

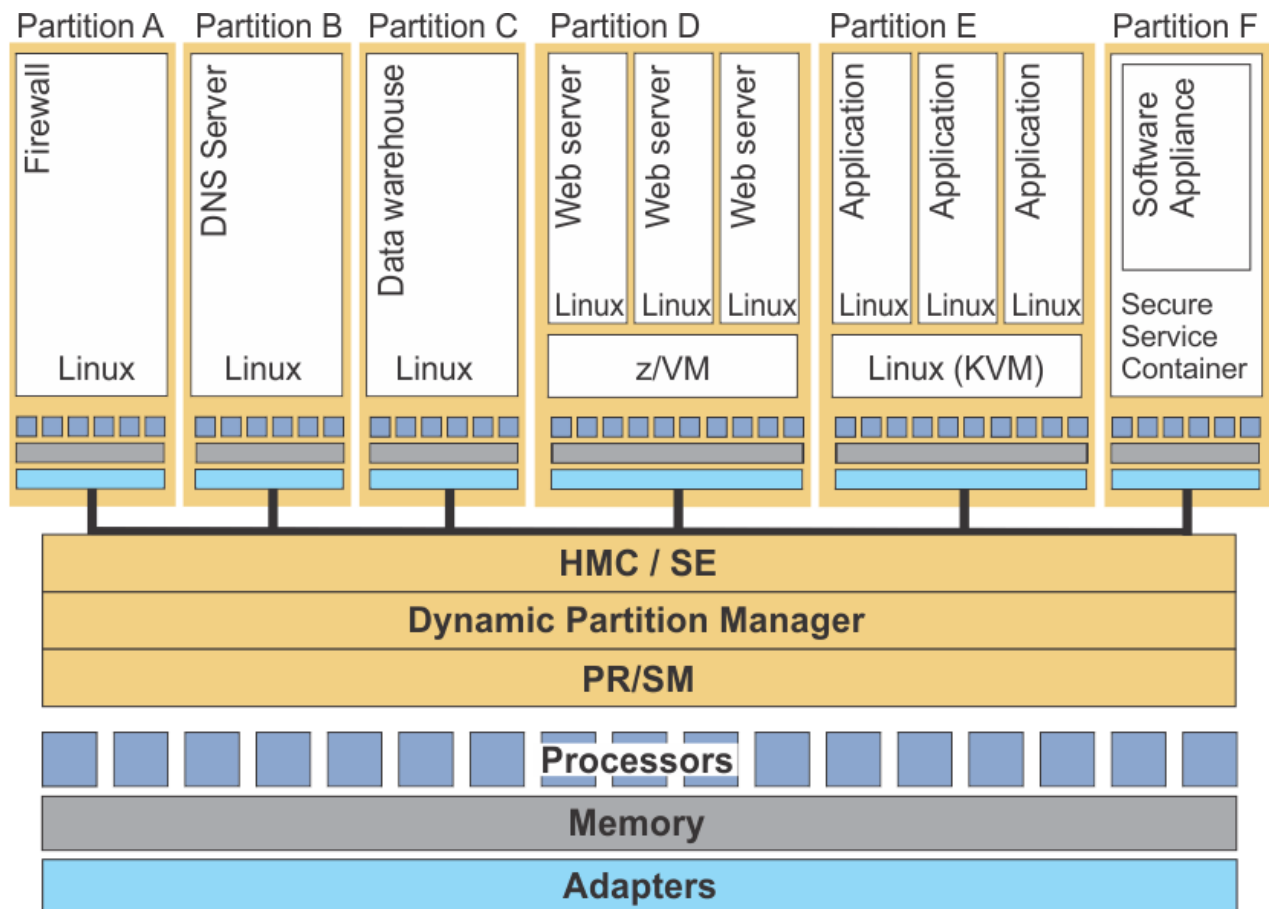


Figure 12. Example: Linux servers configured in partitions on a DPM-enabled system

- 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 Linux hypervisor (for example, Ubuntu KVM) and 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 84](#).

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” on page 26](#)
- [“Boot options for z/VM partitions” on page 27](#)
- [“Boot option for Secure Service Container partitions” on page 29](#)

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. You can also use the **Secure Boot** option to have DPM verify that the software signature matches the signature from the distributor. This option is enabled only when:

- The partition has a partition type of Linux.
- The system that hosts the partition supports the Secure Boot for Linux function.
- You are booting the Linux operating system from a volume in an FCP or NVMe storage group.

For more detail about boot options, see step [“9” on page 45](#) in [“Creating a new partition” on page 39](#), 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 an internal or external storage device. This option is available only when at least one storage group is attached to the partition. When you specify this option, you need to select a boot volume in an attached storage group; you can select a boot volume in any attached FCP or FICON or NVMe storage group. If you plan to boot the operating system or hypervisor from a volume in an NVMe storage group, check the product documentation for the operating system or hypervisor to make sure that it supports booting from NVMe storage.

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. Possible drive selections are **CD/DVD drive** or **USB flash memory drive**, depending on what media drives are installed in the HMC.

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” on page 28 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 in IBM Documentation at <https://www.ibm.com/docs/en/zvm/7.1>

1. Prepare to access the installation files. Note that, depending on the system, a DVD drive might not be available on the HMC. In this case, you can use an alternate HMC media drive instead of a DVD drive, or use an electronic deliverable.
 - 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. Possible drive selections are **CD/DVD drive** or **USB flash memory drive**, depending on what media drives are installed in the HMC.

When you select this option:

- i) If more than one type of media drive is available on the HMC, select the radio button for the media drive on which the INS file resides. Otherwise, skip to the next step.
- ii) Either enter the fully qualified name (relative to the mount point) of an INS file, or complete the following steps.
 - a) Select **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.
 - b) 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:

- i) Select **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.
 - ii) 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 the appropriate edition of *Secure Service Container User's Guide*, 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.

Configuring secure execution for a Linux hypervisor

If the IBM Secure Execution for Linux feature is enabled on the system, you can configure a Linux hypervisor for secure execution, which isolates and protects any guests that run on a hypervisor by restricting host access to guest workloads and data.

- To determine whether the IBM Secure Execution for Linux feature is enabled on the system, go to the **General** section of the **System Details** task and check the values displayed for the Secure Execution field.
- To configure Linux for secure execution, see the product documentation for the Linux distribution that you are using as a hypervisor. This feature supports Linux (KVM) hypervisors.
- To determine whether the Linux hypervisor that runs on a partition is configured for secure execution, go to the **General** section of the **Partition Details** task and check the value displayed for the Secure Execution field.

Where to find more information about operating systems and hypervisors

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 Documentation, at: <https://www.ibm.com/docs/en/linux-on-systems?topic=linuxone-library-overview>
- For the z/VM hypervisor, see the documentation for the version of z/VM that you are installing. The z/VM library is available in IBM Documentation at <https://www.ibm.com/docs/en/zvm/7.1>

Note: DPM supports z/VM with the limitations listed in “Requirements for the hypervisor or operating system” on page 84, 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 the appropriate edition of *Secure Service Container User's Guide*, 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 33](#)
- [Chapter 7, “Creating and starting a new partition,” on page 37](#)
- [Chapter 8, “Dynamically modifying the resources of an active partition,” on page 55](#)
- [Chapter 9, “Summary of tasks for managing systems, adapters, and partitions,” on page 71](#)

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 its 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 on page 33 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 39.


| 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: |
|---|---|--|
| | <p>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:</p> <ul style="list-style-type: none"> • Partition access • Counter facility authorization • Sampling facility authorization | <p>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.</p> |
| | <p>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.</p> <p>If you want to enable simultaneous multithreading for this partition, you must select the IFL processor type.</p> | <p>Amount: _____</p> <p>Type: IFL CP</p> <p>Shared: Yes No</p> |
| | <p>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.</p> <p>If you are creating a Secure Service Container partition, you must specify an initial amount of at least 4096 MB (4 GB).</p> | <p>Amount: _____ MB GB TB</p> <p>Maximum amount: _____ MB GB TB</p> |

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

| ✓ | Planning item or decision | Your notes: |
|---|---|---|
| | <p>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.</p> <p>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.</p> <p>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.</p> | <p>NIC name: NIC device ID: VLAN ID: MAC address: Adapter type: HiperSocket OSA RoCE</p> <p>NIC name: NIC device ID: VLAN ID: MAC address: Adapter type: HiperSocket OSA RoCE</p> <p>NIC name: NIC device ID: VLAN ID: MAC address: Adapter type: HiperSocket OSA RoCE</p> |
| | <p>Storage: Attach one or more storage groups or tape links to enable a partition to access storage resources that are connected to the system. Use the attributes and other information to determine which storage groups or tape links to select.</p> <p>A <i>storage group</i> is a logical group of storage volumes that share certain attributes. FICON and FCP storage groups can be shared by multiple partitions, and multiple storage groups can be attached to one partition. In contrast, only one partition can use an NVMe storage group at any given time; an NVMe storage group cannot be shared. However, a partition that has attached an NVMe storage group can also attach FICON and FCP storage groups or FCP tape links.</p> <p>A <i>tape link</i> defines the attributes of a connection that one or more partitions can use to access one FCP tape library in the SAN. These connection attributes include storage resources such as system adapters, world wide port names (WWPNs), and the number of partitions that can share the connection. Note that you can also attach a tape link to a selected partition through the Configure Storage task, when you use it to either create a new tape link or modify an existing one.</p> | <p>Storage group name: Type: FCP FICON NVMe</p> <p>Storage group name: Type: FCP FICON NVMe</p> <p>Storage group name: Type: FCP FICON NVMe</p> <p>Tape link name:</p> <p>Tape link name:</p> <p>Tape link name:</p> |

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

| ✓ | Planning item or decision | Your notes: |
|---|---|---|
| | <p>Operating system / Hypervisor: For a list of supported operating systems and hypervisors, see Chapter 5, “Operating systems and hypervisors,” on page 25. 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.</p> | <p>Operating system Hypervisor:</p> <p>Location: SAN Net boot FTP HMC local drive</p> <p>ISO file name:</p> <p>.INS file name:</p> |
| | <p>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. If you are booting the Linux operating system from a volume in an FCP or NVMe storage group, you can select Secure Boot to have DPM verify that the software signature matches the signature from the distributor.</p> <p>DPM automatically sets the boot option for the first-time start of a Secure Service Container partition.</p> <p>For a list and descriptions of the boot options, see the online help for the New Partition task.</p> | <p>Boot loader time-out value: _____ seconds</p> <p>Storage device (SAN): Storage group boot volume:</p> <p>FTP server: Host name: User name / password: Protocol: FTP FTPS SFTP</p> <p>HMC removable media: Type of media drive: CD/DVD USB (Note that, depending on the system, a DVD drive might not be available on the HMC.)</p> |

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” on page 37](#)
- [“Creating a new partition” on page 39](#)
- [“Starting a partition and its operating system or hypervisor” on page 49](#)

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 on page 37 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 | — | ✓ |

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

| Partition property | Basic mode | Advanced mode |
|---|--|---|
| Acceptable partition status values | — | ✓ |
| Controls: <ul style="list-style-type: none"> • 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, tape links, or HBAs) | ✓ | ✓ (ability to edit device numbers for FCP storage groups and tape links, and to change adapters for FCP storage groups in this mode only) |
| Accelerators (virtual functions) | ✓ (if supported by the system and installed) | ✓ (if supported by the system and installed) |
| Cryptos (security) | ✓ (if installed on system) | ✓ (if installed on system) Permitting AES, DES, or ECC protected key import is available in this mode only. |
| Boot options, including Secure Boot for Linux | ✓ | ✓ |

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. In other words, switching from advanced mode to basic mode wipes out all changes that you made in advanced mode, and restores the changes that you made in basic mode.

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 37.

- 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 33 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 a page overlay that highlights key task controls on the window.

- c) Click the **Okay, got it** button to remove the page overlay.

The Welcome page is displayed.

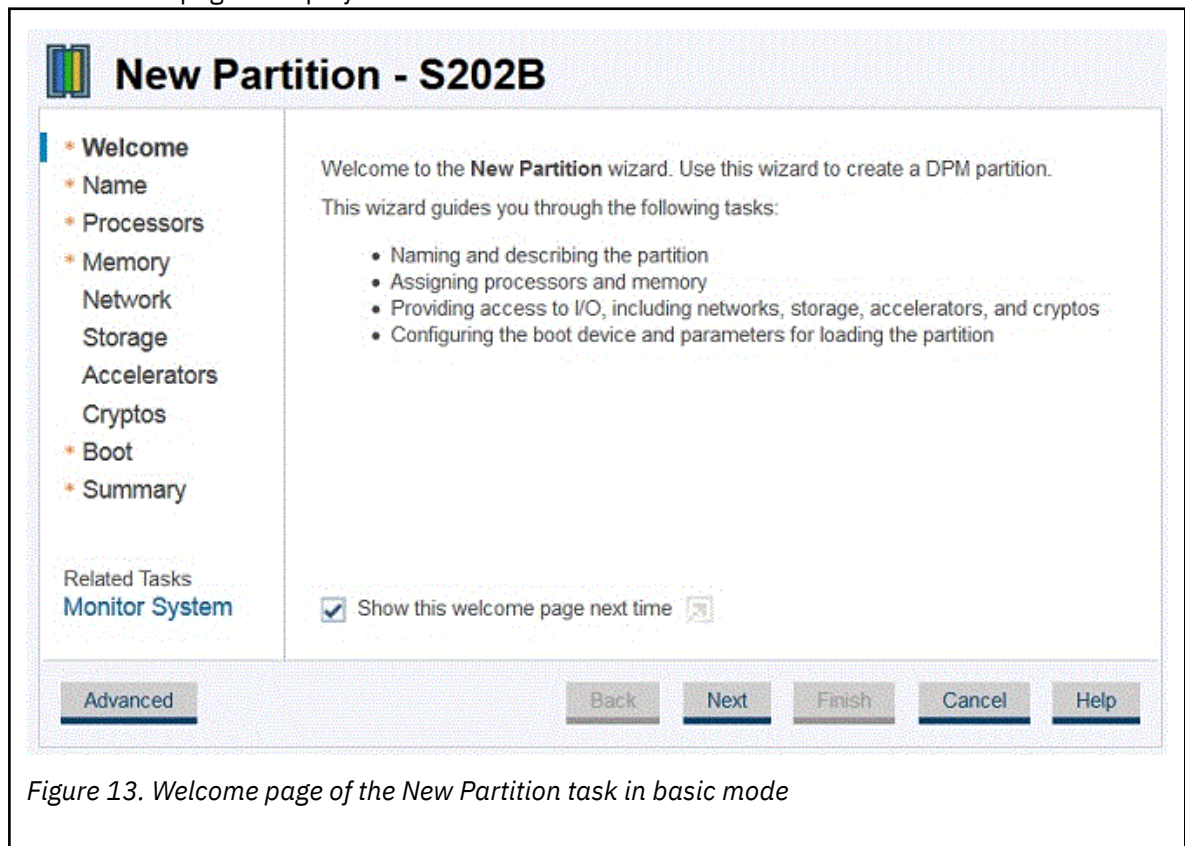


Figure 13. Welcome page of the New Partition task in basic mode

- 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.

- 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.
- Optionally, specify a description for the partition. The description can be up to 1024 characters in length.
- 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 as a single operating system, or as a hypervisor for multiple guests.

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.

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

- 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.

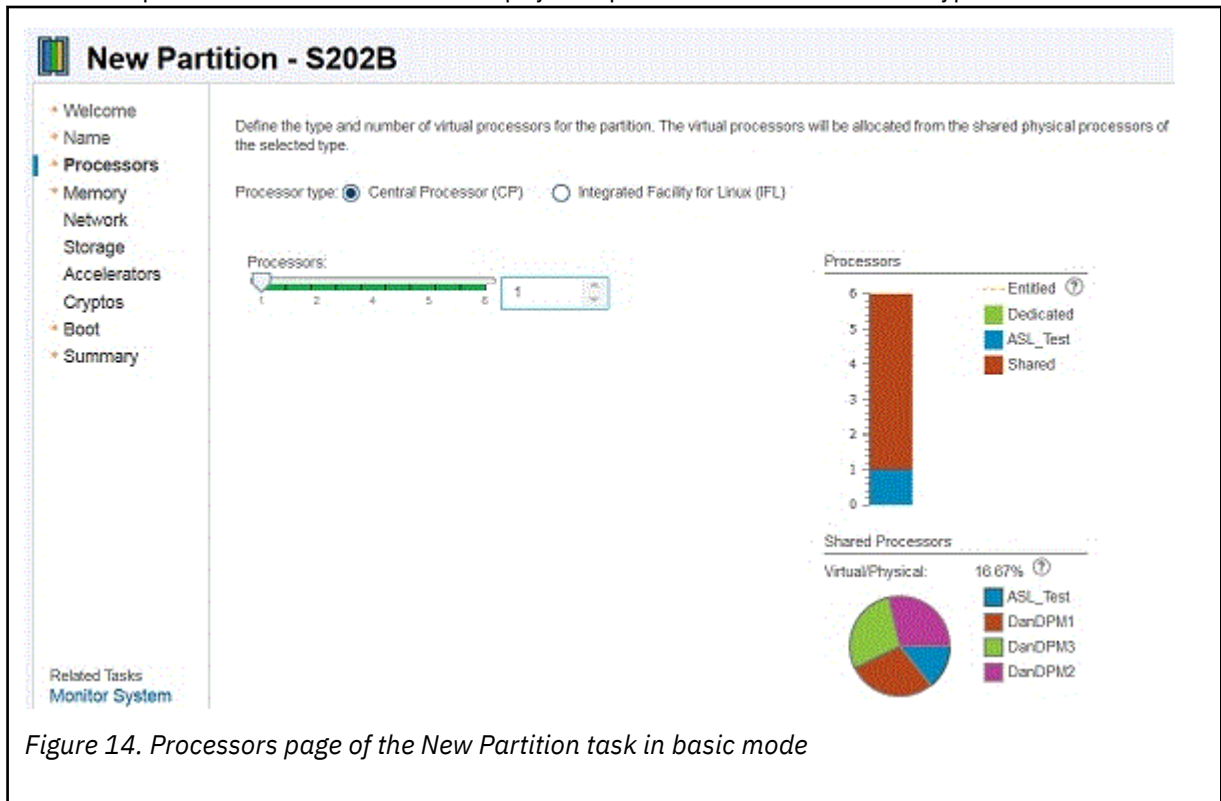


Figure 14. Processors page of the New Partition task in basic mode

- 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.
- 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.
- 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 in 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 in this new partition.

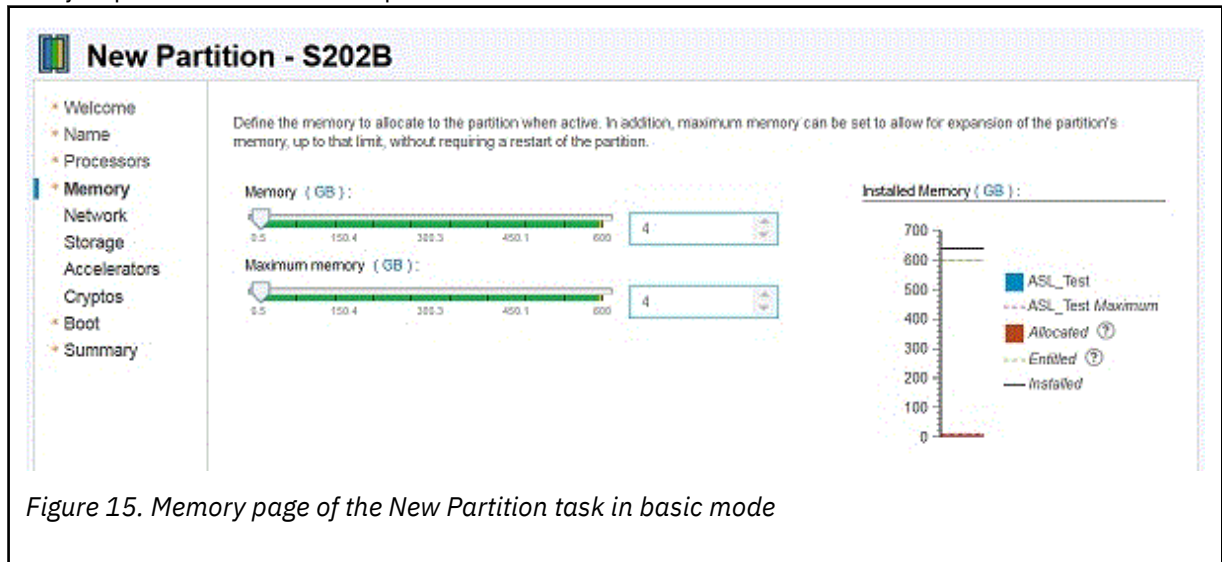


Figure 15. Memory page of the New Partition task in basic mode

- 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 an initial amount of at least 4096 MB (4 GB).
- c) To understand how your selection affects the availability of memory resources on the system, review the updated Installed Memory bar chart.
- 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 44 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.

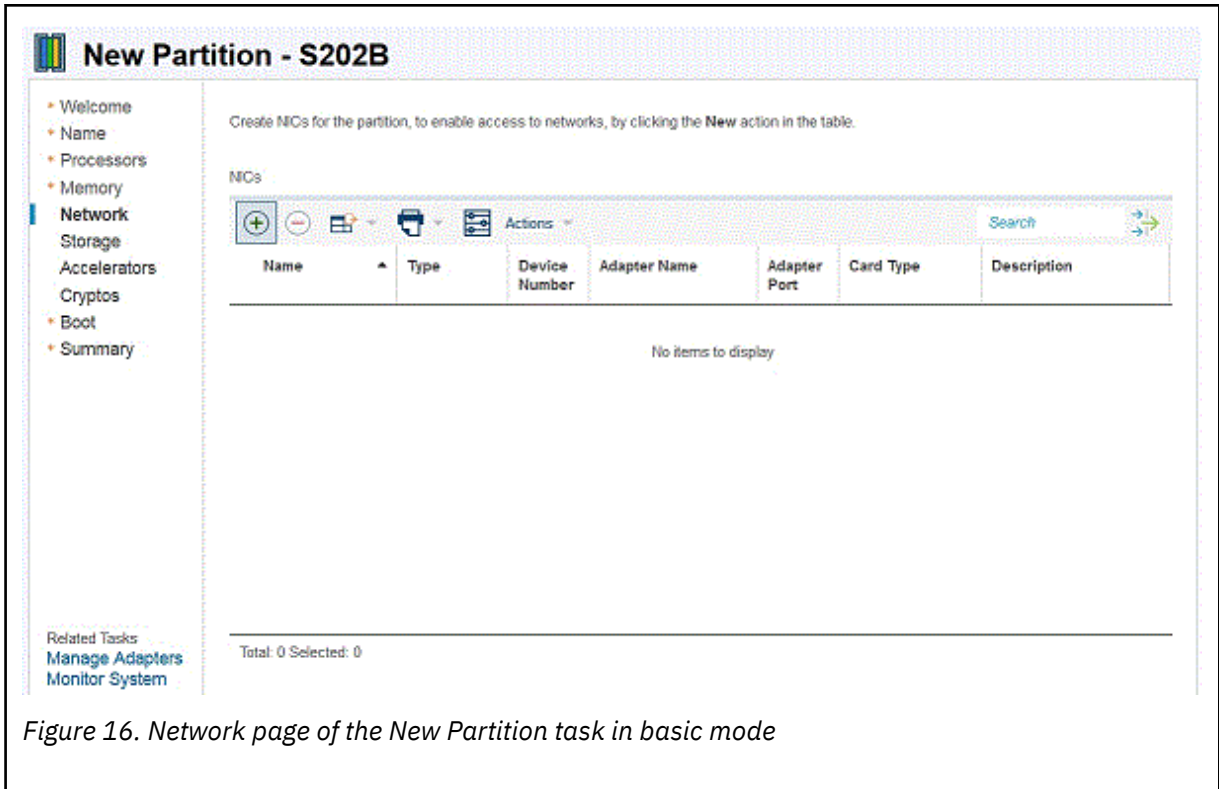


Figure 16. Network page of the New Partition task in basic mode

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

- From the Actions list in the NICs table, select **New** to open the **New Network Interface Card** window.
- Enter a unique, meaningful name and, optionally, a description of the new NIC.
- 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.

- Review the entries in the Adapter Ports and Switches table to determine which network adapters are configured on the system.
 - 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.
 - 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.
 - 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 49.](#)
- 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 tape links that enable the partition to access storage networks and hardware that is connected to the DPM-enabled system.

The Storage section contains a Storage Groups table and a Tape Links table, each with controls that you can use to attach storage groups or tape links to the partition.

If you are creating a partition only to familiarize yourself with the process, you do not need to attach any storage groups 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 45](#) of these instructions. Otherwise, continue with this step.

System administrators create storage groups and tape links 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. A *tape link* defines the attributes of a connection that one or more partitions can use to access one FCP tape library in the SAN.

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 and Tape Links table. Select the plus icon in the table toolbar to open the **Attach Storage Groups** or **Attach Tape Links** 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.
 - 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.

When you have finished selecting storage groups to attach, select **OK** to close the **Attach Storage Groups** window.

- c. On the **Attach Tape Links** window, select one or more tape links listed in the table to attach to this partition.
 - The suggested practice is to select tape links that are in the Complete fulfillment state, but you can select any tape link 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 links in states other than Complete, some storage might not be available for use when you start the partition.
 - Use the additional information in the table, as necessary, to decide which tape links to attach.

When you have finished selecting tape links to attach, select **OK** to close the **Attach Tape Links** window.

- d. Check the entries for the storage groups or tape links that you selected, which are now displayed in the Storage Groups table or Tape Links table in the Storage section. If necessary, you can use the minus icon in the table toolbar to remove a storage group or tape link from the table.

e. When you have finished, review another section or click **OK** to save the partition definition.

The next page to open might be either Accelerators, Cryptos, or Boot, depending on the system configuration. The Accelerators page is displayed only when a system that supports accelerators is managed through this HMC, and is enabled only for systems that support accelerators.

7. If the system supports accelerators and has one installed, use the Accelerators page to enable the new partition to use accelerators that it requires.

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” on page 45. 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 45 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 45. 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 in 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.

Use the **Secure Boot** option to have DPM verify that the software signature matches the signature from the distributor. If the signature does not match, the boot process ends. This option is enabled only when:

- The partition has a partition type of Linux.
- The system that hosts the partition supports the Secure Boot for Linux function.
- You are booting the Linux operating system from a volume in an FCP or NVMe storage group.

If you have selected the partition type **Secure Service Container**, note that option set in the "Boot from" menu is **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.

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

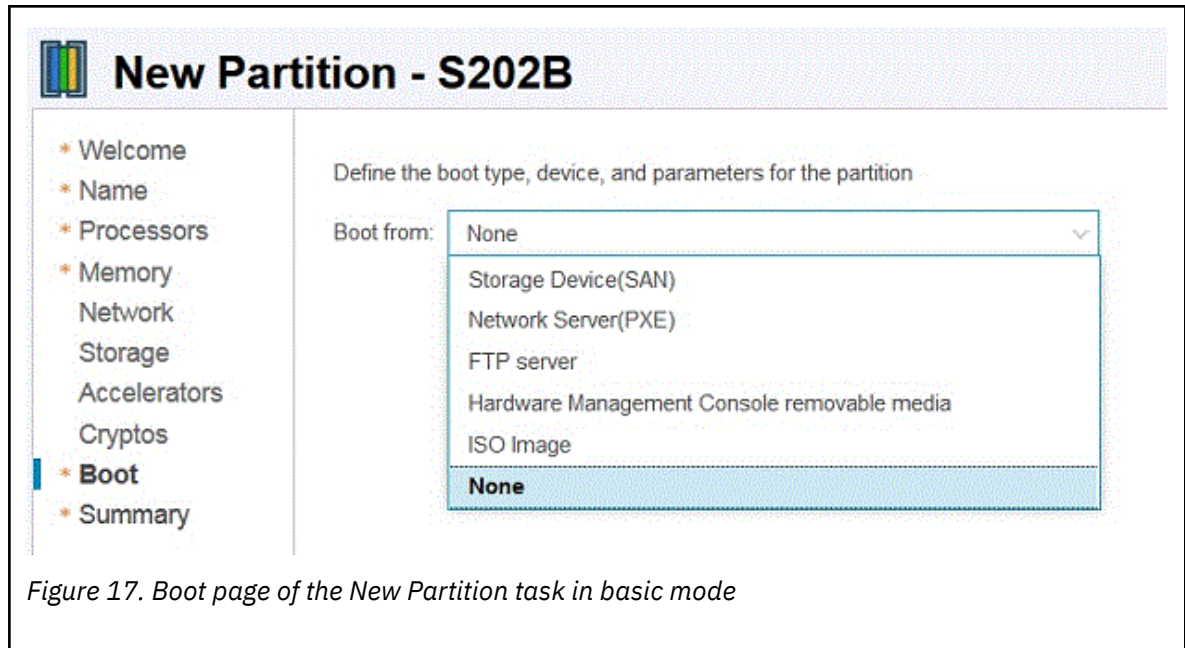


Figure 17. Boot page of the New Partition task in basic mode

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 an internal or external storage device. This option is available only when at least one storage group is attached to the partition. When you specify this option, you need to select a boot volume in an attached storage group.

The Storage Groups table displays the available FCP or FICON or NVMe 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.
- For each boot volume in an NVMe storage group, the Boot Volume table provides the boot volume serial number and capacity, along with a user-supplied description, if any. When you select an NVMe volume, note that NVMe namespace management is not supported, so you can boot programs only from namespace ID=1.

Select the boot volume on which the operating system or hypervisor image resides. If you plan to boot the operating system or hypervisor from a volume in an NVMe storage group, check the product documentation for the operating system or hypervisor to make sure that it supports booting from NVMe storage.

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. Possible drive selections are **CD/DVD drive** or **USB flash memory drive**, depending on what media drives are installed in the HMC.

When you select this option:

- i) If more than one type of media drive is available on the HMC, select the radio button for the media drive on which the INS file resides. Otherwise, skip to the next step.
- ii) Either enter the fully qualified name (relative to the mount point) of an INS file, or complete the following steps.
 - a) Select **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.
 - b) 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:

- i) Select **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.
 - ii) 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.

What to do next

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

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 appropriate edition of *Secure Service Container User's Guide* 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 161](#).

Procedure

1. 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.
2. On the **Partitions** tab for the selected system, find the table entry for the partition that you want to start, and select it.
3. 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.
 - a) 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 in this partition, or to upload the required files to initialize the hypervisor or operating system when the partition itself is started.
 - b) Select **OK** to apply any changes and close the Partition Details window.
4. Use the **Start** task to start the selected 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.
 - b) Expand the **Daily** group, and select **Start**.

If 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 messages that indicate the outcome.

- **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 FICON or FCP storage groups are visible to the Linux operating system

After you start the new partition, you might need to enter Linux commands to make FICON or FCP storage groups available to the operating system that the partition hosts. NVMe storage groups are automatically detected by the operating system, so you do not need to enter Linux commands to make that type of storage group available to the operating system. Similarly, the tape devices that are available through attached tape links are automatically detected by the operating system, so you do not need to enter Linux commands for tape devices either.

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 18. 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 |
|---------------------------------------|-----------|---------|-----------|---------|
| 0x0000000000000000-0x00000001ffffffff | 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 19. 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 20. 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 21. Sample displays resulting from the **lschp** command

To display information about adapters that use the PCI Express (PCIe) protocol, such as Non-Volatile Memory Express (NVMe) adapters, use the **lspci** command, as shown in Figure 22 on page 53.

```
localhost:/home/dpm # lspci
```

```
0101:00:00.0 Non-Volatile memory controller: Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb (rev 01)
0201:00:00.0 Non-Volatile memory controller: Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb (rev 01)
0202:00:00.0 Non-Volatile memory controller: Intel Corporation NVMe Datacenter S SD [3DNAND, Beta Rock Controller]
```

Figure 22. Sample displays resulting from the **lspci** command

To display information about FCP tape storage, use the **lscss** command to display the HBA device numbers that the partition is using for tape links, and the **lstape** command to list the available tape drives. Figure 23 on page 53 shows sample results from the **lstape** command.

```
((root@localhost))> lstape --scsi-only --verbose
```

| Generic | Device | Target | Vendor | Model | Type | State |
|---------|----------|--------------------|--------|------------------|---------|---------|
| | HBA | WWPN | | Serial | | |
| sg0 | st0 | 0:0:0:0 | IBM | ULT3580-HH6 | tapedrv | running |
| | 0.0.0003 | 0x2002000e11159c32 | | 90WT800375 | | |
| sg1 | sch0 | 0:0:0:1 | IBM | 3573-TL | changer | running |
| | 0.0.0003 | 0x2002000e11159c32 | | 00L2U78Z8185_LL0 | | |
| sg2 | st1 | 0:0:1:0 | IBM | ULT3580-HH6 | tapedrv | running |
| | 0.0.0003 | 0x2005000e11159c32 | | 1068012586 | | |
| sg3 | st2 | 1:0:0:0 | IBM | ULT3580-HH6 | tapedrv | running |
| | 0.0.0004 | 0x2002000e11159c32 | | 90WT800375 | | |
| sg4 | sch1 | 1:0:0:1 | IBM | 3573-TL | changer | running |
| | 0.0.0004 | 0x2002000e11159c32 | | 00L2U78Z8185_LL0 | | |
| sg5 | st3 | 1:0:1:0 | IBM | ULT3580-HH6 | tapedrv | running |
| | 0.0.0004 | 0x2005000e11159c32 | | 1068012586 | | |

Figure 23. Sample displays resulting from the **lstape** 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
```

Configuring secure execution for a Linux hypervisor

If the IBM Secure Execution for Linux feature is enabled on this system, you can configure a Linux operating system that functions as a hypervisor for secure execution, which isolates and protects any guests that run on a hypervisor by restricting host access to guest workloads and data.

- To determine whether the IBM Secure Execution for Linux feature is enabled on the system, go to the **General** section of the **System Details** task and check the values displayed for the Secure Execution field.
- To configure Linux for secure execution, see the product documentation for the Linux distribution that you are using as a hypervisor.
- To determine whether the Linux hypervisor that runs on a partition is configured for secure execution, go to the **General** section of the **Partition Details** task and check the value displayed for the Secure Execution field.

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 Documentation at <https://www.ibm.com/docs/en/linux-on-systems?topic=overview-device-drivers-features-commands>
- 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 in IBM Documentation at <https://www.ibm.com/docs/en/zvm/7.1>

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” on page 55](#)
- [“Example: Adding more memory resources” on page 57](#)
- [“Example: Adding a new network device” on page 59](#)

Use the following instructions to dynamically add or manage storage resources for a partition.

- [“Instructions: Attaching storage groups or tape links” on page 62](#)
- [“Instructions: Detaching and reattaching an FCP storage group” on page 64](#)
- [“Instructions: Changing the adapters that are assigned to an FCP storage group” on page 65](#)
- [“Instructions: Changing the adapters that are assigned to an FCP tape link” on page 66](#)

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.

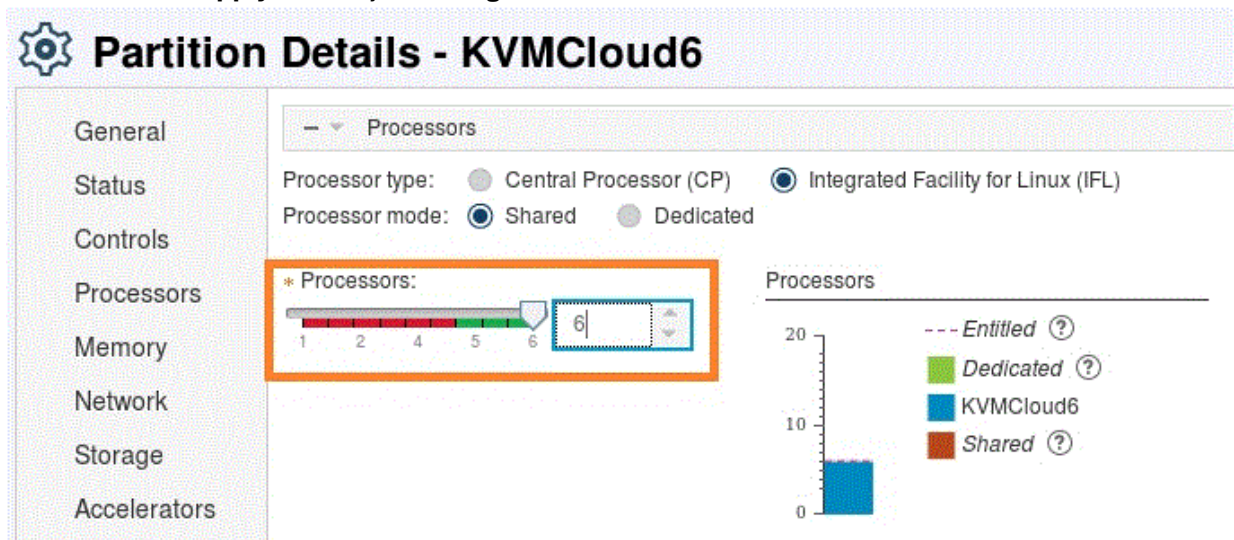


Figure 25. Sample screen of the Processors section in the Partition Details window

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@lpskvm4 ~]# 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@lpskvm4 ~]#

```

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.

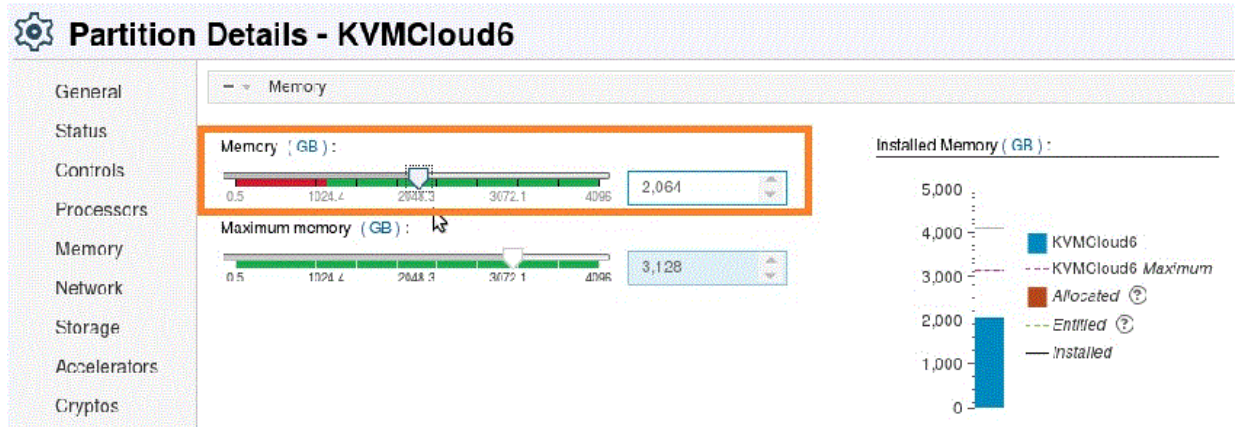


Figure 29. Sample screen of the Memory section in the Partition Details window

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

```
[root@lpskvm4 /]# time chmem -e 8g
```

```

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

```

```
[root@lpskvm4 /]# 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@lpskvm4 /]#

```

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).

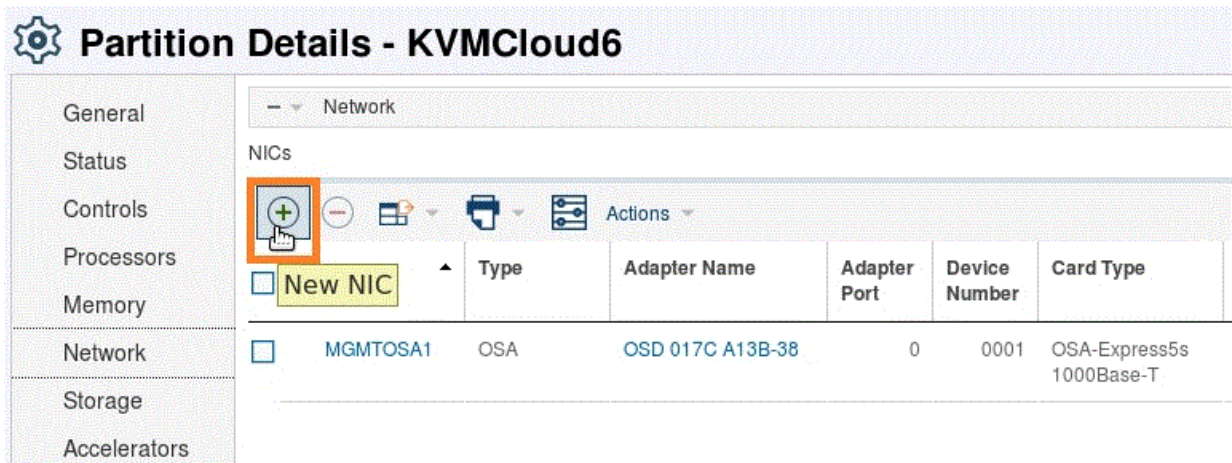


Figure 32. Sample screen of the Network section with the New NIC icon highlighted

2. Select the adapter for the new network connection.


```

[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

Notes: 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.
- The device might be included in a list of devices to be ignored by Linux. If **cio_ignore** is active, the list is automatically adjusted to accommodate all devices that are configured on the HMC interface **only** when you reboot the operating system. To dynamically add the device, use **cio_ignore -r** to remove the device from the list, then set the device online.

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 62](#) 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 or tape links

Through the Storage section of the **Partition Details** task, you can attach one or more FICON, FCP, or NVMe storage groups or FCP tape links to an existing partition. FICON and FCP storage groups can be shared by multiple partitions, and multiple storage groups can be attached to one partition. FCP tape links also can be either shared by many partitions or dedicated to one. In contrast, only one partition can use an NVMe storage group at any given time; an NVMe storage group cannot be shared.

Note: In addition to using the **Partition Details** task to attach an FCP tape link, you can also attach a tape link to a selected partition through the **Configure Storage** task, when you use it to either create a new tape link or modify an existing one.

In the **Partition Details** task, the Storage section contains a Storage Groups table and a Tape Links table. To attach new storage groups or tape links to the partition, complete the following steps.

1. Select the plus icon in the table toolbar to open the **Attach Storage Groups** or **Attach Tape Links** 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.
 - 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.

When you have finished selecting storage groups to attach, select **OK** to close the **Attach Storage Groups** window.

3. On the **Attach Tape Links** window, select one or more tape links listed in the table to attach to this partition.
 - The suggested practice is to select tape links that are in the Complete fulfillment state, but you can select any tape link 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 links in states other than Complete, some storage might not be available for use when you start the partition.
 - Use the additional information in the table, as necessary, to decide which tape links to attach.

When you have finished selecting tape links to attach, select **OK** to close the **Attach Tape Links** window.

4. Check the entries for the storage groups or tape links that you selected, which are now displayed in the Storage Groups table or Tape Links table in the Storage section. If necessary, you can use the minus icon in the table toolbar to remove a storage group or tape link 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. NVMe storage groups are automatically detected by the operating system, so you do not need to enter Linux commands to make that type of storage group available to the operating system. Similarly, the tape devices that are available through attached tape links are automatically detected by the operating system, so you do not need to enter Linux commands for tape devices either. The actions required for FCP or FICON storage groups depend on the type and fulfillment state, and whether the storage group contained the boot volume for the operating system.

Note: If you issue the **lscss** command and the device number of the newly attached storage group is not listed, the device might be included in a list of devices to be ignored by Linux. If **cio_ignore** is active, the list is automatically adjusted to accommodate all devices that are configured on the HMC interface **only** when you reboot the operating system. To dynamically add the device, use **cio_ignore -i** to remove the device from the list, then set the device online.

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.
 - i) Expand the storage group table entry to show the Host Bus Adapters (HBA) table, and record the device number and the assigned adapter.
 - ii) 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.
 - i) 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.
 - ii) 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).
 - iii) Select **Apply** to save your changes.
 - e. Go to the **Boot** section of the **Partition Details** task.
 - i) Select **Storage device (SAN)** as the Boot from option.
 - ii) 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

In the **Configure Storage** task, when you open the Storage Group details page for an FCP storage group, you can review the adapters assigned to a storage group 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.

Instructions: Changing the adapters that are assigned to an FCP tape link

In the **Configure Storage** task, when you open the Tape Link details page for an FCP tape link, you can select the **ADAPTERS** tab to review the adapters that are assigned to the tape link. To change the adapters for a tape link, select **Modify** to open the Tape Link details page in Modification mode. Modification mode is indicated by a large blue bar that contains the Modification label, plus the **CANCEL** and **SAVE** buttons. In this mode, you can add or delete adapters only by changing the number of connecting paths. Remember that changes to this connectivity setting have an impact on bandwidth, performance, and redundancy. You also can replace currently assigned adapters with other available adapters, and resolve adapter mismatches when the fulfillment state of the tape link is Pending with mismatches.

When you save your changes, the fulfillment state of the tape link changes to Pending, and DPM automatically generates a request that you can send to one or more storage administrators; this email includes zoning instructions for the storage administrators to follow to fulfill your request.

The following topics describe the methods for changing the adapters for a tape link, and factors that you need to consider when using each method.

Modify the number of connecting paths


Changes that you make through the **Number of connecting paths** slider, text entry field, or spin button are reflected in the adapters table. The changes vary, depending on whether you are decreasing the number of paths, increasing the number of paths, or both, during a single modification session.


Only decrease the number of connecting paths

When you decrease the number of paths, DPM automatically grays out the appropriate number of adapter rows from the table, selecting the adapters to remove in the following order.

1. Adapters to be selected by the storage administrator (placeholder rows) that DPM assigned or that you selected to replace assigned adapters.
2. Unzoned adapters that are already assigned to the tape link.
3. Zoned ("Matching") adapters that are already assigned to the tape link.

4. Unzoned adapters that you selected to replace assigned adapters.
5. Zoned ("Matching") adapters that you selected to replace assigned adapters.


Each adapter that DPM selects for removal is marked with a recover icon () to the right of the table row. If you prefer to remove different adapters, select the recover icon for any marked adapter, and the table display changes to show each adapter row with a trash can icon, so you can select which assigned adapters to delete. An inline message tells you how many adapters you need to select for deletion. When you click the trash can icon, DPM marks the adapter to be removed with one of the following marks, to the left of the table row:

- A red dot to indicate the adapter that was most recently selected for removal.
- A warning icon () , when this change is disruptive to active partitions that are using the tape link.

Only increase the number of connecting paths

When you increase the number of connecting paths, DPM adjusts the Adapters table by adding the corresponding number of table rows; for example, if you increase the number of paths by 2, the table contains two new table rows.

- DPM preselects adapters that are connected to the tape library already, and that provide optimal redundancy. Optimal redundancy for adapters is based on the following factors, which are listed in priority order from highest to lowest: on the location in the I/O drawers, on the drawer domain, on the current allocation, and on the connection to SAN fabrics.
- If the total number of preselected adapters is less than the number of connecting paths that you selected, placeholder rows represent the remaining number of adapters for the storage administrator to assign. If you do not supply a specific adapter for a placeholder row, the storage administrator selects the adapter for you, as part of fulfilling your modification request.

If more adapters are available for selection, you can select the exchange icon () to replace a preselected adapter.



Increase and decrease the number of paths in a single modification session

If you increase and decrease the number of connecting paths within a single modification session, DPM automatically adds or removes Adapter table rows in a specific order.

- When you increase the number of connecting paths, DPM automatically adds the appropriate number of adapters, preselecting them in the following order.
 1. Adapters that you already deleted in this modification session, if any.
 2. Zoned ("Matching") adapters that were not assigned previously to this tape link.
 3. Adapters that were previously assigned to this tape link, but were replaced through the exchange dialog.
 4. Adapters to be selected by the storage administrator (placeholder rows).
- When you decrease the number of connecting paths, DPM automatically removes the appropriate number of adapter rows from the table, selecting the adapters to remove in the following order.
 1. Adapters to be selected by the storage administrator (placeholder rows) if you increased the paths in this modification session.
 2. Zoned ("Matching") adapters that DPM preselected if you increased the paths in this modification session.
 3. Remaining placeholder rows, if any, including placeholder rows that DPM assigned or that you selected to replace assigned adapters.
 4. Unzoned adapters that are already assigned to the tape link.
 5. Zoned ("Matching") adapters that are already assigned to the tape link.
 6. Unzoned adapters that you selected to replace assigned adapters.
 7. Zoned ("Matching") adapters that you selected to replace assigned adapters.

Replace assigned adapters

To replace an assigned adapter (or add an adapter to a placeholder row), complete the following steps. For the best results, try to assign adapters that reside in different system I/O drawers and in different domains; that are not on the same card; and that are connected to different fabrics. Inline messages provide warnings if the selected adapters do not meet these criteria.

1. Select the exchange icon () . A dialog opens and displays information about the assigned adapter or placeholder row that you want to exchange, and available adapters to replace it.
2. If you want a storage administrator to select an adapter for you, select **Adapter to be assigned by the storage administrator** under the Available adapters heading. Otherwise, use the information in the Available adapters table to select a replacement adapter. Note that any adapter with an existing error condition is marked with an incomplete icon () . The suggested practice is to avoid selecting such adapters.

MATCH

If an adapter is already connected to the tape library that you are using for this tape link, this column contains the label "Matching" in the adapter row. Otherwise, this column is empty.

NAME

Specifies the name of the FCP adapter. DPM assigns a default adapter name in the form *adapter_type pchid partial_location*, which can help you determine whether you are selecting adapters that, for optimal redundancy, reside in different drawers and different domains. For example, in the sample default name FCP 0171 Z22B-11:

- FCP is the type.
- 0171 is the physical channel path identifier (PCHID).
- Z22B is the plug location of the I/O drawer, with the first letter denoting the frame in which the drawer resides.
- 11 is the slot in the drawer in which the adapter is plugged.

FABRIC ID

Displays the worldwide name (WWN) of the uplink Fibre Channel switch. For optimal redundancy, use this value to select adapters that are connected to different fabrics.

ADAPTER ID

Specifies the physical channel ID of the adapter; this ID is a four-character hexadecimal number.

TYPE




Indicates the type of adapter card, which varies depending on the adapter cards that the system supports. Valid values include specific FICON Express adapter names

LOCATION

Specifies the physical location of the adapter in the I/O drawer of the system.

ALLOCATION

Indicates the percentage of host bus adapters (HBAs) that are currently allocated to this adapter, shown in a bar graph and in numeric format. The displayed percentage includes HBAs only for started and reserved partitions. If the percentage is high (for example, 90%), consider assigning a different adapter.

3. Select **Replace** to save your selection and close the dialog window. On the **ADAPTERS** tab, the table row for the adapter that you just selected is marked with one of the following icons.
 - A pending icon () indicates that zoning is required before the fulfillment state of the tape link can change to Complete.
 - An incomplete icon () indicates an existing error condition for the adapter. The suggested practice is to replace such an adapters with different ones.
 - A warning icon () , when this change is disruptive to active partitions that are using the tape link.

4. Repeat these steps, as necessary, to replace any more assigned adapters.

Resolve pending mismatches

When the fulfillment state of the tape link is Pending with mismatches, the **ADAPTERS** tab display contains three separate sections through which you can change the number of connecting paths, view and delete unrequested and zoned adapters, or view and replace assigned adapters. You can use any one of those three actions, or a combination of them, to resolve the pending mismatches.

For example, suppose that you requested a tape link with two connecting paths, and you selected two specific, available adapters for your tape link. If a third adapter is already zoned for the tape library that is selected for your tape link, the fulfillment state of the tape link is Pending with mismatches because three zoned adapters exceeds the requested number of two connecting paths. On the **ADAPTERS** tab, the third adapter is listed in the table of Unrequested and zoned adapters, which is displayed above the assigned adapters table. Through Modification mode, you can resolve this mismatch through one of the following actions.

- Increase the number of connecting paths to 3. In this case, the unrequested adapter becomes an assigned adapter, and no zoning changes are required.
- Select the unrequested adapter for deletion. In this case, when you save your changes, the storage administrator receives zoning instructions to remove the host WWPN for this unrequested adapter.
- Replace one of the assigned adapters with the unrequested adapter. This action can be useful if one of the assigned adapters is pending (not zoned yet).

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,” on page 71 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 161.
- 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 or a later DPM version is applied to the system. |
| 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. |

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 |
|---|---|--------------|--|
| Enable Dynamic Partition Manager |  | SE only | Use the Enable Dynamic Partition Manager task to enable DPM for the system and identify the dedicated OSA-Express 1000BASE-T Ethernet adapters for system management (OSM). |
| 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. If the DPM R3.1 storage management feature or a later DPM version is applied to the system, use the Configure Storage task to manage storage adapters and devices. |
| 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. |

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 |
|--|---|-----------------------------|---|
| 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. |
| 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 37. |
| 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 or Linux 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 or Linux systems.

Topics covered in this part are:

- [Chapter 10, “Prerequisites for using Dynamic Partition Manager,” on page 77](#)
- [Chapter 11, “Migrating DPM configuration data to a new system,” on page 87](#)
- [Chapter 12, “Adapter configuration,” on page 97](#)
- [Chapter 13, “Secure Execution for Linux feature,” on page 101](#)
- [Chapter 14, “Enabling FCP tape support for machine types 3906 and 3907,” on page 103](#)

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 and Linux 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 z15 (z15): machine types 8561 and 8562
- IBM z14 (z14): machine types 3906 and 3907
- IBM z13 (z13) or IBM z13s (z13s)
- IBM LinuxONE III: machine types 8561 and 8562
- 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 2.15.0. For additional details, see [“DPM versioning” on page 77](#). For information about using prior versions of DPM on other systems, see *IBM Dynamic Partition Manager Guide*, SB10-7170.

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 on page 77](#).

| Table 4. Required feature codes for DPM | |
|---|---|
| Feature code | Description |
| 0016 | Hardware for DPM feature |
| Feature codes for OSA adapters (codes vary, depending on the system configuration) | <p>Two dedicated OSA-Express 1000BASE-T Ethernet adapters are required for internal use by DPM, for primary and backup connectivity. For availability, the suggested practice is to locate the adapters on different I/O domains.</p> <p>The feature codes vary, depending on the OSA adapters that you are using for the DPM-enabled system. For example, you might use one OSA-Express5S (feature code 0417) and one OSA-Express6S (feature code 0426) to fulfill the requirement for two dedicated OSA adapters. For a list of the OSA adapters that are supported on a specific system, and their feature codes, see the appropriate system technical guide on the IBM Redbooks web site at http://www.redbooks.ibm.com/.</p> <p>For system connectivity to data center networks, you must configure additional network adapters. For more information, see “About network adapters” on page 82.</p> |

DPM versioning

DPM is available through specific engineering changes (EC) or machine change levels (MCL) for the mainframe and Linux systems that support it. MCLs are packaged and delivered in separate bundles for the HMC and SE. [Table 5 on page 78](#) lists the available DPM versions and the corresponding bundle that

must be applied to use any new functions or updates. The machine types indicate the systems on which you can apply the bundles.

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 versions

| DPM version / release | Machine type | HMC/SE version | Description |
|---------------------------|----------------------------|--|--|
| Version 4.3 (R4.3) | 8561, 8562, 3906, and 3907 | HMC/SE Version 2.15.0 (Driver 41) with the following bundles: <ul style="list-style-type: none"> • H25 or later • S38 or later | <p>Release of DPM R4.3 for the z15 and IBM LinuxONE III systems (machine types 8561 and 8562), which includes enhancements to the base FCP tape support that was previously released for machine types 3906 and 3907. (See the R4.3 table entry for the 3906 and 3907, below.) These enhancements include the following items.</p> <p>For machine types 8561, 8562, 3906, and 3907:</p> <ul style="list-style-type: none"> • Through the Storage Overview and Tape Link details pages in the Configure Storage task, the Resend Request action automatically generates a request that includes zoning instructions for incomplete or pending storage resources. • HMC Web Services APIs are available for all tape link functions. For more information, see the latest edition of <i>Hardware Management Console Web Services API</i>, which is available through the Library link on IBM Resource Link at http://www.ibm.com/servers/resourcelink <p>For machine types 8561 and 8562 only:</p> <ul style="list-style-type: none"> • On the Tape Link details page in the Configure Storage task, the HISTORY tab lists the actions that users have taken for the tape link. • For storage adapters that are in the Degraded state, the Status section of the Partition Details task displays more information and links to the details page of any storage groups or tape links that are affected. <p>Note: Although you can create and start a Secure Service Container partition that has tape links defined for access to FCP tape storage, Secure Service Container appliances do not support the use of tape storage devices.</p> |
| Version 4.3 (R4.3) | 8561 and 8562 | HMC/SE Version 2.15.0 (Driver 41) with the following bundles: <ul style="list-style-type: none"> • H20 or later • S34 or later | <p>Updated version of DPM for the z15 and IBM LinuxONE III systems. DPM R4.3 provides enhancements to the System Details task for managing the global and host keys for the IBM Secure Execution for Linux feature.</p> <p>Specifically, the Manage Secure Execution Keys task replaces the Import Secure Execution Key task. For more information, see Chapter 13, “Secure Execution for Linux feature,” on page 101.</p> |

Table 5. Summary of DPM versions (continued)

| DPM version / release | Machine type | HMC/SE version | Description |
|---------------------------|---------------|--|--|
| Version 4.3 (R4.3) | 3906 and 3907 | HMC/SE Version 2.15.0 (Driver 36) with H37 and S50 or later HMC/SE bundles. For instructions, see Chapter 14, “Enabling FCP tape support for machine types 3906 and 3907,” on page 103. | <p>Release of DPM R4.3 for machine type 3906 and 3907 systems only. DPM R4.3 provides system or storage administrators with a simplified interface to configure access to FCP tape libraries in the storage area network (SAN). This support replaces the temporary tape access method described in Appendix D, “Configuring and accessing FCP tape drives on a DPM-enabled system,” on page 169.</p> <p>The following Hardware Management Console (HMC) tasks are updated for FCP tape support.</p> <ul style="list-style-type: none"> • Through the Configure Storage task, you can create and manage tape links, which define the attributes of a connection that one or more partitions can use to access one FCP tape library in the SAN. You can also use this task to select partitions to which you want the tape link attached, and to manage FCP tape libraries in the DPM environment. • Through the Manage Adapters task, you can use the Connections section of Adapter Details to view the Tape Links table, which lists any tape links that are associated with the target adapter. • Through the New Partition task (basic and advanced modes) and Partition Details task, you can use the Storage section to select one or more tape links to attach to, or detach from, a specific partition. |

Table 5. Summary of DPM versions (continued)

| DPM version / release | Machine type | HMC/SE version | Description |
|---------------------------|---------------|--|--|
| Version 4.2 (R4.2) | 8561 and 8562 | HMC/SE Version 2.15.0 (Driver 41) with the following bundles: <ul style="list-style-type: none"> • H18 or later • S29 or later | <p>Updated version of DPM for the z15 and IBM LinuxONE III systems. Technical enhancements include support for Non-Volatile Memory Express (NVMe) storage adapters, which are available on IBM LinuxONE (LinuxONE) systems only. These systems must have one or more IBM Adapter for NVMe1.1 features enabled. An NVMe adapter consists of two pieces of hardware: an IBM-supplied carrier card installed in a system I/O drawer, and the solid state drive (SSD) that customers purchase.</p> <p>The following Hardware Management Console (HMC) tasks are updated to support the use of NVMe adapters. Note that the HMC Web Services API provide equivalent NVMe support.</p> <ul style="list-style-type: none"> • Through the Configure Storage task, you can view details about any NVMe adapters that are installed in the system; create an NVMe storage group for a partition to use; and view or modify details about the storage group. • Through the Dump task, you can select an SSD volume in an NVMe storage group as the boot volume for a dump program. • Through the Manage Adapters task, you can view information about NVMe storage adapters only when NVMe SSDs are installed in carrier cards in the system I/O drawers. • Through the New Partition (basic and advanced modes) and Partition Details tasks, you can attach one or more NVMe storage groups, and select an SSD volume in an attached NVMe storage group as the boot volume for an operating system or hypervisor. If you plan to boot an operating system or hypervisor from a volume in an NVMe storage group, check the product documentation for the operating system or hypervisor to make sure that it supports booting from NVMe storage. Note that only one partition can use an NVMe storage group at any given time; an NVMe storage group cannot be shared. <p>In addition to the NVMe support, DPM R4.2 also delivers performance improvements that significantly reduce the time required to discover logical units (LUNs) during storage management task processing.</p> |

Table 5. Summary of DPM versions (continued)

| DPM version / release | Machine type | HMC/SE version | Description |
|---------------------------|---------------|---|--|
| Version 4.1 (R4.1) | 8561 and 8562 | HMC/SE Version 2.15.0 (Driver 41) with the following bundles: <ul style="list-style-type: none"> • H14 or later • S22b or later | <p>Updated version of DPM for the z15 and IBM LinuxONE III systems. Technical enhancements include the IBM Secure Execution for Linux feature, which supports Linux (KVM) hypervisors. Secure execution must be configured in the Linux operating system and enabled on the system; for more details, see Chapter 13, “Secure Execution for Linux feature,” on page 101.</p> <p>System Details task The General section of the System Details task displays a new field, Secure Execution, which indicates whether the IBM Secure Execution for Linux feature is enabled on this system. For the feature to be enabled, a global key and host key must be installed on the Support Element.</p> <p>This indicator is also available on the Systems tab on the HMC Systems Management view, but the Secure Execution column on that tab is not displayed in the predefined default table view. To display the Secure Execution column, select the Manage Views icon in the work pane table toolbar to customize the table view.</p> <p>Partition Details task The General section of the Partition Details task displays a new field, Secure Execution, which indicates whether the operating system that runs on a partition is configured for secure execution, which isolates and protects any guests that run on a hypervisor by restricting host access to guest workloads and data.</p> <p>Configure Storage task For FICON connections, the Connect Adapter Ports dialog has enhanced search functions to help you locate the adapter ports that connect a system to a switch or storage subsystem.</p> |

Table 5. Summary of DPM versions (continued)

| DPM version / release | Machine type | HMC/SE version | Description |
|--------------------------------------|---------------|--|--|
| Version 4.0 (R4.0) | 8561 | HMC/SE Version 2.15.0 (Driver 41) with the following bundles: <ul style="list-style-type: none"> • H07 or later • S10 or later | <p>Version of DPM for the z15 and IBM LinuxONE III systems. Usability, simplification, and technical enhancements include the following items.</p> <p>Configure Storage task</p> <ul style="list-style-type: none"> • An exportable file of cabling details that you can use to physically connect the system to SAN hardware. • New navigation aids and displays for configuring storage adapter cards that are installed in a multiple-frame system. • A new page to view and resolve device number conflicts between base and alias volumes in a FICON storage group. • A new page to view and copy information for enabling an operating system or installer to access a specific volume. • A new UI display through which users can select one or more ranges of logical control units (LCUs) when configuring system connections to storage devices. • The Resend Request option to create a modification request that identifies actions for a storage administrator to perform to change the fulfillment state of a storage group to Complete. <p>New Partition and Partition Details tasks</p> <ul style="list-style-type: none"> • A new option that enables applications running on the partition to generate and manage Elliptic Curve Cryptography (ECC) protected keys through the CP Assist for Cryptographic Functions (CPACF) feature. • A new option, Secure Boot, through which administrators can request DPM to verify the software signature of the Linux operating system that is to run in a partition. • Detection of features that are not available on all supported systems: HMC DVD drives and zEnterprise Data Compression (zEDC) accelerators. |
| Version 3 (multiple releases) | 3906 and 3907 | HMC/SE Version 2.14 (Driver 36) | <p>For descriptions of the support added with each Version 3 release, see the versioning topic in <i>IBM Dynamic Partition Manager Guide</i>, SB10-7170. The Version 3 releases include:</p> <p>R3.2 R3.1 R3.0</p> |
| Version 2 (multiple releases) | 2964 and 2965 | HMC/SE Version 2.13 (Driver 27) | <p>For descriptions of the support added with each Version 2 release, see the versioning topic in <i>IBM Dynamic Partition Manager Guide</i>, SB10-7168. The Version 2 releases include:</p> <p>R2.1 R2.0</p> |

About network adapters

In addition to the dedicated OSA-Express 1000BASE-T Ethernet adapters that are required by DPM feature code 0417, you must configure additional network adapters for system connectivity to data center networks. 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) adapters, 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 adapters. These adapters provide high speed, low latency data transfer over Ethernet networks.

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 technical guide on the IBM Redbooks web site at <http://www.redbooks.ibm.com/>.

DPM automatically detects OSA and RoCE adapters that are installed in the system. Use the HMC **Manage Adapters** task to manage these adapters. 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 FICON and FCP storage adapters

Mainframe and Linux systems use Fibre Channel connections (FICON) to provide high-speed data transfer to external storage devices. Through cables, FICON Express features (adapter cards) connect a DPM-enabled system to the devices in the storage area network (SAN). DPM automatically detects FICON Express adapters that are installed in the system. 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 technical guide 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. Support for FCP mode is available with all DPM versions.
 - Note that, because DPM runs all FCP adapters in N Port Identifier Virtualization (NPIV) mode, the SAN must support NPIV.
 - In releases prior to R4.3, DPM detects but does not display FCP tape drives, so administrators cannot directly manage them through DPM tasks on the HMC. Starting with DPM R4.3, however, administrators can use the **Configure Storage** task to configure and manage access to tape storage.
- FICON native (FC or FICON) mode provides access to extended count key data (ECKD) 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). Support for FICON mode is available with DPM R3.1 and later DPM versions, 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 drives.

Starting with DPM R3.1, the Hardware Management Console (HMC) **Configure Storage** task replaces the use of the **Manage Adapters** task to configure and manage FCP and FICON storage adapters.

About NVMe storage adapters

Starting with DPM R4.2, your company can order IBM Adapter for NVMe1.1 features, which are available on IBM LinuxONE (LinuxONE) systems only. Non-Volatile Memory Express (NVMe) storage adapters use the PCI Express (PCIe) protocol to provide high-speed storage within a system.

Each NVMe adapter consists of two pieces of hardware: an IBM-supplied carrier card installed in a system I/O drawer, and the solid state drive (SSD) that customers purchase. Although you can use the **Configure Storage** task and the **Manage Adapters** task to view information about NVMe storage adapters, you cannot use those tasks to configure them; reconfiguration requires properly removing the carrier card and its SSD from the drawer and reinstalling them in a different physical location, as instructed by a service representative.

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 Documentation, at: <https://www.ibm.com/docs/en/linux-on-systems?topic=linuxone-library-overview>

- 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. Note that, depending on the system, a DVD drive might not be available on the HMC. In this case, you can use an alternate HMC media drive or use an electronic deliverable instead of a DVD.
 - 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 or **ISO image** 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 25. For information about z/VM, see the product documentation. The z/VM library is available in IBM Documentation at <https://www.ibm.com/docs/en/zvm/7.1>

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 z15, the 8561DEVICE or 8562DEVICE PSP bucket
 - For a LinuxONE III, the 8561DEVICE or 8562DEVICE PSP bucket
 - 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 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. Only specific migration paths are supported for an MES frame upgrade.

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” on page 87](#).

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. For instructions, see [“Migration instructions for a system replacement” on page 91](#).

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 77](#)
 - [“Where to find hardware planning and corequisite software information” on page 84](#)
- [Table 6 on page 88](#) indicates whether a specific DPM migration path is supported for an MES frame upgrade.

- Each DPM version corresponds to particular machine types; for more information about DPM versions and releases, see [“DPM versioning”](#) on page 77. Your service representative can provide the existing and new machine types for which MES frame upgrades are supported.
- Downgrades from one version or release to a prior version or release are not supported options. For example, you cannot downgrade from Version 3 Release 1 (R3.1) to Version 2 Release 1 (R2.1), or from R3.2 to R3.1.
- Migrating to the latest release of a given version is the only supported option; for example, you cannot migrate from DPM R2.0 to R3.1 because R3.2 is the latest Version 3 release.

Table 6. Supported DPM migration paths for MES frame upgrades

| From ↓ | To → | DPM R3.2 | DPM R4.0 or later |
|-----------------|------|---|---|
| DPM R2.0 | | Yes (with HBAs converted to FCP storage groups) | Yes (with HBAs converted to FCP storage groups) |
| DPM R2.1 | | Yes (with HBAs converted to FCP storage groups) | Yes (with HBAs converted to FCP storage groups) |
| DPM R3.1 | | No | Yes |
| DPM R3.2 | | — | Yes |

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 a hardware message for the new system. This message includes the following information:

- A list of partitions that were defined as 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 using accelerator adapters. This list is included only for an MES frame upgrade to DPM R4.0 or later. Accelerator adapters are not supported on the systems on which you can run DPM R4.0 or later, so DPM deletes the accelerator virtual functions 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 on page 89](#) shows a sample hardware message for the automated DPM process.

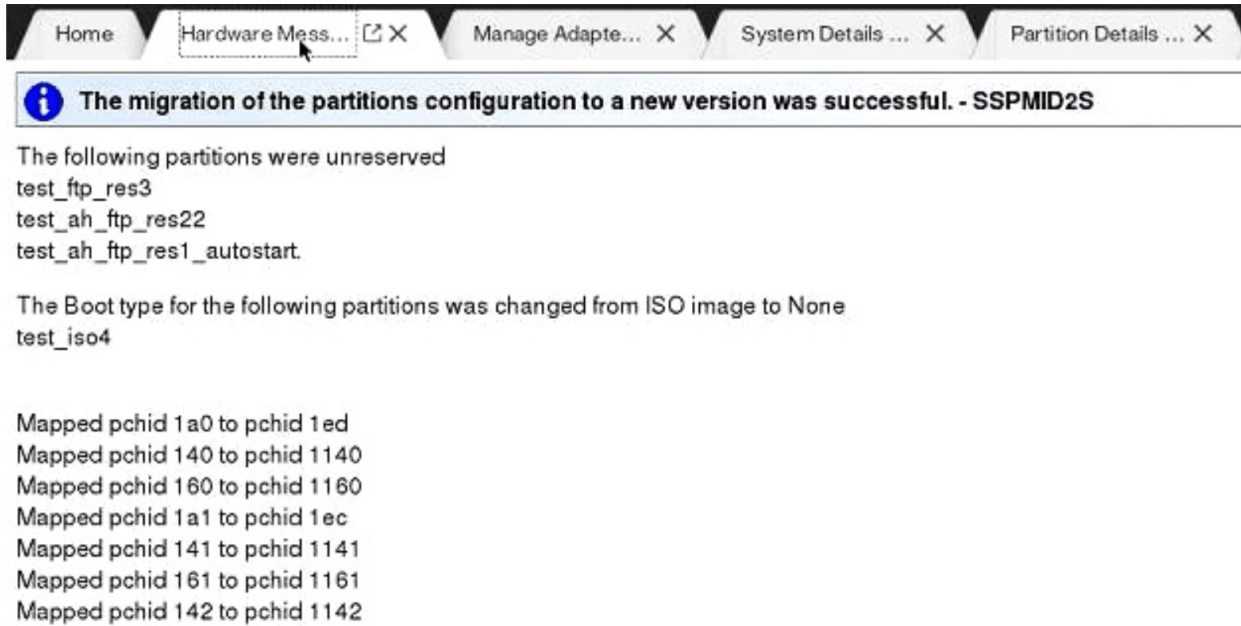


Figure 37. Sample hardware message that shows the results of a successful DPM conversion of migrated configuration information

Procedure

1. If you are using two new dedicated OSA-Express 1000BASE-T Ethernet adapters (required feature 0417) on the new system, configure these new system management (OSM) adapters.
These two adapters are required for internal use by DPM, for primary and backup connectivity. For availability, the suggested practice is to locate the adapters on different I/O domains.
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.

- i) 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.
- ii) Select **Reassign Devices** from the **Actions** list.
- iii) On the **Reassign Devices** window, select the new backing adapter for the devices, and click **Reassign**.
- iv) 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.
 - i) 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.
 - ii) 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.
 - iii) 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.
 - i) Select or scroll to the **Boot** section, and select **ISO image** from the **Boot from** list.
 - ii) 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.
5. Review any new or existing storage groups in the DPM configuration on the new system.

Open the **Configure Storage** task to the Storage Overview and check the fulfillment state of each storage group. Take any actions that might be required to change the state to Complete. For more information, see the online help for the task or go to [“Viewing and managing storage groups and tape links” on page 143](#).

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.

Before you begin

- Table 7 on page 91 indicates whether a specific DPM migration path is supported for a system replacement.
 - Each DPM version corresponds to particular machine types; for more information, see “DPM versioning” on page 77.
 - Downgrades from one version or release to a prior version or release are not supported options. For example, you cannot downgrade from Version 3 Release 1 (R3.1) to Version 2 Release 1 (R2.1), or from R3.2 to R3.1.
 - Migrating to the latest release of a given version is the only supported option; for example, you cannot migrate from DPM R2.0 to R3.1 because R3.2 is the latest Version 3 release.

Table 7. Supported DPM migration paths for a system replacement

| From ↓ | To → | DPM R3.2 | DPM R4.0 or later |
|-----------------|------|--|---|
| DPM R2.0 | | Yes (with HBAs converted to FCP storage groups) | Yes (with HBAs converted to FCP storage groups) |
| DPM R2.1 | | Yes (with HBAs converted to FCP storage groups) | Yes (with HBAs converted to FCP storage groups) |
| DPM R3.0 | | No; this path is only an upgrade on the same machine type. | Yes (with HBAs converted to FCP storage groups) |
| DPM R3.1 | | No; this path is only an upgrade on the same machine type. | Yes |
| DPM R3.2 | | — | Yes |

- When you migrate from DPM R3.0 or an earlier version to DPM R3.2 or a later version, each host bus adapter (HBA) that a partition uses is 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:
 - `zhmccclient`, 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/zhmccclient>.
 - `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/resourceLink/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.

- If you are migrating from DPM R3.1 or later, the configuration data also includes storage groups and templates, as well as defined FICON connections, such as sites, subsystems, fabrics, volumes, and paths.
- If you are migrating from DPM R4.3 or later, the configuration data also includes FCP tape links.

When you migrate from any supported version of DPM to a later version, adapter IDs are changed. Additionally, when you migrate from DPM R3.0 or an earlier version to a target system with DPM R3.2 or later, each HBA that a partition uses is converted into a dedicated FCP storage group. You need to review these adapter ID changes and new storage groups, and modify them, if necessary.

- To understand the differences between storage access with DPM R3.0 or an earlier version and storage access with DPM R3.2 or later, see [Figure 38 on page 93](#).
- For more information about storage management starting with DPM R3.1, see [Part 6, “Topics for storage administrators,” on page 117](#).



Before migration**After migration**

Figure 38. Partition access to storage when migrating from an early DPM version to DPM R3.1 or later

In Figure 38 on page 93:

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 94 in the following procedure).
5. Note that, on a target system with DPM R3.1 or later, 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.

Procedure

1. For your script, install the latest version of the `zhmccclient` 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. For variable descriptions and a sample generated file from the `exportDpmResourcesToFile.py` script, see the accompanying PDF on the download page in Resource Link: <https://www.ibm.com/servers/resourceLink/lib03020.nsf/pages/zHmcWebServicesExportDpmConfig?OpenDocument>

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 the latest edition of *Hardware Management Console Web Services API*, 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.

If you migrated from DPM R3.1 or later to a target system with a later DPM version

If you are using an auto-start list to start partitions on the target system, and one or more of those partitions are booted from a storage volume, the start process might fail if the attached storage group is not yet in the Complete fulfillment state. To reduce the possibility of partition start failures, DPM assigns highest priority to storage groups that contain at least one boot volume, so that automatic discovery of the logical unit numbers (LUNs) begins as quickly as possible on the target system.

If you migrated from DPM R3.0 or earlier to a target system with DPM R3.2 or later



- 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 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 93](#) 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 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 detected 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**.

Check the fulfillment state of the storage groups 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 [“The Storage Group details page” on page 144](#) (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 64](#).

Chapter 12. 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 appropriate HMC task to review details about the adapters, and make any adjustments that might be necessary.

- For FICON Express storage adapters, use the **Configure Storage** task to configure storage adapters as either FCP or FICON, and to specify which adapters connect the system to storage devices.
- To view any installed Non-Volatile Memory Express (NVMe) storage adapters, use the **Configure Storage** task. Note that you cannot use this task to configure NVMe adapters; reconfiguration requires properly removing the carrier card and its SSD from the drawer and reinstalling them in a different physical location, as instructed by a service representative.
- To view details and modify network or crypto adapters, use the **Manage Adapters** task.
- If HiperSockets are required on the system, you need to define them. See [“Creating a HiperSockets adapter”](#) on page 99 for an overview of the procedure.

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 8 on page 97](#) provides examples of the default names that DPM assigns to installed adapters. In the table, PCHID refers to the physical channel path identifier and, for some adapter types, Location includes jack and slot information.

| Table 8. 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 |
| NVMe | — | 0160 | NVMe | A23B-LG12 | NVMe 0160 A23B-12 |
| 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 |
| zEDC¹ | — | 0188 | zEDC Express | Z15B-LG03 | zEDC 0188 Z15B-03 |

Note:

1. This sample entry represents an accelerator adapter. Only specific systems support accelerators. To determine whether a specific system supports accelerators, see the appropriate system technical guide on the IBM Redbooks web site at <http://www.redbooks.ibm.com/>.

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 39](#) on page 98.

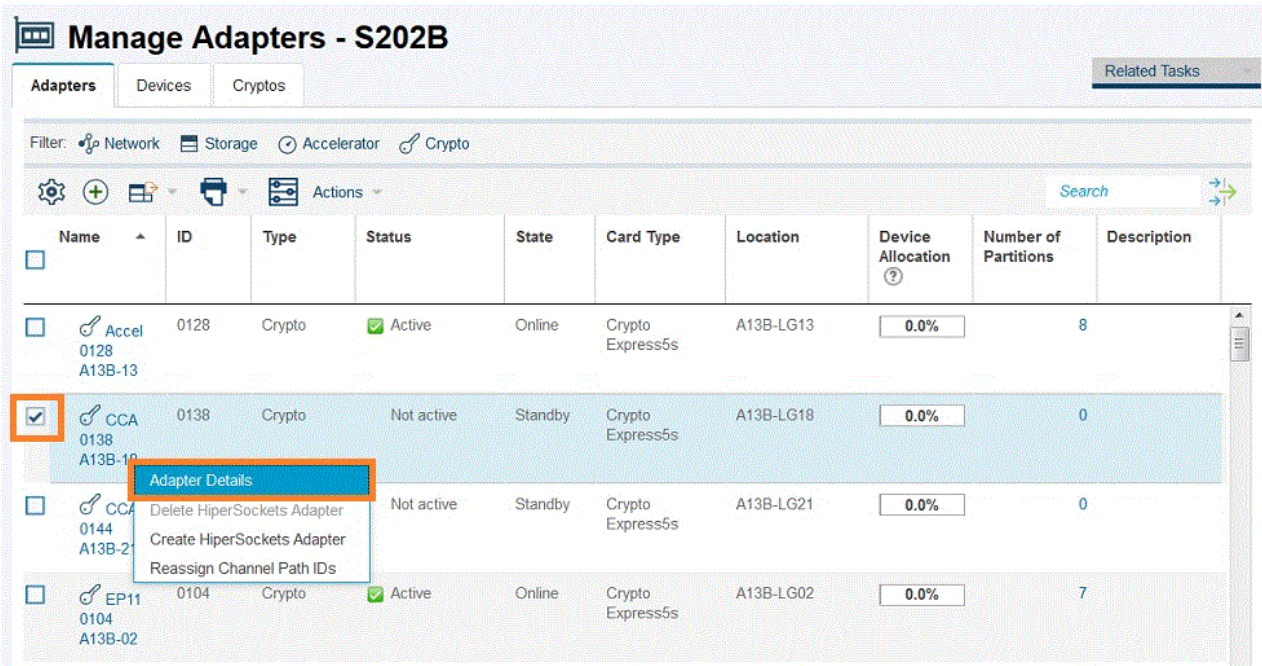


Figure 39. Sample screen of the Manage Adapters task

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 40](#) on page 99 shows the adapter details for a cryptographic adapter. After you modify the name or add a description, click **Apply** to save your changes.

Creating a HiperSockets adapter

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 |
|--|------------|----------|------------------|----------------|-------------------|----------------------|-------------|
| <div> <div> <div></div> <div> <div>Adapter Details</div> <div>Delete HiperSockets Adapter</div> <div>Create HiperSockets Adapter</div> <div>Reassign Channel Path IDs</div> <div>Configure Options</div> <div>Select All</div> <div>Deselect All</div> <div>Expand All</div> <div>Collapse All</div> <div>Clear Sort</div> </div> </div> <div> <div></div> <div> <div>F10131</div> <div>A13B</div> </div> </div> </div> | Not active | Reserved | FICON Express16s | A13B-D216-J.01 | 0% | 0 | |
| <div> <div> <div></div> <div> <div>Adapter Details</div> <div>Delete HiperSockets Adapter</div> <div>Create HiperSockets Adapter</div> <div>Reassign Channel Path IDs</div> <div>Configure Options</div> <div>Select All</div> <div>Deselect All</div> <div>Expand All</div> <div>Collapse All</div> <div>Clear Sort</div> </div> </div> <div> <div></div> <div> <div>F10134</div> <div>A13B</div> </div> </div> </div> | Not active | Reserved | FICON Express16s | A13B-D117-J.01 | 0% | 0 | |
| <div> <div> <div></div> <div> <div>Adapter Details</div> <div>Delete HiperSockets Adapter</div> <div>Create HiperSockets Adapter</div> <div>Reassign Channel Path IDs</div> <div>Configure Options</div> <div>Select All</div> <div>Deselect All</div> <div>Expand All</div> <div>Collapse All</div> <div>Clear Sort</div> </div> </div> <div> <div></div> <div> <div>F10135</div> <div>A13B-17</div> </div> </div> </div> | Not active | Reserved | FICON Express16s | A13B-D217-J.01 | 0% | 0 | |

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:

Description:

MTU/Frame size (KB):

OK Cancel Help

Figure 42. 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:

Description:

Object ID: A56C8508-F06F-11E5-9051-9ABE94228451

System: S223B

Status: Not active

Detailed status: Not active

State: Standby

NIC allocation:

Adapter ID: 07C0

Adapter type: Internal queued direct communications (HiperSockets)

Channel path ID (CHPID): 48

Card type: HiperSockets

MTU/Frame size (KB):

Figure 43. Sample screen of the new HiperSockets adapter in the Adapter Details window

Chapter 13. Secure Execution for Linux feature

With the IBM Secure Execution for Linux feature, you can deploy multiple secure, isolated virtual environments on a single system. Secure Execution isolates and protects any guests that run on a hypervisor by restricting host access to guest workloads and data. It protects the confidentiality and integrity of guest data from both internal and external attacks, but does not protect against a denial of service.

When you order the IBM Secure Execution for Linux feature, IBM either installs the required global and host key bundles on a new machine before delivery, or electronically sends them to your company for the service representative to import to an already installed system. For a new machine with the feature and keys installed, Secure Execution is enabled on the system, and hypervisors that are configured to use the feature can successfully do so. For an existing system at your site, however, you first need to complete the following steps before Secure Execution is enabled and ready for use.

1. Order the feature.
2. Install the feature, using the **Perform Model Conversion** task.
3. Import the key bundles.

Importing the key bundles requires logging in to the Support Element using the SERVICE user ID or an ID with equivalent permissions, and opening the **System Details** task to the General section. This section includes a button through which you can open a separate task to view or manage the keys. Depending on the latest service that you have applied, the button label is either **Manage Keys** or **Import**.

With the latest bundle for DPM R4.3 or a later release

Manage Keys opens the **Manage Secure Execution Keys** task on the HMC or SE in read-only mode, through which you can view details about each key. If you are logged in to the SE using the SERVICE user ID or an ID with equivalent permissions, you can import the required global key or host key bundle; view the hashes for the existing global or host key; or clear the global or host key, which immediately prevents further usage by any partition on the system, and deletes the corresponding key bundle file from the system.

With DPM R4.2 and R4.1

On the SE only, **Import** opens the **Import Secure Execution Keys** task, through which you can import either the global key bundle or host key bundle. To import keys, you must be logged in to the SE using the SERVICE user ID or an ID with equivalent permissions.

Generally, you need to import the keys only once, because keys are reimported during initial microcode load (IML).

- If possible, do not change the bundle file names when you receive the files.
- The bundle content is created for a specific system only. Make sure that you import the bundles to the system for which they were created, and do not attempt to import those same bundles to any other system.
- You can import only one key bundle at a time, from an FTP server or a removable media device.

For more information about the IBM Secure Execution for Linux feature, see *Introducing IBM Secure Execution for Linux 1.1.0*, SC34-7721, which is available in IBM Documentation at: <https://www.ibm.com/docs/en/linux-on-systems?topic=overview-introducing-secure-execution-linux>.

- To determine whether a system has this feature installed, open the **System Details** task to the General section, and check the Secure Execution indicator. This indicator is also available on the **Systems** tab on the HMC **Systems Management** view, but the Secure Execution column on that tab is not displayed in the predefined default table view. To display the Secure Execution column, select the **Manage Views** icon in the work pane table toolbar to customize the table view.
- To configure Linux for secure execution, see the product documentation for the Linux distribution that you are using as a hypervisor. This feature supports Linux (KVM) hypervisors.

- To determine whether the Linux hypervisor that runs on a partition is configured for secure execution, go to the **General** section of the **Partition Details** task and check the value displayed for the Secure Execution field.

Chapter 14. Enabling FCP tape support for machine types 3906 and 3907

FCP tape support requires Hardware Management Console (HMC) hardware with Version 2.15.0 and the minimum H20 service bundle applied. If you want to exploit the FCP tape support for workloads that run on machine type 3906 or 3907 systems that are managed by HMCs with Version 2.14.1, you need to upgrade the firmware for at least one of those Version 2.14.1 HMCs to Version 2.15.0 with correct service bundle. This upgrade is firmware only; a hardware upgrade is not required.

If you already have Version 2.15.0 HMCs because you ordered them for your 8561 or 8562 systems, you can use those HMCs to manage the machine type 3906 and 3907 systems as well. In this case, you do not have to upgrade any HMCs; installing the minimum H20 service bundle is sufficient.

Otherwise, if the HMCs that you are using have Version 2.14.1 firmware, work with a service representative (or with the assistance of the support system) to upgrade the firmware level of at least one HMC.

1. Upgrade all HMCs to Version 2.15.0 with the H20 service bundle for driver 41.
2. If a subset of the HMCs that manage the earlier systems cannot be upgraded to driver 41, install the H37 service bundle for driver 36 to avoid feature incompatibilities. Even after installing this service bundle, make sure that you **do not use** these remaining Version 2.14.1 HMCs to manage the 3906 or 3907 systems that are connected to FCP tape libraries.
3. Apply the S50 service bundle for driver 36 to the Support Elements (SEs) of 3906 and 3907 systems through which you want to connect to FCP tape libraries in the storage area network (SAN).

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 15, “Customizing the network boot server environment,” on page 107](#)
- [Chapter 16, “Securing FTP operations,” on page 111](#)
- [Chapter 17, “Using VLAN IDs to secure network communications,” on page 113](#)

Chapter 15. 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";
}
```

- ii) 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" on page 108.

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 **tfpd**. 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=zkvm251 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:
 - i) Select the **Network server (PXE)** option from the "Boot from" menu list.
 - ii) 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 zkvm251 {
    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 16. 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” on page 111](#) or [“Setting up SFTP” on page 111](#).

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**.
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.
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 17. 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 44 on page 113, 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.

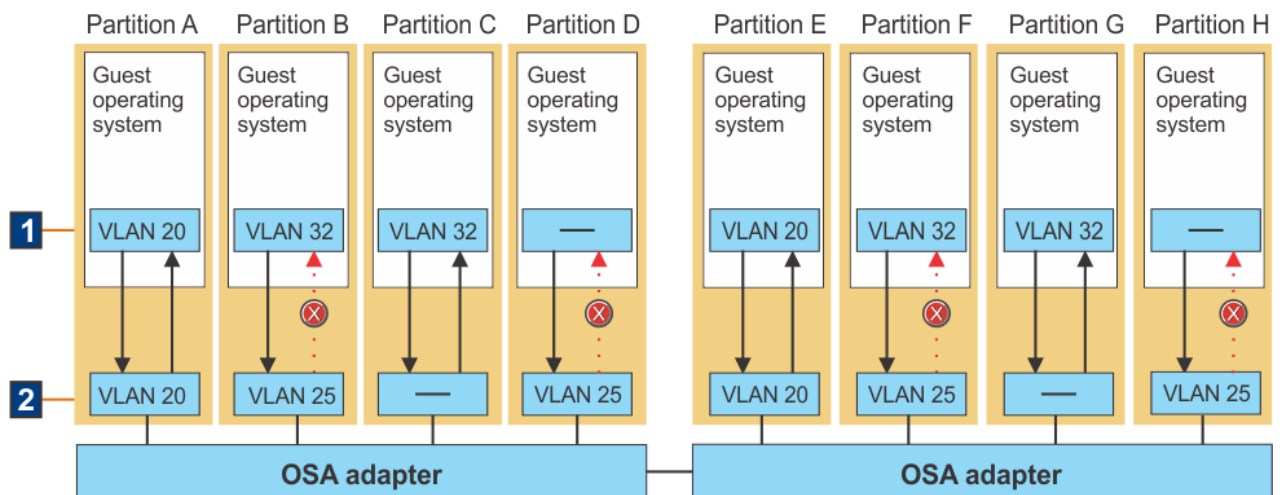


Figure 44. Sample VLAN configuration in partitions and guest operating systems

- 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-OSD`, 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-0SD
TYPE=vlan
ONBOOT=no
VLAN_ID=700
IPV6INIT=no
```

Part 6. Topics for storage administrators

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 mainframes or Linux systems. 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 the **Manage Adapters** task to configure and manage storage adapters. Although multiple systems can share the same storage resources, note that the scope of this task is *one single system only*.

Topics in this part provide information about storage-related concepts and tasks that are specific to working with DPM-enabled systems. These topics apply to supported types of Fibre Connection (FICON) extended count key data (ECKD) direct-access storage devices (DASD), and Fibre Channel Protocol (FCP) Small Computer System Interface (SCSI) disk storage devices and FCP tape storage devices.

- [Chapter 18, “The Configure Storage task: a simplified and improved approach to storage management,” on page 119](#)
- [Chapter 19, “Setting up storage resources for the first time,” on page 129](#)
- [Chapter 20, “Managing storage resources after the initial configuration is complete,” on page 143](#)

Chapter 18. The Configure Storage task: a simplified and improved approach to storage management

Starting with R3.1, DPM introduced a simplified and improved way of configuring and managing storage. For an introduction to storage management through the HMC **Configure Storage** task, see the following topics.

- [“A simplified, visual process for connecting storage to one system” on page 119](#)
- [“Integrated requests and notifications” on page 122](#)
- [“Access to FICON and FCP disk storage through storage groups” on page 123](#)
- [“Access to FCP tape storage through tape links” on page 124](#)
- [“Access to NVMe storage through storage groups” on page 126](#)
- [“Storage group templates” on page 127](#)
- [“Automatic detection and assignment of FICON and FCP adapters” on page 127](#)
- [“Device number allocation for FICON storage groups” on page 127](#)
- [“Storage Overview and details” on page 128](#)

A simplified, visual process for connecting storage to one system

To accomplish the initial setup of storage resources for a DPM-enabled system, system and storage administrators use the **CONNECT TO STORAGE** wizard of the **Configure Storage** task. For access to external storage devices, 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 storage hardware devices in the SAN and their FICON connections to this system. These hardware devices include disk storage subsystems, fabrics, and switches. 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.

A simplified, visual process for connecting storage to one system

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. For the **Configure Storage Cards** page, DPM automatically detects and displays all installed storage adapters, which include the following types.

- FICON Express adapter cards connect the system to devices in the storage area network (SAN) through Fibre Channel connections. Each FICON adapter card is configured to operate in a specific mode, which determines the type of external storage device that can be connected to the adapter card.
 - Fibre Channel Protocol (FCP) mode provides access to Small Computer System Interface (SCSI) devices, which include tape libraries as well as disk storage.
 - FICON native (FC or FICON) mode provides access to extended count key data (ECKD) devices (also known as direct-access storage devices, or DASD).
- Non-Volatile Memory Express (NVMe) storage adapters use the PCI Express (PCIe) protocol to provide high-speed storage within a system. These storage adapters are available on IBM LinuxONE (LinuxONE) systems only, when those systems have IBM Adapter for NVMe1.1 features installed.

[Figure 45 on page 120](#) shows a sample visual display of one drawer on the **Configure Storage Cards** page.

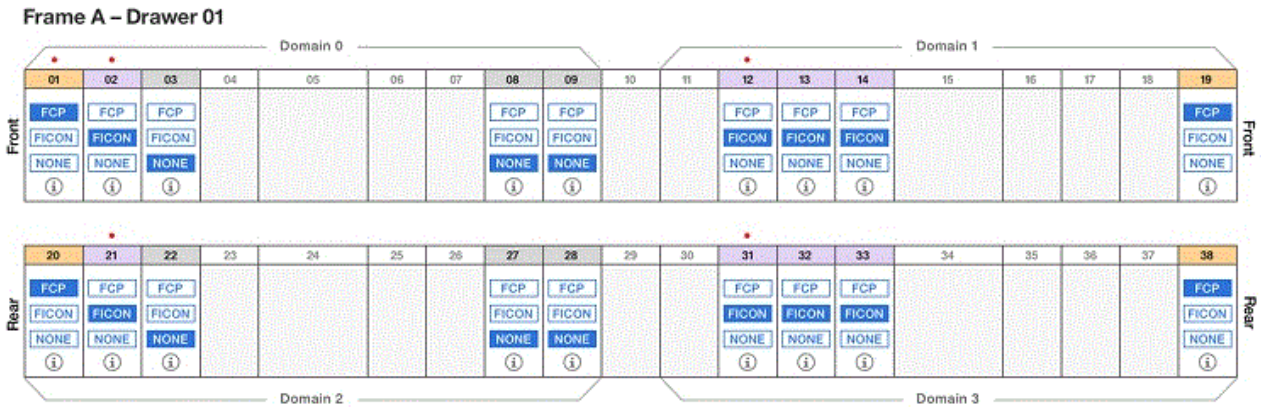


Figure 45. Sample frame and card display on the **Configure Storage Cards** page

For a comparison, look at Figure 46 on page 120, 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 45 on page 120 mirror the size and location of actual slots in the photograph.



Figure 46. Photograph of a physical adapter card in a system I/O drawer

In the **Configure Storage Cards** page shown in Figure 45 on page 120, the adapter cards that you can configure contain selectable labels. You can manually configure cards by clicking the FCP or FICON label to select the protocol. Or, for an automated approach, you can use other controls on the page to provide a total number of FICON or FCP cards 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. Red dots in the drawer display indicate the most recently configured FICON or FCP adapter cards.

Depending on the configuration of the system, you might need to use navigation controls to view all of the frames, drawers, and cards on the **Configure Storage Cards** page. For single-frame systems, use the scroll bar or expand/collapse controls to view adapters in the frame drawers. For multiple-frame systems, use the overview map to change the viewport display, as shown in Figure 47 on page 121.

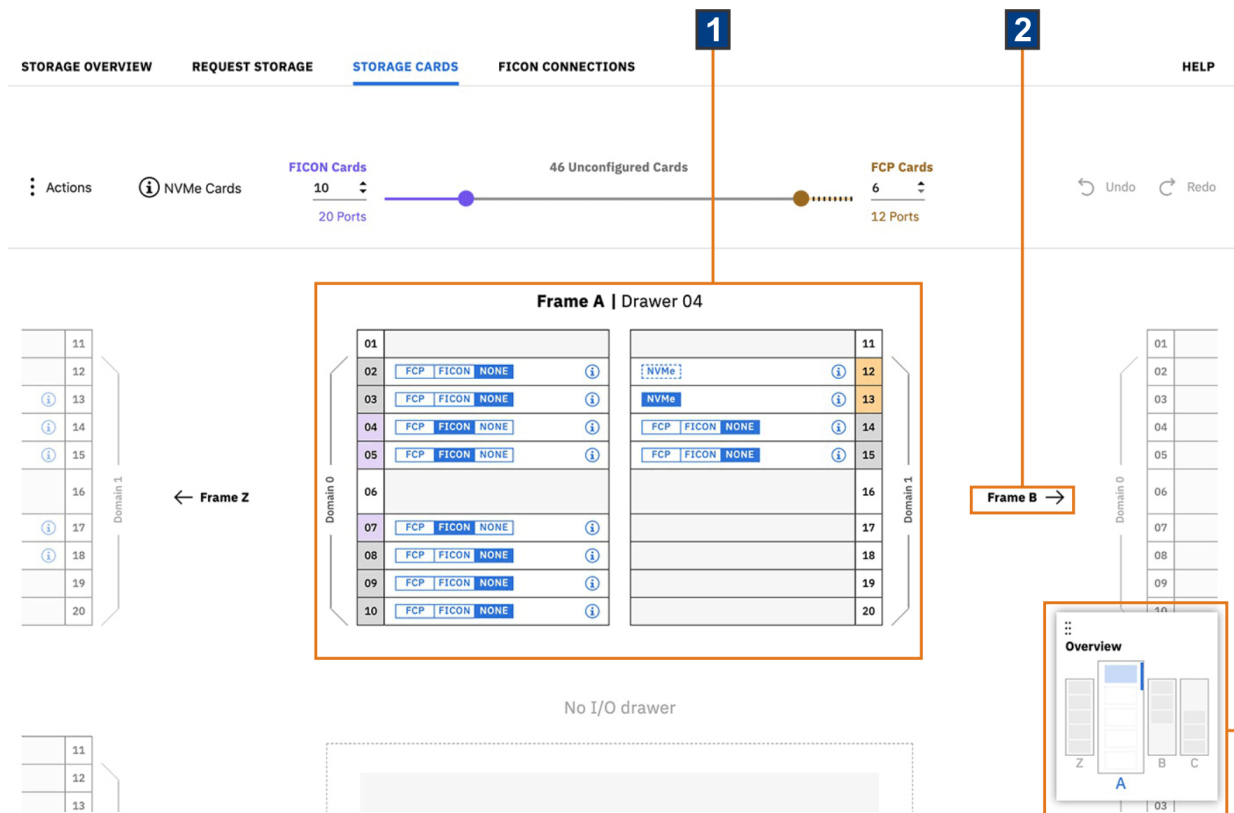


Figure 47. Overview map and other viewport navigational controls

1. Figure 47 on page 121 shows Frame A centered in the viewport.
2. Selectable frame buttons on either side of Frame A provide a way to change the viewport to either the previous or next frame in the system. These buttons are displayed only when the system is configured with additional frames to the left or right of the current frame.
3. The overview map shows how many frames are in the system, how many I/O drawers are in each frame, and highlights the I/O drawer that is currently displayed in the viewport. Red dots in the overview map indicate recent changes to a specific frame that contains an I/O drawer with storage adapter cards. To change the display, select a different frame in the overview map or use the frame buttons. Scroll up or down to view the drawers within the current frame.

Note that, in Figure 47 on page 121, Frame A contains NVMe cards as well as FICON and FCP adapter cards. If the system has one or more IBM Adapter for NVMe1.1 features, the frame display also indicates the current location of any installed NVMe carrier card. Although you can view NVMe cards in the page display, you cannot configure them.

After configuring FICON or FCP adapter cards, the next step is to configure the FICON storage devices in the SAN that are or will be connected to the DPM system, such as disk storage subsystems, fabrics, and switches. (Do not define tape libraries as part of the FICON connections in this step; DPM generates the required connections to tape storage through tape links instead.) 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 122 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.

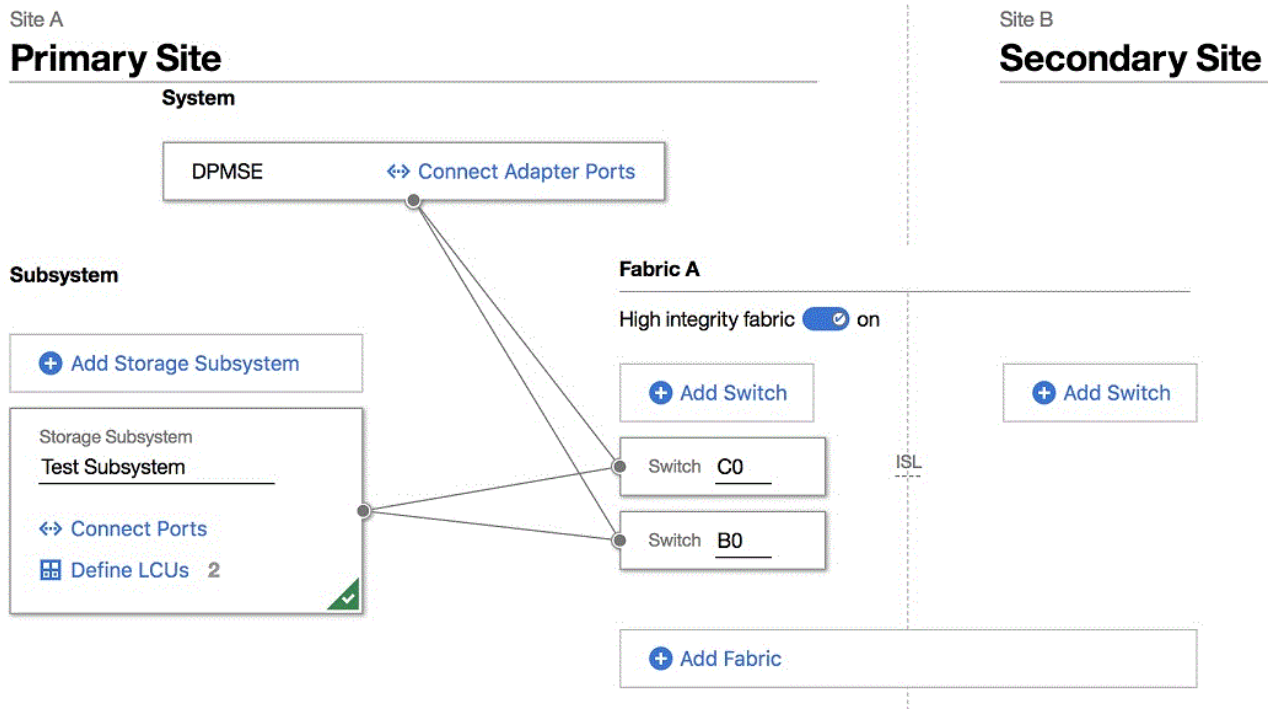


Figure 48. Sample display of two sites with physical storage hardware configured for the primary site

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.
- An exportable file of an FCP and FICON adapter cabling plan that you can use to physically connect the system to SAN hardware. The file is in Comma Separated Values (CSV) format that you can view in a spreadsheet application. For a sample file, see [Appendix B, "Sample System Cabling Plan,"](#) on [page 165](#).

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 129](#).
- For information about setting up SMTP, see the online help for the **Monitor System Events** task.

Access to FICON and FCP disk storage through storage groups

Administrators use the **Configure Storage** task to configure access to Fibre Connection (FICON) extended count key data (ECKD) direct-access storage devices (DASD), and Fibre Channel Protocol (FCP) Small Computer System Interface (SCSI) disk storage 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; another 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 knowledge of mainframes or Linux systems; however, a storage administrator must fulfill any requests for storage. You can create one or more storage groups through a single request by defining one storage group and specifying the number of duplicates.

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 123](#) shows several storage groups with key attributes, and information about the physical devices that an administrator has configured to fulfill them.

DPM-enabled system

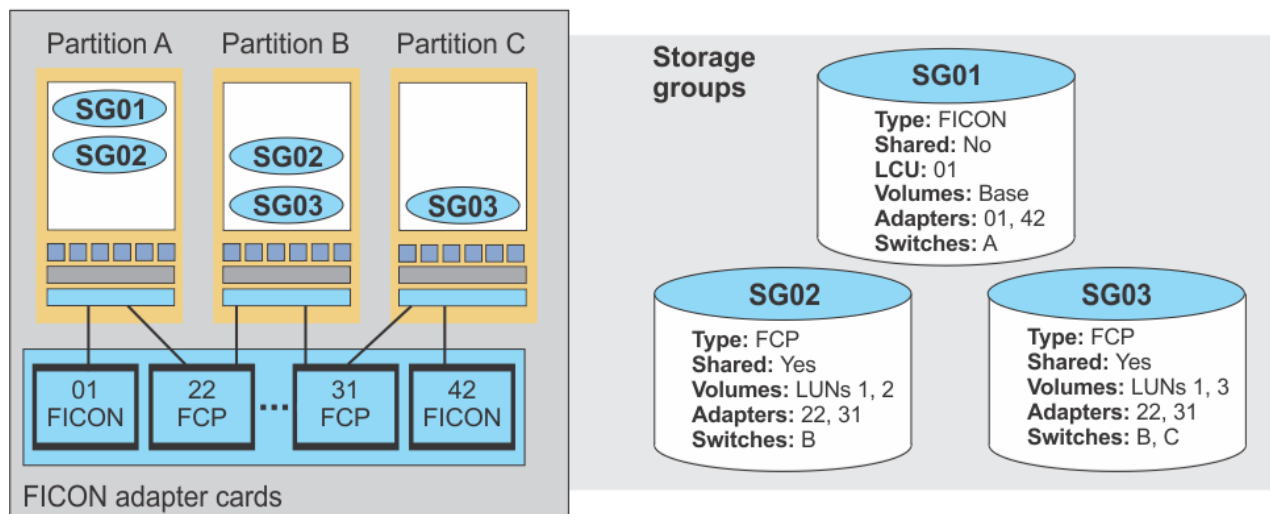


Figure 49. Sample FICON and FCP storage groups that are attached to partitions on the system

In [Figure 49 on page 123](#), 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.

Access to FCP tape storage through tape links

Administrators use the **Configure Storage** task to configure access to Fibre Channel Protocol (FCP) Small Computer System Interface (SCSI) tape storage devices. After system adapters are configured as FCP mode, and the system is connected to physical storage hardware by cables, system administrators can enable partitions to access and use those resources by creating tape links.

A *tape link* defines the attributes of a connection that one or more partitions can use to access one FCP tape library in the SAN. These connection attributes include storage resources such as system adapters, world wide port names (WWPNs), and the number of partitions that can share the connection. This connection is ready for use after the storage administrator completes zoning tasks through tools for managing SAN switches. *Zoning* is the process of grouping two sets of WWPNs into one zone: the host WWPNs for the system and the target WWPNs for the tape drives in a tape library. This grouping enables communication between the system and the tape library, while preventing communication with other systems or SAN devices.

To create a tape link, select **REQUEST TAPE LINK** or **CREATE TAPE LINK**. (The option that you can select depends on the authorization of your user ID.) Creating a tape link can be as easy as providing a name for the tape link, and checking the default value for the number of connecting paths. In this case, the storage administrator selects the tape library and the system adapters to use for the tape link.

Figure 50 on page 125 provides a graphical illustration of the physical and logical resources that are required for a tape link, and the process flow for requesting or creating one.

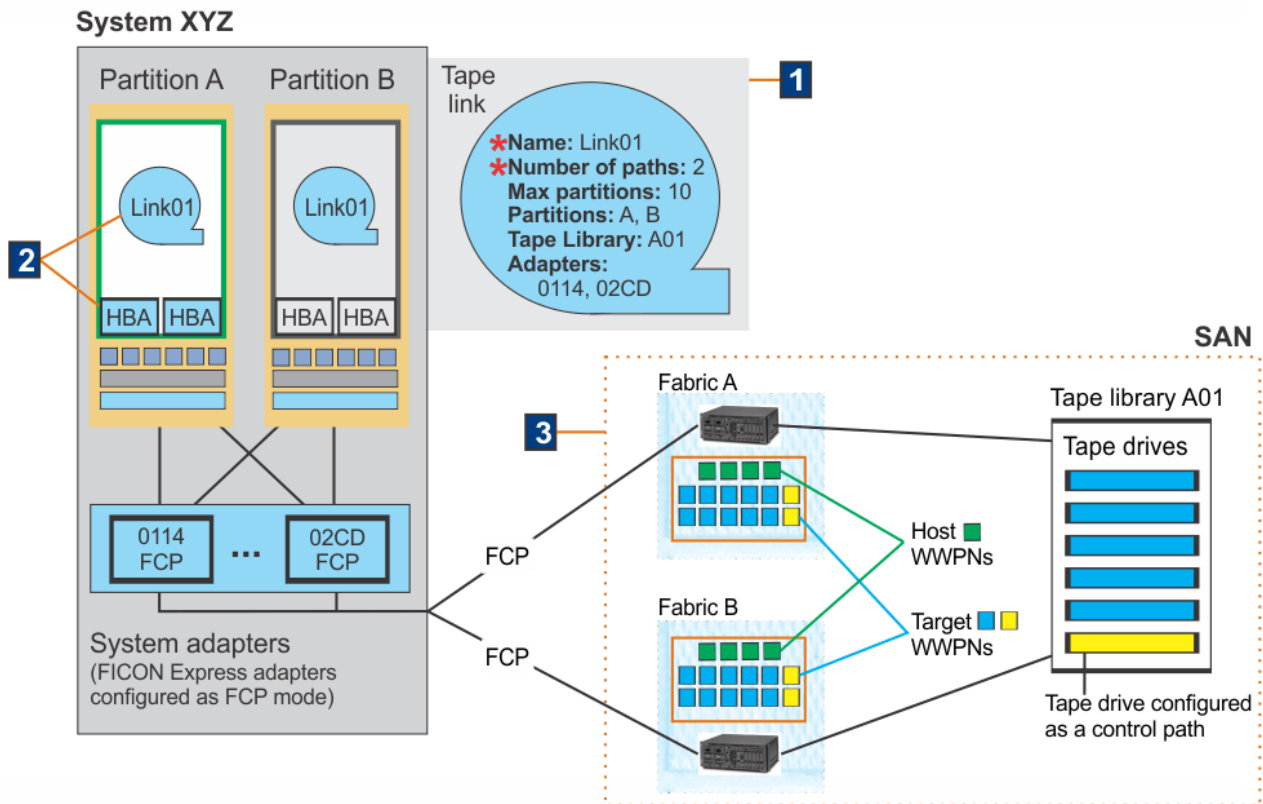


Figure 50. A tape link connects one or more partitions to one FCP tape library in the SAN

In Figure 50 on page 125:

1. You can request or create a tape link by providing a name for the tape link, and checking the default value for the number of connecting paths. You can change this value to increase the number of connecting paths (decreasing the value to 1 is possible, but not a suggested practice).

If you want additional control over the resources, you can optionally select specific partitions to which DPM attaches your tape link; set the maximum number of partitions that share the tape link; select a tape library; and select the system adapters.

If you select a tape library, DPM preselects adapters that are connected to the tape library already, and that provide optimal redundancy. Optimal redundancy for adapters is based on the following factors, which are listed in priority order from highest to lowest: on the location in the I/O drawers, on the drawer domain, on the current allocation, and on the connection to SAN fabrics. (If no adapters meet those criteria, DPM assigns placeholders for the storage administrator to select.)

After you confirm your request:

- DPM automatically generates the WWPNS that storage administrators use to fulfill the tape link request.
 - DPM also generates zoning instructions for the storage administrators who receive the request.
2. After you send your request, DPM creates the tape link, sets its fulfillment state to Pending, and starts to run automated background checks every 10 minutes to determine when the required storage resources for this tape link are available. In these background checks, DPM tries to detect the tape drives (logical units or LUNs) for the WWPNS that are assigned to this tape link.

Also, if you selected any specific partitions as part of your request, DPM asynchronously attaches the tape link to those partitions, automatically generating the virtual host bus adapters (HBAs) that the partitions need for access to the tape library. This attachment process might take some time, depending on the status and number of the selected partitions. The attachment to stopped partitions is relatively quick, but attachment to active partitions can take longer because these partitions might be busy.

3. To complete the connection from the system to the tape library, the storage administrator uses the generated instructions to configure the SAN devices. To properly configure a tape link connection to a tape library, the storage administrator:

- Selects the tape library and the number of tape drives that are accessible through the tape link.
- Ensures that one tape drive in the library is configured as a control path.
- Zones the host and target WWPNNs in each fabric.

Depending on the zoning and the status of assigned resources, the fulfillment state for the tape link can change to Complete, Incomplete, or Pending with mismatches. Only when the fulfillment state of the tape link changes to Complete, DPM performs dynamic I/O operations to complete the connection to the tape library for started partitions, such as Partition A. For stopped partitions, such as Partition B, DPM performs dynamic I/O operations only when the stopped partition is started.

More information about fulfillment states is available through the **Tape Link details** page and its associated online help; this information includes possible user actions to resolve potential problems. To access the **Tape Link details** page, go to the **Storage Overview** page, scroll to the Tape Links table, and select the table row of the tape link that you want to review.

Access to NVMe storage through storage groups

Administrators use the **Configure Storage** task to view any installed Non-Volatile Memory Express (NVMe) storage adapters, which are available on IBM LinuxONE (LinuxONE) systems only. Each NVMe adapter consists of two pieces of hardware: an IBM-supplied carrier card installed in a system I/O drawer, and the solid state drive (SSD) that customers purchase. IBM service representatives install the NVMe SSDs in the carrier cards after the system is delivered to the customer site.

Through **Connect Storage** and **Storage Cards** in the **Configure Storage** task, you can view NVMe adapters in the visual display of system frames and I/O drawers, and view specific details about each adapter; for example, you can tell whether the carrier card has an SSD installed. You cannot use this task to reconfigure NVMe storage adapters; reconfiguration requires properly removing the carrier card and its SSD from the drawer and reinstalling them in a different physical location, as instructed by a service representative. However, you can download an exportable file of the NVMe adapter plan, which lists details about the installed adapters. The file is in Comma Separated Values (CSV) format that you can view in a spreadsheet application. For a sample file, see [Appendix C, “Sample NVMe adapter plan,” on page 167](#).

System administrators can enable partitions to access and use NVMe storage by creating one or more NVMe storage groups, and attaching them to the partitions.

- To create an NVMe 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.)
 - You can create a storage group containing one or more available NVMe SSDs that are installed in carrier cards. You can define these SSDs as either boot or data volumes. You can create only one NVMe storage group at a time.
 - Because NVMe storage groups do not require the involvement of a storage administrator, DPM automatically creates the storage group, which is immediately ready for use.
 - Note that you cannot use a template to define an NVMe storage group.
- To enable partitions to access NVMe storage, use the **New Partition** task or **Partition Details** task to attach available NVMe storage groups. You can attach one or more NVMe storage groups, and select an SSD volume in an attached NVMe storage group as the boot volume for an operating system or hypervisor. If you plan to boot an operating system or hypervisor from a volume in an NVMe storage group, check the product documentation for the operating system or hypervisor to make sure that it supports booting from NVMe storage.

Only one partition can use an NVMe storage group at any given time; an NVMe storage group cannot be shared. A partition that has attached NVMe storage groups can also have attached FICON and FCP storage groups, and FCP tape links.

Storage group templates


To make storage requests even easier to complete, administrators can create a template for requesting 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. (Note that you cannot use a template to define an NVMe storage group.)

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 FICON and FCP adapters

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 FICON Express 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.

When you use the **Configure Storage** task to create a FICON or FCP storage group, DPM automatically assigns configured FICON or FCP adapters to the storage group. In prior DPM releases, you had to directly assign adapters through the **New Partition** or **Partition Details** task. In contrast, with DPM R3.1 or later releases, adapters are assigned through the storage group, when you attach it to a partition. For FCP storage groups only, you have the opportunity to change the assigned adapters when you attach an FCP storage group to a partition. FCP adapter reassignment is available through the **Partitions** tab in Storage Group details in the **Configure Storage** task, and the **Storage** section of the **New Partition** and **Partition Details** tasks.

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 and details

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, and the Tape links table, which contains one row for each tape link. To view more details about a specific storage group or tape link, click anywhere in a table row to open the Storage Group details page or Tape Link details page. The fulfillment state indicates whether the storage group or tape link is available for use.

If no storage groups or tape links exist, the table includes a selectable tile through which you can request or create a new storage group or tape link.

Depending on the permissions that are associated with your user ID, you can select actions (such as modify or delete) for a specific storage group or tape link. The connection report for a selected FCP or FICON storage group can be useful for diagnosing configuration problems. Similarly, the **ENVIRONMENT** tab display on the Tape Link details page can also pinpoint zoning problems and other connection errors.

Chapter 19. Setting up storage resources for the first time

When your company orders a 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. 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 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” on page 129](#)
- [“Connecting a system to external storage devices in the SAN” on page 130](#)
- [“Requesting FICON or FCP storage for partitions to use” on page 134](#)
- [“Requesting NVMe storage for partitions to use” on page 137](#)
- [“Requesting tape storage for partitions to use” on page 139](#)

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 external storage devices in the SAN

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-enabled 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 in the storage area network (SAN) 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.

For connections to tape storage in the SAN, use the **CONNECT TO STORAGE** wizard only to define system adapters as FCP adapters. Do not define FCP tape libraries as storage subsystems in the FICON Connections section of the wizard. Instead, create a tape link, which defines the attributes of a connection that one or more partitions can use to access one FCP tape library in the SAN. For more information, see [“Requesting tape storage for partitions to use” on page 139](#).

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 129](#).
 - 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 129](#).

- If the system has one or more IBM Adapter for NVMe1.1 features, you cannot configure those adapters through this task, but you can view them: DPM detects any NVMe carrier card and its solid state drive (SSD) and displays them in the frames and drawers where they are installed. From the **Configure Storage Cards** page, you also can download an exportable file of the NVMe adapter plan, which lists details about the installed adapters. The file is in Comma Separated Values (CSV) format that you can view in a spreadsheet application.

An overview of the initial FCP and FICON 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 FCP and FICON 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 FICON Express adapters 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 51 on page 131 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 other types of adapter cards (for example, network or crypto adapter cards). Red dots indicate the most recent changes to storage adapter cards in the drawer.

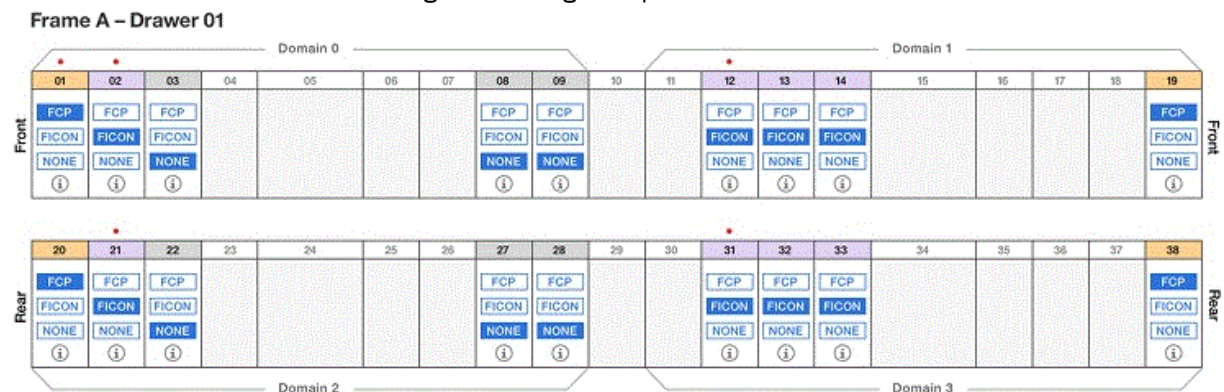


Figure 51. Sample frame and card display on the **Configure Storage Cards** page

Depending on the configuration of the system, the I/O drawers with storage adapter cards can be spread across multiple frames. For multiple-frame systems, the **Configure Storage Cards** page provides an interactive overview map through which you can easily navigate between frames to view the I/O drawers and cards installed in the system. Red dots in the overview map indicate recent changes to a specific frame that contains an I/O drawer with storage adapter cards.

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

The second task is building a visual copy of the storage hardware devices in the SAN and their FICON connections to this 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. (**Reminder:** in this step, do not define any FCP tape libraries as storage subsystems.)

Figure 52 on page 132 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.

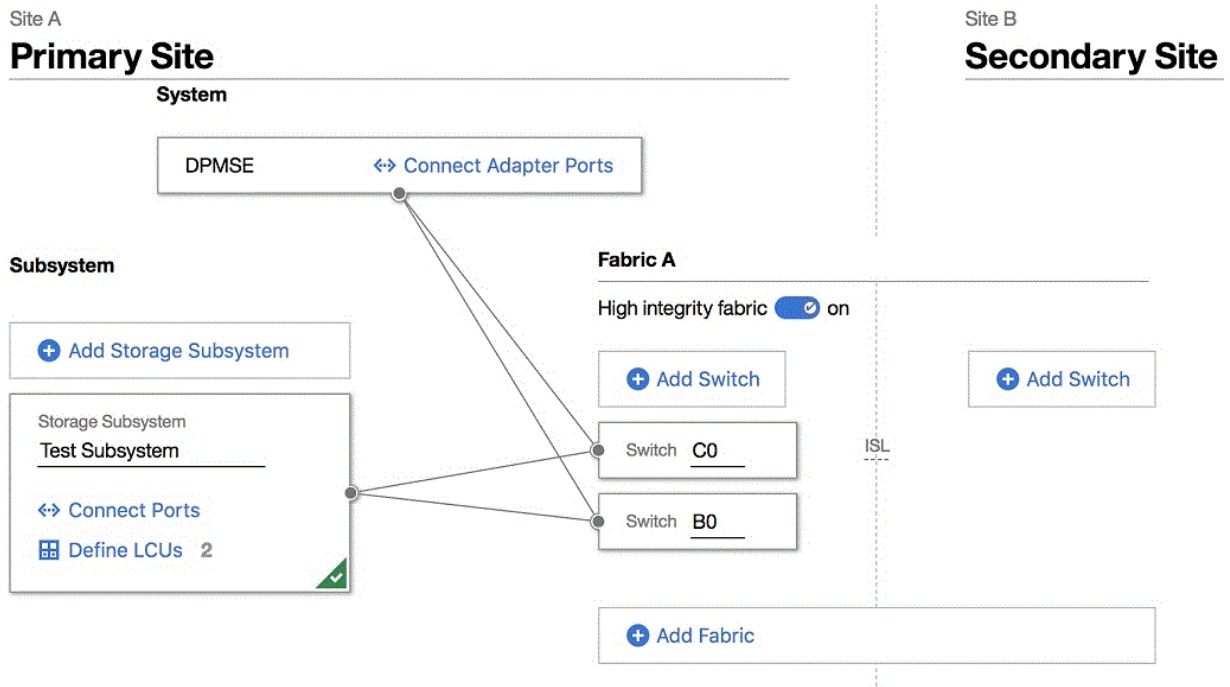


Figure 52. Sample display of two sites with physical storage hardware configured for the primary site

System administrators can define physical elements, such as switches, fabrics, and disk 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 FICON Express adapters 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 alphanumerics, 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.
- An exportable file of an FCP and FICON adapter cabling plan that you can use to physically connect the system to SAN hardware. The file is in Comma Separated Values (CSV) format that you can view in a spreadsheet application. For a sample spreadsheet, see [Appendix B, “Sample System Cabling Plan,” on page 165](#).

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 FICON or FCP storage for partitions to use” on page 134](#).

- 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 156.

Requesting FICON or FCP storage for partitions to use

Review this topic to prepare for using the **Request Storage** task to create FICON or FCP storage groups for use by partitions on a DPM-enabled system. You can access the **Request Storage** task by selecting **REQUEST STORAGE GROUP** 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 knowledge of mainframes or Linux systems; however, a storage administrator must fulfill any requests for storage. 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

- If you want to request Non-Volatile Memory Express (NVMe) storage, see the topic [“Requesting NVMe storage for partitions to use”](#) on page 137.
- If you want to request access to FCP tape storage, see [“Requesting tape storage for partitions to use”](#) on page 139.
- 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 129.
 - 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 129.

An overview of the Request Storage process for FCP and FICON storage groups

To request storage for partitions to use, you define one or more FCP or FICON storage groups. A *storage group* is a logical group of storage volumes that share certain attributes. FICON and FCP storage groups can be shared by multiple partitions, and multiple storage groups can be attached to one partition.

You can define FCP or FICON 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, even if no storage adapter cards of that type are configured for the system.

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.

- **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:** Select this attribute when you want to enable the direct assignment of host bus adapters (HBAs) so an operating system or its guests can access an FCP 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.

Although the controls in the **Configure Storage** task allow you to select this attribute only for a dedicated (not shared) FCP storage group, you can optimize 2nd level virtualization for separate partitions so they can share the same storage disks. For more information, see [“Optimize 2nd level virtualization and share the same FCP disks across partitions”](#) on page 136.

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.

Optimize 2nd level virtualization and share the same FCP disks across partitions

When you request storage for partitions to use, you can select the **Optimized for 2nd level virtualization** attribute when you want to enable the direct assignment of host bus adapters (HBAs) so an operating system or its guests can access an FCP storage group. Although the controls in the **Configure Storage** task allow you to select this attribute only for a dedicated (not shared) FCP storage group, you can optimize 2nd level virtualization for separate partitions so they can share the same storage disks.

For example, suppose that you have created two partitions, ZVM1 and ZVM2, in which you plan to install z/VM to host multiple Linux images, and you want both of these partitions to share the same set of 10 storage volumes. To accomplish this goal, you need to request two different FCP storage groups, and request that your storage administrator configure the same logical unit numbers (LUNs) in the storage controller for the two storage groups. To do so, perform the following steps.

1. Open the **Configure Storage** task and, depending on your user ID authorization, select either **REQUEST STORAGE** or **CREATE STORAGE GROUP**.
2. On the **Specify Storage Attributes** page, define the attributes for the storage group.
 - For **Type**, select FCP.
 - For **Shareability**, select Dedicated.
 - For **Connectivity**, specify the number of paths to be available for use by each operating system with access to this storage group.

- Select **Optimized for 2nd level virtualization**, and specify the number of additional connections (HBAs) that can be assigned directly to the operating system or its guests.

Select **NEXT** to continue.

3. On the **Add Storage Volumes** page, specify the capacity and type of volumes that you require. For this example, define 10 storage volumes. When you have finished defining the volumes, select **NEXT** to continue.
4. On the **Name and Duplicate** page:
 - a. Provide a name for the first storage group; for example, FCP-STORAGE-1.
 - b. Enter 1 as the number of duplicates and select **DUPLICATE**.
 - c. Provide a unique name for the second storage group; for example, FCP-STORAGE-2.
 - d. Select **NEXT** to continue.
5. On the **Confirm** page, review the summary of your storage request; then select **NEXT** to continue.
6. Review the automatically generated storage request, and add instructions for the storage administrator to configure the same logical unit numbers (LUNs) in the storage controller for both storage groups. Then send it to one or more storage administrators for fulfillment.

For FCP storage groups, 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.
7. Attach the storage groups and start the partitions. Note that you can perform these tasks before the storage groups are fulfilled.
 - a. Open the **Partition Details** task for partition ZVM1, and go to the **Storage** section to attach FCP storage group FCP-STORAGE-1. Save your changes and close the task.
 - b. Open the **Partition Details** task for partition ZVM2, and go to the **Storage** section to attach FCP storage group FCP-STORAGE-2. Save your changes and close the task.
8. Use the **Start** task to start the ZVM1 and ZVM2 partitions.

Requesting NVMe storage for partitions to use

Review this topic to prepare for using the **Request Storage** task to create a Non-Volatile Memory Express (NVMe) storage group for use by a partition on a DPM-enabled system. NVMe storage is available only when the system has one or more IBM Adapter for NVMe1.1 features. Depending on the authorization of your user ID, you can access the **Request Storage** task by selecting **REQUEST STORAGE** or **CREATE STORAGE GROUP** in the **Configure Storage** task. 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 an NVMe storage group

- If you want to request FICON or FCP disk storage, see the topic [“Requesting FICON or FCP storage for partitions to use”](#) on page 134.
- If you want to request access to FCP tape storage, see [“Requesting tape storage for partitions to use”](#) on page 139.
- To request a storage group through the **Configure Storage** task, you can use 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.
- If you plan to boot an operating system or hypervisor from a volume in an NVMe storage group, check the product documentation for the operating system or hypervisor to make sure that it supports booting from NVMe storage.

- If you plan to boot an operating system, hypervisor, or dump program from an NVMe SSD volume, note that NVMe namespace management is not supported, so you can boot programs only from namespace ID=1.

An overview of the procedure for creating an NVMe storage group

In contrast to FICON and FCP adapters that provide access to external storage devices in the storage area network (SAN), NVMe storage adapters provide high-speed storage within a system. Each NVMe adapter consists of two pieces of hardware: an IBM-supplied carrier card installed in a system I/O drawer, and the solid state drive (SSD) that customers purchase. IBM service representatives install the NVMe SSDs in the carrier cards after the system is delivered to the customer site.

To request NVMe storage for partitions to use, you define one or more storage groups that contain one or more SSD volumes. The following steps describe the process of creating an NVMe storage group.

Step 1. Specify the SSD volumes to add to the storage group

On the **Add Storage Volumes** page, review the entries in the Available Volumes table to determine which NVMe SSDs to add to the storage group. For each volume, the table includes the following information:

CAPACITY

Indicates the size of the SSD volume in gibibytes (GiB).

SERIAL NUMBER

Specifies the serial number of the SSD volume.

LOCATION

Specifies the physical location of the NVMe carrier card in the I/O drawer of the system.

Step 2. Optionally, change the attributes of each SSD volume

In the Available Volumes table, you can change the type, specify a device number, or enter a description of each volume that you want to add to the storage group.

- For Type, select either **Data** or **Boot**. Select **Boot** only if this volume is to contain bootable programs, such as the image of the operating system to be installed in a partition. You can specify only one type for each volume, but you can define more than one volume of each type for the storage group.
- Enter a unique, four-digit hexadecimal device number in the range 0001 - ffff for this volume; otherwise, DPM automatically assigns a device number when the storage group is first attached to a partition. To avoid conflicts, the suggested practice is to have DPM automatically assign device numbers.

Step 3. Add the SSD volume to the storage group

Select **ADD** in the table row of each SSD that you want to add to the NVMe storage group. After you select an SSD, the SSD entry moves to the Volumes to be added table.

Step 4. Name the storage group

On the **Name and describe** page, specify the name of your new storage group. Optionally, provide a description. For the name of the storage group, specify a value that is 1 - 64 characters in length. Supported characters are alphanumerics, 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.

Step 5. Create the storage group

On the **Confirm and create** page, review the summary of your storage group. When you are satisfied with the details on the **Confirm and create** page, select **CREATE** to create the NVMe storage group.

Results of using the Request Storage task for an NVMe storage group

DPM creates the NVMe storage group and changes the screen display to the **Storage Overview**, where you can view the new storage group in the list. Note that the fulfillment state is Complete so the new storage group is ready for use.

What to do next

- 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. Only one partition can use an NVMe storage group at any given time; an NVMe storage group cannot be shared.
- You can view and modify the details of the NVMe storage group and its volumes through the **Storage Group details** page in **Storage Overview**. Note that the modifiable attributes (type, device number, and description) specified for any NVMe SSD volume persist, even after you delete the volume from the storage group, or delete the storage group itself.

Requesting tape storage for partitions to use

Review this topic to prepare for creating a tape link for use by partitions on a DPM-enabled system. Depending on the authorization of your user ID, select **REQUEST TAPE LINK** or **CREATE TAPE LINK** to provide access for partitions to one FCP tape library in the storage area network (SAN). 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 a tape link

- If you want to request FICON or FCP disk storage, see the topic [“Requesting FICON or FCP storage for partitions to use”](#) on page 134.
- If you want to request Non-Volatile Memory Express (NVMe) storage, see the topic [“Requesting NVMe storage for partitions to use”](#) on page 137.
- For information about using this task to manage tape libraries, see [“Viewing and managing FCP tape libraries”](#) on page 158.
- Before you can successfully request or create a tape link, physical storage hardware (subsystems, switches, and so on) must be connected by cables, and storage cards must be configured. If the **REQUEST TAPE LINK** or **CREATE TAPE LINK** option is disabled, none of the system adapter cards are configured as FCP. Select **STORAGE CARDS** to view and configure the available system adapters.
- 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 129.
 - For information about setting up SMTP, see the online help for the **Monitor System Events** task.
- For **REQUEST TAPE LINK**, you can use the default SYSPROG or SERVICE user ID, or any user IDs that an access administrator has authorized to the **Configure Storage** task through customization controls in the **User Management** task.
- For **CREATE TAPE LINK**, you need a user ID with the same permissions as both the default SYSPROG and STORAGEADMIN user IDs.

An overview of the procedure for creating a tape link

A *tape link* defines the attributes of a connection that one or more partitions can use to access one FCP tape library in the SAN. These connection attributes include storage resources such as system adapters, world wide port names (WWPNs), and the number of partitions that can share the connection. This connection is ready for use after the storage administrator completes zoning tasks through tools for managing SAN switches.

Creating a tape link can be as easy as providing a name for the tape link, and checking the default value for the number of connecting paths. In this case, the storage administrator selects the tape library and the system adapters to use for the tape link. The following steps describe how to create a tape link that has more attributes than the name and number of connecting paths.

Step 1. Specify the tape link name and optional description

On the **Name and partitions** page, specify the name of your new tape link and an optional description.

Tape link name

For the name of the tape link, 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 distinguish this tape link from all other tape links that are defined for this system.

Description (optional)

For the optional description, use up to the maximum of 200 characters.

Step 2. Optional: select partitions and set the maximum number of partitions

On the **Name and partitions** page, select partitions to which you want the tape link attached, or define the maximum number of partitions that can attach the tape link, or both.

- To identify the partitions to which you want DPM to attach the tape link:
 1. Select **SELECT PARTITIONS** to open a dialog through which you can view details about available partitions on the system: the partition name; its current status (Active, Stopped, and so on); and the user-supplied description, if any.
 2. Use the check box to select each partition, then select **ADD** to close the dialog and populate the partitions table on the **Name and partitions** page.
- To set the maximum number of partitions to which the tape link can be attached, use the spinner or type the number that you want to set as the maximum. If you select partitions to attach the tape link, this number adjusts to match the number of selected partitions.

The maximum value can exceed the number of selected partitions, if any; however, the maximum value cannot exceed the system limit for concurrently active partitions. The text for this field indicates the limit for concurrently active partitions on your system.

Step 3. Optional: select a tape library

On the **Library and paths** page, select a tape library. To select a tape library, select an entry for a specific library in the Tape library list. If no tape libraries are shown, or you do not select a specific tape library, the storage administrator selects a tape library when fulfilling your request for the tape link.

Step 4. Check the number of connecting paths for the tape link

On the **Library and paths** page, decide whether to change the default number of connecting paths from the system to the tape library in the SAN. This connectivity setting has an impact on bandwidth, performance, and redundancy.

- The default value is 2, which you can change by using the "Number of connecting paths" slider, text entry field, or spin button.
- The maximum value is set to the number of system adapters that are currently configured as FCP adapters.

You can also select specific system adapters to use for this tape link.

- When you select a tape library, DPM preselects adapters that are connected to the tape library already, and that provide optimal redundancy. Optimal redundancy for adapters is based on the following factors, which are listed in priority order from highest to lowest: on the location in the I/O drawers, on the drawer domain, on the current allocation, and on the connection to SAN fabrics. (If no adapters meet those criteria, DPM assigns placeholders for the storage administrator to select.)
- You can use **Advanced path settings** to exchange any assigned adapters with other available adapters, to select an available adapter to fill in a placeholder row, or to replace an assigned adapter with a placeholder row. If you decide to select adapters yourself, note that the number of system adapters must match the number of connecting paths. For the best results, try to assign adapters that reside in different system I/O drawers and in different domains; that are not on the same card; and that are connected to different fabrics. Inline messages provide warnings if the selected adapters do not meet these criteria.

Step 5. Send the storage group request to a storage administrator for fulfillment

Depending on the authorization of your user ID, DPM automatically generates email that you can either send to a storage administrator for fulfillment, or use as a reminder of the zoning tasks that you need to complete for this tape link. Depending on the SMTP configuration, you can either make use of the integrated email and notifications, or download the generated request to send it manually.

Results of creating a tape link

DPM creates the FCP tape link, changes the screen display to the **STORAGE OVERVIEW** page, and displays a message confirming that the requested tape link was created. In the Tape Links table, the fulfillment state for the tape link remains as Pending until the following conditions are met:

- The storage administrator completes the configuration through tools for managing SAN switches. Note that, depending on the zoning and the status of assigned resources, the fulfillment state for the tape link can change to Complete, Incomplete, or Pending with mismatches.
- DPM asynchronously attaches the tape link to the partitions, if any, that you selected as part of your tape link request. This attachment process might take some time, depending on the status and number of the selected partitions. The attachment to stopped partitions is relatively quick, but attachment to active partitions can take longer because these partitions might be busy.

Only when the fulfillment state of the tape link changes to Complete, DPM performs dynamic I/O operations to complete the connection to the tape library for started partitions. For any stopped partition, DPM performs dynamic I/O operations only when the stopped partition is started.

What to do next

- To view more details about a specific tape link, click anywhere in a table row to open the Tape Link details page.
- To view, discover, or remove tape libraries from the DPM environment, select **REQUEST TAPE LINK** or **CREATE TAPE LINK**. For more information, see [“Viewing and managing FCP tape libraries” on page 158](#).
- If you did not select any specific partitions as part of your tape link request, use the **Partition Details** task to attach one or more tape links to an existing partition, or attach them when you create a new partition through the **New Partition** task. Through those tasks, you can view but cannot attach tape links that are in a pending fulfillment state (Pending or Pending with mismatches). For more information, see the online help for the appropriate task.

Chapter 20. Managing storage resources after the initial configuration is complete

After the initial storage configuration for the system is complete, and storage groups and tape links 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 or tape links. You can also use either the **New Partition** task or **Partition Details** task to attach storage groups or tape links to a partition, or use the **Configure Storage** task to attach tape links to a partition.

For overviews of specific storage management tasks, see the following topics.

- [“Viewing and managing storage groups and tape links” on page 143](#)
- [“Attaching storage resources to a partition” on page 152](#)
- [“Modifying the storage configuration of the system” on page 156](#)
- [“Viewing and managing FCP tape libraries” on page 158](#)

Viewing and managing storage groups and tape links

Review this topic to become familiar with the content of the **Storage Overview** and **Storage Group details** and **Tape Link details** pages, through which administrators can view and modify storage groups and tape links. More detailed information is available through the online help for the **Configure Storage** task.

At-a-glance view of storage: the Storage Overview page

Use the **Storage Overview** to view information about all of the storage groups and tape links that are defined for a DPM-enabled system.

Depending on the permissions that are associated with your user ID, you can select actions (such as modify or delete) for a specific storage group or tape link. 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.

Storage Groups table

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 or NVMe.

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 to open the Connection Report, 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.

Tape Links table

The Tape Links table contains the following information about each tape link.

NAME

Specifies the user-defined name of the tape link.

TYPE

Specifies the type of tape link: FCP.

PARTITIONS

Specifies the number of partitions to which the tape link is attached.

SHAREABLE

Specifies whether the tape link can be shared among partitions, or whether it is dedicated to only one partition.

TAPE LIBRARY

Specifies the serial number of the tape library that partitions can access through this tape link. If the value is "Not specified", the storage administrator has not yet selected a tape library for this tape link.

DESCRIPTION

Specifies the user-provided description, if any, of this tape link. The description can be up to 200 characters in length.

FULFILLMENT STATE

Identifies the current state of the tape link. DPM runs a background check of storage resources for FCP tape links and, if necessary, changes the fulfillment state. These checks are more frequent (every 10 minutes) for fulfillment states other than Complete (every 24 hours).

ACTIONS

In any table row, select the ellipsis (***) to display a selectable list of actions that you can take for the tape link. The listed actions vary, depending on the permissions that are associated with your user ID.

The Storage Group details page

Use the Storage Group details page to view or modify information about a specific 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 on the same system, or migrated to another system that has DPM R3.1 or a later DPM version applied. This fulfillment state indicates that DPM is checking the logical and physical elements that support a 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. For DPM to recheck the storage group and change the fulfillment state, 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.

For more details about the migration process, see the appropriate topic in [Chapter 11, “Migrating DPM configuration data to a new system,”](#) on page 87.

Complete

The storage group is ready for use.

Incomplete

One or more volumes or adapters that are used for a storage group are marked as incomplete. DPM periodically checks the availability of storage volumes or adapters for 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 (❗).

Pending

A system administrator has sent a request to create or modify a FICON or FCP 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

For an FCP storage group, a system administrator sent a request to create or modify that storage group, and the storage administrator fulfilled that request, but with an amount of storage that does not exactly match the original request. For an NVMe storage group, as part of a repair, one or more NVMe SSDs were replaced with SSDs of a different size. 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 Group details page.

VOLUMES

The **VOLUMES** tab display varies, depending on the type of storage group. For an FCP or NVMe 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 table also includes a link through which you can view information that you need to enable a partition's operating system or installer to access a specific volume.

The **VOLUMES** tab display also varies depending on the current fulfillment state of the storage group. [“The Storage Group details page”](#) on page 144 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 9. Effect of fulfillment status on the Volume tab display

| Fulfillment state | Volume tab display | Possible action |
|---|---|--|
| Checking migration <p>This fulfillment state indicates that DPM is checking the logical and physical elements that support a storage group it created during a system migration or firmware upgrade process.</p> | The display contains an empty Volumes table. | <p>When DPM completes the check, it changes the fulfillment state to Complete or Pending with mismatches.</p> <p>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. For DPM to recheck the storage group and change the fulfillment state, 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 (↻).</p> |
| Complete <p>This fulfillment state indicates that DPM has successfully detected all of the logical and physical elements that support the volumes in this storage group.</p> <p>Note that DPM can mark a FICON storage group as Complete even if some requested alias devices were not included because of device number conflicts with base volumes.</p> | The display contains a table entry with complete information for all volumes in this storage group. | <p>None.</p> <p>If alias volumes are excluded from a FICON storage group, scroll to view the Alias Volumes table; if some alias volumes were not included, an inline message prior to the table indicates the number of alias volumes that were excluded, and provides a link through which you can open a new window to view and resolve specific device number conflicts.</p> |
| Incomplete <p>This fulfillment state indicates that one or more volumes are incomplete. Volumes can be marked as incomplete under the following conditions.</p> <ul style="list-style-type: none"> • When DPM can no longer detect 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 or a later release. • When DPM detects that an NVMe adapter has become degraded or the NVMe SSD was incorrectly removed from its carrier card. | An incomplete icon (❗) is displayed in the table entry of each incomplete volume. | <p>To diagnose the problem for an FCP or FICON storage group, select the Connection Report action icon (🔗).</p> <p>When an NVMe storage group is marked as Incomplete, go to the HISTORY tab and check entries in the Actions table to find more details about the specific error.</p> <ul style="list-style-type: none"> • You can remove the volume that is marked with the incomplete icon, but you do not need to take any action as long as the NVMe storage group contains other volumes that are usable. The Incomplete fulfillment status and incomplete icons continue to be displayed for the SSD until a repair is completed under the direction of a service representative. • If you do not remove the volume and the NVMe SSD in error is later reinstalled or replaced, DPM detects the repair, and either changes the fulfillment state to Complete or to Pending with mismatches. Note that, if the SSD in error is replaced by a different SSD, DPM automatically changes the serial number of the volume. |

Table 9. Effect of fulfillment status on the Volume tab display (continued)

| Fulfillment state | Volume tab display | Possible action |
|--|---|---|
| Pending (creation request) When a creation request is in progress, this fulfillment state indicates that 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. |
| Pending (modification request) This fulfillment state indicates that 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) This fulfillment state indicates that 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 This fulfillment state indicates that DPM has detected changes to the number, size, or accessibility of volumes in the storage group. The following items describe the various conditions under which DPM assigns this fulfillment state to an FCP or NVMe storage group. <ol style="list-style-type: none"> 1. DPM detected more volumes for an FCP storage group than the number that was originally requested. 2. DPM detected the correct number of volumes in an FCP storage group, 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 an FCP storage group. 4. As part of a repair, an NVMe SSD was replaced by an SSD with a larger or smaller capacity. | 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 previous and current sizes are shown in the CAPACITY column, along with a message that explains the mismatch. | <ul style="list-style-type: none"> • The first two conditions for an FCP storage group 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 for an FCP storage group 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. • For a replacement NVMe SSD that is not the same size as the removed SSD, you can either keep the replacement SSD as a volume in the storage group, or modify the storage group to remove the replacement SSD or add different SSDs. |

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. Note that only one partition can use an NVMe storage group at any given time; an NVMe storage group cannot be shared.

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 ADAPTERS** to review the adapters assigned to a storage group 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. For FICON storage groups, volume device numbers must be in the range 0000 - ffff, and must be unique within a storage group and across all attached partitions and partition resources.

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.

The Tape Link details page

Use the Tape Link details page to view or modify information about a specific tape link on a DPM-enabled system. The Tape Link 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. You also can use this page to delete a tape link, or to resend a zoning request.

The actions that are displayed depend on the fulfillment state of the tape link and the permissions that are associated with your user ID. DPM runs a background check of storage resources for FCP tape links and, if necessary, changes the fulfillment state. These checks are more frequent (every 10 minutes) for fulfillment states other than Complete (every 24 hours). In these background checks, DPM tries to detect the tape drives (logical units or LUNs) for the WWPNs that are assigned to this tape link. Users can manually start a background check by selecting the **Update tape link environment** icon () on the **ENVIRONMENT** tab.

For specific fulfillment states, note that various summary fields, tabbed sections, environment elements, and table entries have a pending, incomplete, warning, or error icon to alert you to details that might need your attention or action.

The following list provides a summary of each fulfillment state. For more details about fulfillment states and possible actions that you can take to correct errors or mismatches, see [Table 10 on page 149](#).

Complete

All of the storage resources listed in a create or modify request are available, properly configured and zoned, and DPM detects only those resources.

Incomplete

One or more storage resources for the tape link are marked as incomplete because the resource is missing, or in an error or degraded condition. Because DPM periodically checks the availability of storage adapters, switches, and tape libraries that are in use for a tape link, resources that were functioning properly can become incomplete.

Pending

One or more requested storage resources are not yet available or zoned correctly, or the tape link is not yet attached to all partitions that were specified in the original create request or a modify request.

Pending with mismatches

DPM detects system adapters that do not match the original create request or a modify request. Either the number of system adapters does not match the number of connecting paths, or the detected adapters do not match specific adapters that were assigned to the tape link.

The tabbed sections provide more details related to the tape link. Note that the following icons indicate details that might require your attention or action: pending icons (🟡 or ⚠️), and incomplete icons (❗).

The following list describes the tabbed sections on the Tape Link details page.

ENVIRONMENT

The **ENVIRONMENT** tab displays the current zoning and other information about the system adapters, SAN fabrics and switches, the tape library and tape drives for an FCP tape link. You can update (or refresh) the display, which includes a 24-hour timestamp for the last update, in the format yyyy.mm.dd hh:mm:ss. You also can download a snapshot of the **ENVIRONMENT** tab contents.

| Table 10. Interpreting the ENVIRONMENT tab display for different tape link fulfillment states | | |
|--|--|-----------------|
| Fulfillment state | ENVIRONMENT tab display | Possible action |
| Complete All of the storage resources listed in a create or modify request are available, properly configured and zoned, and DPM detects only those resources. | The display contains complete icons (✅) next to the three section headings, and connections between the elements are shown as solid black lines. | None. |

Table 10. Interpreting the ENVIRONMENT tab display for different tape link fulfillment states (continued)

| Fulfillment state | ENVIRONMENT tab display | Possible action |
|--|---|---|
| <p>Incomplete</p> <p>The following conditions cause the fulfillment state for a tape link to be Incomplete.</p> <ul style="list-style-type: none"> • One or more system adapters are unreadable, in an error state, or no longer detected. • An error occurred, or an unrequested reconfiguration of storage resources in the SAN affects the requested resources. • The storage administrator has not configured a tape drive as the control path to the requested tape library. • The original tape link request included a specific library but a different library is detected, or more than one tape library is detected, through the current zoning. The tape link environment can include only one tape library. • The tape library can no longer be reached due to an error with cables or with the tape library itself. | <p>The display contains incomplete icons (❗) next to any section heading or individual elements (adapters, fabrics, or tape library) for which DPM detected an error. Also, connections between elements with errors are shown as solid red lines, with a slash (/) indicating where the connection is broken because of the error.</p> <p>If the errors are detected for system adapters or host WWPNs, an error icon is displayed next to the ADAPTERS or WWPNS tabs.</p> | <p>Hover your cursor over any incomplete icon to view more information about the specific error condition.</p> <ul style="list-style-type: none"> • If an adapter is in error, go to the ADAPTERS tab and select the adapter name to open the Adapter Details page in the Manage Adapters task. The adapter details can determine how you can resolve the error. The simplest fix is replacing the adapter in error with a different available adapter, which you can do by selecting Modify. • If the errors are due to zoning in the SAN fabrics, cable or optic errors, multiple tape libraries, or errors with the tape library itself, select Resend request to notify one or more storage administrators so they can take the appropriate action. |

Table 10. Interpreting the *ENVIRONMENT* tab display for different tape link fulfillment states (continued)

| Fulfillment state | ENVIRONMENT tab display | Possible action |
|---|--|---|
| <p>Pending</p> <p>One or more requested storage resources are not yet available or zoned correctly, or the tape link is not yet attached to all partitions that were specified in the original create request or a modify request.</p> <p>When the fulfillment state is Pending, DPM does not perform any dynamic I/O operations for partitions to which the tape link is attached.</p> | <p>The display contains pending icons (🟡) next to any section heading or individual elements (adapters, fabrics, or tape library) when connections are expected but not yet available because of pending zoning or configuration tasks.</p> <ul style="list-style-type: none"> If zoning tasks are pending for system adapters or host WWPNs, a pending icon is displayed next to the ADAPTERS or WWPNS tab. If the create or modify request included partitions to which the tape link is to be attached, and one or more asynchronous attachment processes are still in progress, a pending icon is displayed next to the PARTITIONS tab. | <p>The Pending fulfillment state is generally a temporary state, but some tasks require time to complete. To determine whether those tasks are progressing, you can periodically select the Update environment button to refresh the tab display, and note the changes, if any.</p> <ul style="list-style-type: none"> If zoning tasks are pending for a longer time than you expect, you can select Resend request to remind one or more storage administrators of the pending tasks. If tape link attachment operations are in progress, note that the attachment process might take some time, depending on the status and number of the selected partitions. The attachment to stopped partitions is relatively quick, but attachment to active partitions can take longer because these partitions might be busy. <p>If further investigation seems necessary, go to the PARTITIONS tab, view the partitions that are marked as pending, and select the partition name to open the Partition Details task.</p> |
| <p>Pending with mismatches</p> <p>DPM detects system adapters that do not match the original create request or a modify request. Either the number of system adapters does not match the number of connecting paths, or the detected adapters do not match specific adapters that were assigned to the tape link.</p> <p>When the fulfillment state is Pending with mismatches, DPM does not perform any dynamic I/O operations for partitions to which the tape link is attached.</p> | <p>The display contains pending with mismatches icons (🚨) next to the Systems section heading, on one or more adapters in this section, and on the ADAPTERS tab. Also, connections between the mismatched system adapters, the SAN fabrics, and the tape library are shown as solid black lines.</p> | <p>You must resolve all of the mismatched adapters before you can make any other modifications to the tape link. You can resolve mismatches by modifying the assigned adapters. To do so, select Resolve adapter mismatches or the Modify icon for the tape link to change the page display to Modification mode.</p> |

PARTITIONS

The **PARTITIONS** tab lists the partitions to which the tape link is attached in table format. If the tape link has not been attached yet to one of the listed partitions, the **PARTITIONS** tab has a pending icon (🟡), and another pending icon marks the table row for the specific partition.

To attach the tape link to more partitions, or to detach the tape link from one or more of the listed partitions, select the **Modify** icon for the tape link to change the page display to Modification mode. In this mode, you can *** modify_tl_partitions_intro_list ***

For each partition, the Partitions table includes the name, which is a hyperlink through which you can open the **Partition Details** task; operating status; and description (if any). 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.

ADAPTERS

The **ADAPTERS** tab lists the system adapters that are either specified through the original create request for the tape link, or connected to the tape library for this tape link. The adapter attributes include the name, ID, type, location, and current allocation.

- If all adapters are not yet assigned, the **ADAPTERS** tab has a pending icon (🕒) next to it, and the table entry for that adapter is marked with the same icon. The total at the foot of the Adapters table lists how many adapters are assigned to the tape link.
- If the adapters that are zoned do not match the adapters that were specified in the original create request, the **ADAPTERS** tab has a pending with mismatches icon (⚠️) and the tab display includes a table of Unrequested and zoned adapters, all of which are marked with the same pending with mismatches icon.

Note: You must resolve all of the mismatched adapters before you can make any other modifications to the tape link. You can resolve mismatches by modifying the assigned adapters. To do so, select the **Modify** icon for the tape link to change the page display to Modification mode.

- If an existing adapter becomes incomplete, the **ADAPTERS** tab has an incomplete icon (❗) next to it, and the table entry for that adapter is marked with the same icon. You can replace such adapters through Modification mode.

WWPNS

The **WWPNS** tab lists the host worldwide port numbers (WWPNs) that are available for use. The tab display contains two tables: one table that lists each host WWPN that is in use by a partition, and one table that lists unused host 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 tape link. 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. If the tape link 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 or tape links to a partition on a DPM system. To attach a storage group or tape link to a partition, use either the **New Partition** task or the **Partition Details** task. For tape links, you can also attach a tape link to one or more partitions through the **PARTITIONS** tab on the Tape Link details page; for more information, see the online help for the Tape Link details page in the **Configure Storage** task.

Before you attach a storage group or tape link to a partition

- To attach storage groups or tape links 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 37](#).
- 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 and a Tape Links table with controls that you can use to attach storage groups or tape links.

Step 1. Open the Attach Storage Groups or Attach Tape Links window

Select the plus icon in the table toolbar to open the **Attach Storage Groups** or **Attach Tape Links** 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.

- 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.

Step 2. Review and select one or more available tape links

On the **Attach Tape Links** window, select one or more tape links listed in the table to attach to this partition.

- The suggested practice is to select tape links that are in the Complete fulfillment state, but you can select any tape link 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 links in states other than Complete, some storage might not be available for use when you start the partition.
- Use the additional information in the table, as necessary, to decide which tape links to attach. This information includes not only the tape link name and fulfillment state, but also the number of partitions to which the storage group is already attached (if any).

Step 3. Check the table entries for the storage groups or tape links that you selected

Check the entries for the storage groups or tape links that you selected, which are now displayed in the Storage Groups table or Tape Links table in the Storage section. If necessary, you can use the minus icon in the table toolbar to remove a storage group or tape link from the table.

- For FICON storage groups only, you can change the volume device numbers only when the device number input field is active. The factors that determine whether you can change the device number include not only the current state of the partition, but also whether the storage group is shared or dedicated, and whether it is already attached to other partitions.
- For FCP storage groups and FCP tape links 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 or FCP tape link. An error icon is displayed if you try to specify a device number that is already in use.
- For FCP storage groups only, 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 4. Save the partition definition

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 FCP or FICON storage groups, 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 in the SAN.
- For FICON storage groups, DPM assigns device numbers, avoiding conflicts with any other device numbers for network, accelerator, and crypto adapters.
- For FCP tape links, DPM asynchronously attaches each tape link to the partition, and generates the HBAs that are required for the partition, and its operating system, to access the physical tape library in the SAN. This attachment process might take some time, depending on the status of the partition. The attachment to stopped partitions is relatively quick, but attachment to active partitions can take longer because these partitions might be busy.

Only when the fulfillment state of the tape link changes to Complete, DPM performs dynamic I/O operations to complete the connection to the tape library for started partitions. For any stopped partition, DPM performs dynamic I/O operations only when the stopped partition is started.

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 FICON and FCP 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). NVMe storage groups are automatically detected by the operating system, so you do not need to enter Linux commands to make that type of storage group available to the operating system. Similarly, the tape devices that are available through attached tape links are automatically detected by the operating system, so you do not need to enter Linux commands for tape devices either.

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. 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 for FCP or FICON storage groups depend on the type and fulfillment state, 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 system. Note that you cannot use this task to configure Non-Volatile Memory Express (NVMe) storage adapters; reconfiguration requires properly removing the carrier card and its SSD from the drawer and reinstalling them in a different physical location, as instructed by a service representative.

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 FCP or FICON storage adapter cards

The process of modifying installed storage adapters is very similar to Step 1 in [“An overview of the initial FCP and FICON storage configuration process”](#) on page 131.

- 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 adapter card that is already in use.
 - An FCP adapter card cannot be reconfigured when it is in use by a storage group that is attached to a partition, regardless of the partition state; when it is in use by a tape link; or when it is used to reach a tape library. For an FCP adapter that cannot be reconfigured, you can determine which storage groups are using the adapter by completing the following steps:
 1. Select the information icon to open a display that contains a link to open the Adapter Details page.
 2. Select the link and review the Connections section to determine which storage groups or tape links are using the adapter.
 - A FICON adapter cannot be reconfigured when it provides a logical control unit (LCU) path to a switch or storage subsystem. For a FICON adapter that cannot be reconfigured, you can determine which switch or storage subsystem is using that adapter by completing the following steps:
 1. Select the information icon to open a display that contains the adapter ID that you need, and make a note of it.
 2. Go to the FICON Connections page, and select the **Connect Adapter Ports** link in the System box.
 3. On the Connect System to Switches page or Connect System to Subsystem page, use the search to find the adapter ID in the frame display. Enter the adapter ID into the search field, and select the name of the matching adapter port in the drop-down list. The adapter port name in the search field is outlined with gray, which denotes that the adapter port is in use. When you hover your cursor over the adapter port name in the search field, the adapter port label in the frame display shows the ID of the switch or storage subsystem that is using the port. When you hover your cursor over the adapter port in the frame display, the adapter details include a message indicating that the adapter port is in use and cannot be disconnected.

FCP and FICON adapter cards that are in use are shown with disabled (gray) label selections for changing the current configuration.

- When you open the Storage Group details page for an FCP storage group, you can review the adapters assigned to a storage group 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 65.

Planning considerations for modifying FICON connections

The process of modifying the visual copy of the storage hardware environment is very similar to Step 2 in [“An overview of the initial FCP and FICON storage configuration process”](#) on page 131. (**Reminder:** Do not define any FCP tape libraries as storage subsystems.)

- 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.

Viewing and managing FCP tape libraries

Use **REQUEST TAPE LINK** or **CREATE TAPE LINK** to view information about the tape libraries that are connected to this system. You can also use this task to send zoning instructions to storage administrators to start initial zoning, to update zoning, or to remove FCP tape libraries from the DPM environment.

Before you begin

- For **REQUEST TAPE LINK**, you can use the default SYSPROG or SERVICE user ID, or any user IDs that an access administrator has authorized to the **Configure Storage** task through customization controls in the **User Management** task.
- For **CREATE TAPE LINK**, you need a user ID with the same permissions as both the default SYSPROG and STORAGEADMIN user IDs.

An overview of tape library actions

When you open the **Configure Storage** task and select either **REQUEST TAPE LINK** or **CREATE TAPE LINK**, the **Request tape link** or **Create tape link** page opens. The page display includes the FCP tape libraries table, which lists any tape libraries that DPM has discovered in the SAN. Above this table, the display includes either descriptive text or a time stamp indicating the last time, if any, that DPM discovered any correctly zoned FCP tape libraries. For discovered tape libraries, the table includes the tape library name (serial number), model, and state.

To display the table actions, select the ellipsis (•••) in the table header.

Start initial zoning

If the Tape libraries table is empty, select **Start initial zoning**, which causes DPM to generate an email with initial zoning instructions for a storage administrator. The instructions include the system management world wide port name (WWPN), which is a dedicated host WWPN that does not enable access to data in the tape library; its sole purpose is to enable DPM to discover (or detect) tape libraries in the SAN.

The instructions tell the storage administrator to complete the following zoning tasks in fabrics (switches) for *each* tape library to be connected to the system.

- Zone the system management WWPN with the target WWPN of at least one tape drive in the tape library.
- Configure the tape drive associated with the target WWPN as a control path.

When the initial zoning is complete, the time stamp and table content are updated.

Discover libraries

If the Tape libraries table has entries but does not list a specific tape library that you want find, or the time stamp is not recent, select **Discover libraries** to cause DPM to perform a one-time check of the current connections to tape libraries in the SAN. Depending on the number of system adapters that are defined as FCP storage adapters, and the number of target ports that are currently zoned, this discovery check could take some time.

Discover libraries is not available for selection until an administrator either requests initial zoning or requests a new tape link. When the **Discover libraries** check completes, the time stamp and table content are updated.

Available

At least one physical path reaches the tape library. (Note that you cannot delete a tape library that is in this state.)

Not available

No physical path, including the system management WWPN path, reaches the library. This state usually indicates a tape library that was correctly zoned and used for one or more tape links, but is no longer connected.

Remove libraries

To delete a tape library that you no longer want to use in your environment, use the following suggested procedure.

1. Delete all of the tape links that use the specific tape library. To delete tape links, you must use the default SYSPROG user ID or an equivalent user ID with the same permissions.
 - a. To quickly find a list of the tape links that you need to delete, go to **STORAGE OVERVIEW**, and sort the Tape links table rows by the value in the LIBRARY column.
 - b. For each tape link to be deleted, check the **Tape Link details** page to determine whether any active partitions have attached the tape link. If so, use the **Stop** task to stop the active partitions.
 - c. Select the **Delete** icon on the **Tape Link details** page to delete the tape link. The time required to successfully delete the tape link depends on the number of partitions that have attached the tape link.
 - d. Return to **STORAGE OVERVIEW**.
2. Select **REQUEST TAPE LINK** or **CREATE TAPE LINK** and use the **Remove libraries** action in the FCP tape libraries table to switch into table-edit mode. In this mode, a trash can icon is added to each table row.
3. Select the appropriate trash can icon to remove one or more tape libraries, and select **Save**. This action not only removes the table rows for the deleted tape libraries, but also results in an automatically generated email that contains instructions for a storage administrator to unzone the deleted libraries, including the system management WWPN.

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 12 on page 161](#) 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 12 on page 161](#).

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 11 on page 161](#) for the tasks mapping legend.

Table 11. 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 12. DPM task roles mapping

| DPM Tasks | Task Roles | | | | | |
|--|------------|----|----|-----|-----|-----------------------------------|
| | AA | SP | OP | AOP | SER | SA |
| Configure Storage | | O | | | O | O "1" on page 163 |
| - Create, modify, or delete a storage group or tape link | | X | | | | |
| - Attach a storage group or tape link to a partition | | X | | | | |
| - Detach a storage group or tape link from a partition | | X | | | | |
| - Change an HBA device number for an FCP storage group or an FCP tape link | | X | | | | |

Table 12. DPM task roles mapping (continued)

| DPM Tasks | Task Roles | | | | | |
|---|------------|----|----|-----|-----|----|
| | AA | SP | OP | AOP | SER | SA |
| 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 | |
| - 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 | |
| - Manage Secure Execution Keys or Import Secure Execution Keys* | | | | | X | |

Table footnote:

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. Sample System Cabling Plan

When you connect a system to storage by using the **Configure Storage** task, DPM creates an exportable file of an FCP and FICON adapter cabling plan that you can use to physically connect the system to SAN hardware. The file is in Comma Separated Values (CSV) format that you can view in a spreadsheet application.

The following sample shows the content of a CSV file. Note that the headings in the CSV file are all on a single line, instead of wrapping on two lines, as shown here.

```
System Name,Adapter ID,Adapter Port Location,Frame,Drawer,Domain,Slot,Port,Card Type,
Protocol,Optic Type,Switch ID
M01,0115,A15A-D113-J.00,A,A15A,1,D113,0,FICON Express 8s,FICON,short range,10
M01,0116,A15A-D113-J.00,A,A15A,1,D113,1,FICON Express 8s,FICON,short range,11
M01,0117,A15B-D114-J.00,A,A15A,1,D114,0,FICON Express 8s,FICON,long range,10
M01,0118,A15B-D114-J.00,A,A15A,1,D114,1,FICON Express 8s,FICON,long range,11
M01,0119,Z15B-LG13-J.00,Z,Z15B,2,LG13,0,FICON Express 16s,FCP,short range,-
M01,0120,Z15B-LG13-J.00,Z,Z15B,2,LG13,1,FICON Express 16s,FCP,short range,-
M01,0121,Z15B-LG14-J.00,Z,Z15B,2,LG13,0,FICON Express 16s,FCP,long range,-
M01,0122,Z15B-LG14-J.00,Z,Z15B,2,LG13,1,FICON Express 16s,FCP,long range,-
```

Appendix C. Sample NVMe adapter plan

If a DPM-enabled system has one or more IBM Adapter for NVMe1.1 features, you can view Non-Volatile Memory Express (NVMe) storage adapters through the **Configure Storage** task. You also can download an exportable file of the NVMe adapter plan, which lists details about the installed adapters. The file is in Comma Separated Values (CSV) format that you can view in a spreadsheet application.

Each NVMe adapter consists of two pieces of hardware: an IBM-supplied carrier card installed in a system I/O drawer, and the solid state drive (SSD) that customers purchase. The following sample shows the entries for an NVMe carrier card and an installed SSD in a CSV file.

```
System Name,Adapter ID,Adapter Location,Frame,Drawer,Domain,Slot,Card Type
P000T218,01B0,A31B-LG17,A,A31B,0,17,NVMe SSD
P000T218,01A0,A31B-LG14,A,A31B,1,14,NVMe carrier card
```


Appendix D. Configuring and accessing FCP tape drives on a DPM-enabled system

Note: The instructions in this appendix apply only to DPM R4.2 or earlier. With DPM R4.3, system or storage administrators can use a simplified interface to configure access to FCP tape libraries in the storage area network (SAN). DPM R4.3 is the preferred method of tape support. For more information, see the DPM R4.3 entries in [Table 5 on page 78](#).

Starting with R3.1, DPM introduced a simplified way of configuring and managing storage. The Hardware Management Console (HMC) **Configure Storage** task provides an end-to-end user experience for administrators who complete the following tasks:

- Connecting a DPM-enabled system to 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 in the storage area network (SAN).
- Creating storage groups, through which partitions on the system can access and use the connected storage devices. A *storage group* is a logical group of storage volumes that share certain attributes.
- Sending requests for storage groups to storage administrators who fulfill the requests through the storage subsystem management interface.

DPM detects but does not display FCP tape drives, so administrators cannot directly manage them through DPM tasks on the HMC. However, operating systems that are hosted by partitions on a DPM-enabled system can access and use those FCP tape drives. This topic provides instructions for configuring and accessing FCP tape drives on a system on which DPM R3.1 or a later release is installed. The set of instructions you use depends on the configuration of tape drives in your SAN:

- If the tape drives are attached to the system through one or more SAN fabrics that are also used for other types of storage devices (such as disks), see [“Configuring and accessing tape drives through a shared fabric” on page 169](#).
- If the tape drives are attached to the system through one or more SAN fabrics that are dedicated for tape drives only, see [“Configuring and accessing tape drives through a dedicated fabric” on page 172](#).

Configuring and accessing tape drives through a shared fabric

If the tape drives are accessible through one or more SAN fabrics that are also used for other types of storage devices, complete the steps in this procedure to correctly configure the tape drives for use. A system administrator can perform most of these steps, but step [“2” on page 171](#) requires the skills of a storage administrator to fulfill the storage request.

Before you begin

You need a user ID that has authorization to use the **Configure Storage** task and also has permission to access related objects: the CPC, storage groups, and adapters. The suggested practice is to use the **User Management** task to define a user based on the default SYSPROG user ID, and add the following storage roles:

- Storage Administrator Objects
- Storage Administrator Tasks

For instructions, see the topic *Authorizing users of the Configure Storage task*, in either the HMC online help for the **Configure Storage** task or in the *DPM Guide*.

About this task

If the tape drives that you want to use are attached to the system through one or more SAN fabrics that are also used for other types of storage devices, you need to create storage group that contains an empty

volume on one of the other devices, as shown in [Figure 53 on page 170](#). This empty volume can be of minimal size because you are not going to use it to store any data; you need it only for DPM to detect it as part of the storage configuration.

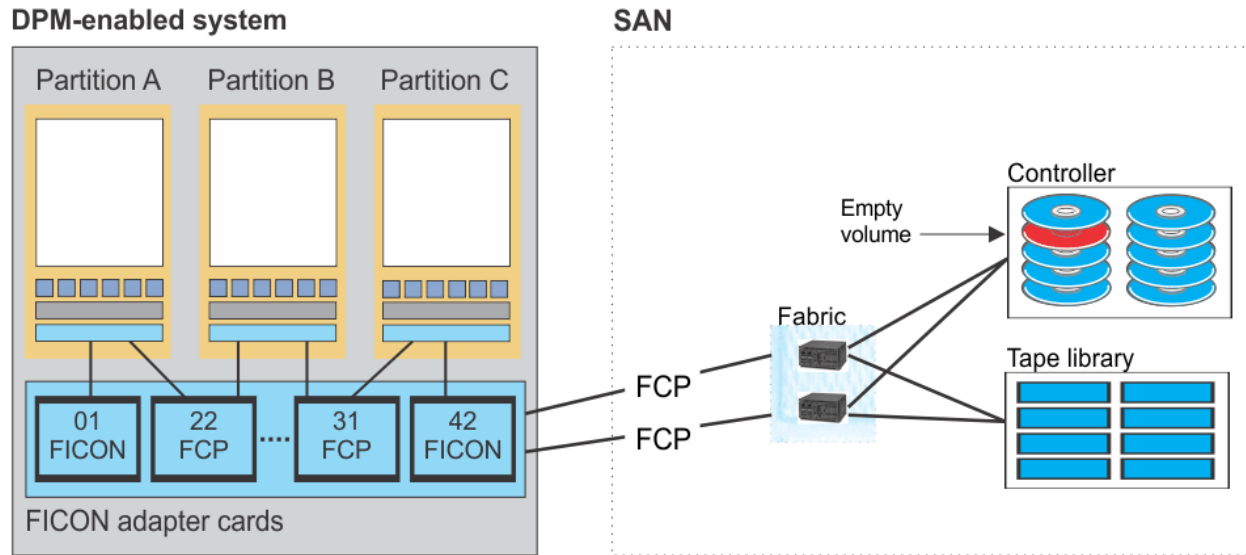


Figure 53. A sample SAN configuration with a shared fabric for access to FCP tape storage drives and other storage devices

- Through the **Configure Storage** task, you create one FCP storage group that contains one volume of minimal size. Through this process, DPM automatically generates the world wide port names (WWPNs) that are required for configuring the tape drive.

Note: For better performance, keep access to the tape drive separate from access to other storage volumes. To achieve this separation, make sure that this storage group does not contain anything other than the empty volume. To provide access to other volumes that reside on the same controller as the empty volume, create a different FCP storage group that does not include this empty volume.

- Depending on the established roles in your company, the storage, tape library, and SAN administrators must use the appropriate tools to zone the WWPNs to access the FCP storage disk and the required tape libraries. They must also restrict the WWPNs to only access the required FCP storage disks or tape drives, wherever possible.

Because DPM can detect the empty volume, the fulfillment state of the storage group changes to Complete.

- Through the **Partition Details** task, you attach the new storage group to one or more partitions that need access to the tape drive. Through this process, DPM automatically generates the host bus adapters (HBAs) that are required for accessing the tape drive. With this last step, the operating systems running on partitions to which the storage group is attached are able to access the tape drive.

Procedure

- On the HMC, use the **Configure Storage** task to create an FCP storage group that contains one volume. The following steps provide basic instructions for using the **Configure Storage** task. If you need any more information, see the online help for the task.
 - Open the **Configure Storage** and select the appropriate link: **REQUEST STORAGE** or **CREATE STORAGE GROUP**. (The link label varies, depending on the authorization of your user ID.)
 - On the **Specify Storage Attributes** page, specify the attributes that you want for the new storage group.

- For **Type**, select **FCP**.
 - For **Shareability**, select either **Dedicated** or **Shared**. If you select **Dedicated**, you can attach this storage group to only one partition.
 - For **Connectivity**, specify the number of paths to be available for use by each operating system with access to this storage group. This number should match your cabling setup (and should reflect how many adapter ports connect to the tape drive).
 - To go to the next page, select **NEXT**.
- c) On the **Add Storage Volumes** page, specify one volume of minimal size. To go to the next page, select **NEXT**.
- d) On the **Name and Duplicate** page, specify the name of your new storage group, and select **NEXT**. The suggested practice is to use a name or add a description indicating that this storage group provides access to tape drives.
- e) On the **Confirm** page, review the summary of your storage request, and select **NEXT**.
- f) On the next page, review the automatically generated storage request and send it to one or more storage administrators.
- The content of this page varies, depending on whether your installation has an SMTP server configured for integrated requests and notifications. If an SMTP server is configured, the page heading is **Send Request**; otherwise, the page heading is **Manually Send Request**. To complete your storage request, follow the steps for the appropriate page.

On the Send Request page

- i) Select one or more storage administrators to receive your request.
- ii) Optional: Add a personal message to the generated request.
- iii) Select **SEND REQUEST** to send the request. If you requested duplicate storage groups, the text on this button indicates the number of requests to be sent.

On the Manually Send Request page

- i) Either download or copy the generated request and send it to one or more storage administrators.
- ii) Select **FINISH** to continue.

When you send or finish your request for the storage group, DPM automatically generates the WWPNS for the storage administrator to use in step “2” on page 171. Also, depending on the value that you specified for **Connectivity**, DPM creates the same number of HBAs for the partitions to which the storage group is attached.

2. Depending on the established roles in your company, the storage, tape library, and SAN administrators must use the appropriate tools to zone the WWPNS to access the FCP storage disk and the required tape libraries. They must also restrict the WWPNS to only access the required FCP storage disks or tape drives, wherever possible.

After the WWPNS are zoned, DPM can detect the empty volume, the fulfillment state of the storage group changes to Complete.

3. On the HMC, use the **Partition Details** task to attach the new FCP storage group to one or more partitions on the system. If you selected **Shared** as the shareability attribute for the storage group, repeat this step for each partition that needs to have access to the tape drive.

The following steps provide basic instructions for using the **Partition Details** task. If you need any more information, see the online help for the task. Note that the partition can be in any state when you complete this step.

- a) Open the **Partition Details** task for a partition and navigate to the **Storage** section.
- b) In the Storage Group table toolbar, select the plus icon to open the **Attach Storage Groups** window.
- c) On the **Attach Storage Groups** window, select the storage group that you just created.
- d) Select **OK** to close the **Attach Storage Groups** window.

e) Select **OK** to save the partition definition.

DPM automatically generates host bus adapters (HBAs) for the partition and assigns them to adapters that enable access to the tape drive.

Results

The storage group configuration is complete. On the **Configure Storage** task pages that display the fulfillment state, the fulfillment state of this storage group is Complete, and operating systems are able to access the tape drive.

What to do next

For any partitions that were in Active state when you attached the storage group in step “3” on page 171, you need to make sure that the HBAs of the storage group are visible to the operating system that each partition hosts. For more details, go to the **What to do next** topic *Ensuring that storage groups are visible to the Linux operating system* in “Starting a partition and its operating system or hypervisor” on page 49, and review the FCP information.

Configuring and accessing tape drives through a dedicated fabric

If your tape drives are accessible through one or more SAN fabrics that are dedicated to tape storage, complete the steps in this procedure to correctly configure them for use. A system administrator can perform most of these steps, but step “3” on page 175 requires the skills of a storage administrator to fulfill the storage request. These configuration steps include working with DPM tasks on the HMC, and also running an IBM-supplied script that uses the Hardware Management Console (HMC) Web Services API. Although the actions that you can perform through the HMC can also be completed through the appropriate API, the suggested practice is to complete the steps using the interfaces as described in this procedure.

Before you begin

- For the best results, upgrade to the most recent HMC/SE bundles for a system that is running with DPM R3.1 or later.
- Download `fcg_tape_config.py`, which is a Python script that contains the API calls that you need to complete the configuration. You can access the download page by selecting the link for *Web Services DPM tape support script* on the IBM Z APIs page in Resource Link: <https://www.ibm.com/servers/resourcelink/lib03020.nsf/pages/zSeriesAPI?OpenDocument&pathID=>

You need to edit and run this script, as described in steps “5” on page 175 and “6” on page 176 of this procedure, from the command line of a machine that has access to the HMC and has Python Version 2.7 installed (the script will not run on later Python versions).

If you need more information about the syntax and descriptions of the APIs that are used in this script, see *Hardware Management Console Web Services API*, which is available on Resource Link at <http://www.ibm.com/servers/resourcelink> on the Library page for the system that you are using.

- To run the `fcg_tape_config.py` script, you need a user ID that has authorization to use the **Configure Storage** task and also has permission to access related objects: the CPC, storage groups, and adapters. The suggested practice is to use the **User Management** task to define a user based on the default SYSPROG user ID, and add the following storage roles:
 - Storage Administrator Objects
 - Storage Administrator Tasks

For instructions, see the topic *Authorizing users of the Configure Storage task*, in either the HMC online help for the **Configure Storage** task or in the *DPM Guide*. After creating the user ID, make sure that you enable this user ID to use the HMC Web Services API.

- Make sure that you have enabled the use of the HMC Web Services API on the system through the HMC **Customize API Settings** task.

About this task

If the tape drives that you want to use are attached to the system through one or more SAN fabrics that are dedicated for tape drives only, as shown in [Figure 54 on page 173](#), you need to use a combination of DPM tasks and Web Services APIs to configure access to the tape drive.

1. Through the DPM tasks, you create and attach one FCP storage group to one or more partitions. The storage group contains only one volume that does not need to be physically configured in the SAN. Through this process, DPM automatically generates the world wide port names (WWPNs) and host bus adapters (HBAs) that are required for configuring and accessing the tape drive.
2. Depending on the established roles in your company, the storage, tape library, and SAN administrators must use the appropriate tools to zone the WWPNs to access the required tape libraries. They must also restrict the WWPNs to only access the required tape drives, wherever possible.
3. Through the `fcplib.py` script, which contains the Web Services API calls, you manually fulfill the storage group by first adding candidate adapter ports to the storage group, and then mapping the volume in the storage group to a dummy logical unit.

Because DPM cannot detect the tape drive and no LUN configurations are done for the volume in the storage group, the storage group remains in Pending fulfillment state. However, the operating systems running on partitions to which the storage group is attached are able to access the tape drive.

Important: Make sure that you complete the steps in this procedure in sequence. You must attach the storage group to all partitions that need access to the tape drive *before*:

- The storage administrator configures the WWPNs in step “3” on [page 175](#).
- You run the script in step “6” on [page 176](#).

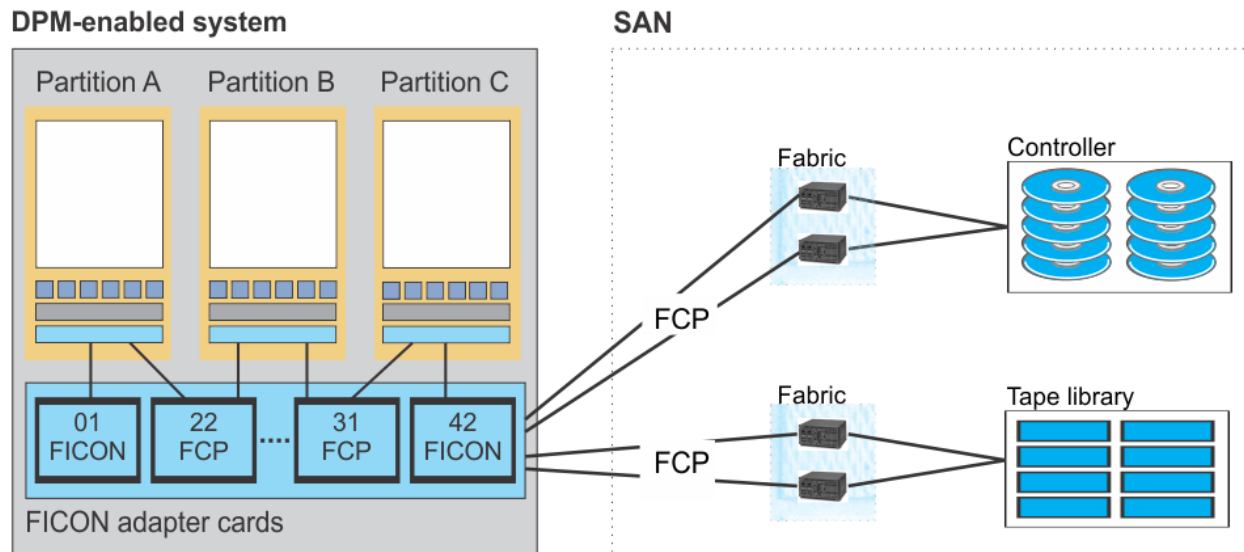


Figure 54. A sample SAN configuration with a dedicated fabric for access to FCP tape storage

Procedure

1. On the HMC, use the **Configure Storage** task to create an FCP storage group that contains one volume. The following steps provide basic instructions for using the **Configure Storage** task. If you need any more information, see the online help for the task.
 - a) Open the **Configure Storage** and select the appropriate link: **REQUEST STORAGE** or **CREATE STORAGE GROUP**. (The link label varies, depending on the authorization of your user ID.)

- b) On the **Specify Storage Attributes** page, specify the attributes that you want for the new storage group.
- For **Type**, select **FCP**.
 - For **Shareability**, select either **Dedicated** or **Shared**. If you select **Dedicated**, you can attach this storage group to only one partition.
 - For **Connectivity**, specify the number of paths to be available for use by each operating system with access to this storage group. This number should match your cabling setup (and should reflect how many adapter ports connect to the tape drive).
 - To go to the next page, select **NEXT**.
- c) On the **Add Storage Volumes** page, specify one boot volume of any random size. (This volume does not need to be physically configured in any storage subsystem.) To go to the next page, select **NEXT**.
- d) On the **Name and Duplicate** page, specify the name of your new storage group, and select **NEXT**.
The suggested practice is to use a name or add a description indicating that this storage group provides access to tape drives.
- e) On the **Confirm** page, review the summary of your storage request, and select **NEXT**.
- f) On the next page, review the automatically generated storage request and send it to one or more storage administrators.

The content of this page varies, depending on whether your installation has an SMTP server configured for integrated requests and notifications. If an SMTP server is configured, the page heading is **Send Request**; otherwise, the page heading is **Manually Send Request**. To complete your storage request, follow the steps for the appropriate page.

On the Send Request page

- i) Select one or more storage administrators to receive your request.
- ii) Add a personal message to the generated request, asking the storage administrator to ignore the FCP disk configuration, and to only zone and map the WWPNs to the tape drives.
- iii) Select **SEND REQUEST** to send the request. If you requested duplicate storage groups, the text on this button indicates the number of requests to be sent.

On the Manually Send Request page

- i) Either download or copy the generated request and add a personal message, asking the storage administrator to ignore the FCP disk configuration, and to only zone and map the WWPNs to the tape drives.
- ii) Send the request to one or more storage administrators.
- iii) Select **FINISH** to continue.

When you send or finish your request for the storage group, DPM automatically generates the WWPNs for the storage administrator to use in step “3” on page 175. Also, depending on the value that you specified for **Connectivity** in step “1.b” on page 174, DPM creates the same number of HBAs for the partitions to which the storage group is attached.

After the storage group is created, go to the **Storage Overview** page, select the storage group, and record the Storage group ID (UUID) value that is displayed in the summary section of the **Storage Group details** page. You need this value as input for step “5” on page 175.

2. On the HMC, use the **Partition Details** task to attach the new FCP storage group to one or more partitions on the system. If you selected **Shared** as the shareability attribute for the storage group, repeat this step for each partition that needs to have access to the tape drive.

The following steps provide basic instructions for using the **Partition Details** task. If you need any more information, see the online help for the task. Note that the partition can be in any state when you complete this step.

- a) Open the **Partition Details** task for a partition and navigate to the **Storage** section.

- b) In the Storage Group table toolbar, select the plus icon to open the **Attach Storage Groups** window.
 - c) On the **Attach Storage Groups** window, select the storage group that you just created.
 - d) Select **OK** to close the **Attach Storage Groups** window.
 - e) Select **OK** to save the partition definition.
- DPM automatically generates host bus adapters (HBAs) for the partition but they are not assigned to any adapters.
3. Send your request to the storage administrator to fulfill the storage group.
Depending on the established roles in your company, the storage, tape library, and SAN administrators must use the appropriate tools to zone the WWPNs to access the required tape libraries. They must also restrict the WWPNs to only access the required tape drives, wherever possible.
 4. Collect the universally unique identifiers (UUIDs) of the adapters that you want to list as candidates for your new storage group.
 - a) On the HMC, open the **Manage Adapters** task to the Adapter Details page.
 - b) If necessary, sort the list of adapters by type to display the FCP adapters in sequence.
 - c) For each adapter that you want to use, record the value in the Object ID column (for example: B0D3C2F0-4FC4-11E9-B8FD-00106F0D81C9).
Use these object ID values to create the adapter port list that you need to run the Add Candidate Adapter Ports to an FCP Storage Group API in step “5” on page 175.

Important: The number of adapters in the list must equal or exceed the value that you specified for **Connectivity** in step “1.b” on page 174.
 5. Edit the `fcg_tape_config.py` script to run it in your environment.
The script uses the following APIs to complete the configuration:
 - Add Candidate Adapter Ports to an FCP Storage Group, which uses the object IDs that you collected in step “4.c” on page 175, and adds them to the storage group that you created in step “1” on page 173.
 - Fulfill FCP Storage Volume, which maps the storage volume to a dummy logical unit. This mapping results in the manual fulfillment of the storage volume. (The *target-wwpn* and *lun-number* input values for this API are provided as part of the script, and do not need to match any real values from your storage configuration.)

Update the Declaration Section of the script with the following information.

| Table 13. Variables to specify in the Declaration Section of the script | |
|---|--|
| Variable | Description |
| <i>hmc_address</i> | The IP address of the HMC. |
| <i>hmc_userid</i> | The username that the API script must use to make the API queries. This user ID must have the authorization described in <i>Before you begin</i> . |
| <i>hmc_password</i> | The password associated with the user ID that you specified in <i>hmc_userid</i> . |
| <i>cpc_name</i> | The name of the system on which the storage group is available. |
| <i>storage_group_uuid_customer</i> | The UUID of the storage group that contains the volume that must be manually fulfilled. To find the UUID, open the Configure Storage task to the Storage Overview , select the storage group, and use the Storage group ID (UUID) value in the summary section of the Storage Group details page. |

Table 13. Variables to specify in the Declaration Section of the script (continued)

| Variable | Description |
|---------------------------------|--|
| <i>adapter_port_list</i> | A list of the UUIDs of each adapter that should be added as a candidate for the storage group. This list consists of the object IDs that you collected in step “4.c” on page 175. As stated in that step, the number of adapters in the list must equal or exceed the value that you specified for Connectivity in step “1.b” on page 174. When you run the script in the next step, DPM selects the first adapter port in this list as input for the Fulfill FCP Storage Volume operation. |

6. After the storage administrator has fulfilled the storage group request, run the script from the command line of a machine that has access to the HMC and has Python Version 2.7 installed (the script will not run on later Python versions). Use this command: `python fcp_tape_config.py`
 - DPM assigns the partition HBAs created in step “1” on page 173 to the adapters that you specified in the *adapter_port_list* variable for the Add Candidate Adapter Ports to an FCP Storage Group API operation.
 - The volume UUID contains a value that is a combination of the *target-wwpn* and *lun-number* values that are provided in the script for the Fulfill FCP Storage Volume API operation. The volume UUID is displayed on the **VOLUMES** tab on the **Storage Group details** page of the **Configure Storage** task.

Results

The storage group configuration is complete. On the **Configure Storage** task pages that display the fulfillment state, the fulfillment state of this storage group remains in Pending state; however, operating systems are able to access the tape drive.

What to do next

For any partitions that were in Active state when you attached the storage group in step “2” on page 174, you need to make sure that the HBAs of the storage group are visible to the operating system that each partition hosts. For more details, go to the **What to do next** topic *Ensuring that storage groups are visible to the Linux operating system* in “Starting a partition and its operating system or hypervisor” on page 49, and review the FCP information.

Note:

If the storage group is detached from a partition for any reason, you need to rerun the `fcp_tape_config.py` before reattaching the storage group to the partition.

Appendix E. 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 or registered trademarks of International Business Machines Corporation, in the United States and/or other countries. Other product and service names

might be trademarks of IBM or other companies. A current list of IBM trademarks is available on <http://www.ibm.com/trademark>.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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

accelerator
 adapter [23](#)
 accessibility
 contact IBM [xvii](#)
 features [xvii](#)
 adapter
 definition [3](#), [17](#)
 network [17](#)
 specifying port number [53](#)
 storage [18](#)
 summary of types [17](#)
 assistive technologies [xvii](#)
 attach a storage group [55](#)

B

boot option
 for Linux operating system [26](#)
 for z/VM hypervisor [26](#)
 network server [107](#)
 PXE [107](#)
 Secure Service Container [26](#)
 boot options
 for operating systems [26](#)
 hypervisor [26](#)
 operating system [26](#)
 bundle
 prerequisites for DPM versions [77](#)

C

cabling plan
 sample [165](#)
 carrier card
 NVMe storage [18](#)
 commands
 Linux operating system [54](#)
 z/VM hypervisor [54](#)
 concepts
 hypervisors [25](#)
 operating systems [25](#)
 configure
 FCP adapter card
 initial setup [130](#)
 FICON connection
 initial setup [130](#)
 configure storage task
 tape link details [149](#)
 Configure Storage task
 benefits of using [117](#)
 connect system to storage hardware
 initial setup [130](#)
 planning considerations [130](#)
 prerequisites [130](#)
 sample cabling plan [165](#)

Configure Storage task (*continued*)
 connect system to storage hardware (*continued*)
 sample NVMe adapter plan [167](#)
 task overview [130](#)
 introduction [117](#), [129](#)
 manage tape library
 planning considerations [158](#)
 prerequisites [158](#)
 task overview [158](#)
 request NVMe storage
 planning considerations [137](#)
 prerequisites [137](#)
 task overview [137](#)
 request storage
 FCP tape storage drives [169](#)
 planning considerations [134](#)
 prerequisites [134](#)
 task overview [134](#)
 request tape link
 planning considerations [139](#)
 prerequisites [139](#)
 task overview [139](#)
 required authorization
 procedure [129](#)
 view FCP tape link
 modify tape link [143](#), [148](#)
 view environment display [148](#)
 view storage group
 map FICON volume [144](#)
 modify storage group [143](#), [144](#)
 view connection report [144](#)
 connection report
 for storage group [143](#), [144](#)
 corequisite software [84](#)
 create a partition
 comparison of New Partition task modes [37](#)
 overview [5](#)
 selecting New Partition task mode [37](#)
 step-by-step instructions [39](#)
 switching New Partition task modes [38](#)
 crypto features [23](#)
 cryptographic adapter [23](#)

D

detach a storage group [55](#)
 disk storage
 request for partition use
 with or without template [134](#)
 DPM
 advantages [4](#)
 enabling for a system [4](#)
 how to get started [5](#)
 introduction [3](#)
 prerequisites [77](#)
 summary of tasks [71](#)
 upgrade [78](#)

DPM (*continued*)
 user interface [3](#)
 DPM-enabled system
 supported types [77](#)
 DVDPRIME command [26](#)
 Dynamic Partition Manager (DPM), *See* DPM

E

echo command [53](#)
 engineering change (EC)
 definition [77](#)
 list for DPM versions [77](#)
 environment display
 for FCP tape link [148](#)

F

FCP adapter card
 change adapter assignment [143](#), [144](#)
 configure
 initial setup [130](#)
 in-use definition [156](#)
 modify [156](#)
 FCP mode)
 storage adapter [18](#)
 FCP storage [117](#)
 FCP tape link
 environment display [148](#)
 modify [148](#)
 view [148](#)
 FCP tape storage drive
 configure and access [169](#)
 feature code
 for Dynamic Partition Manager [77](#)
 for OSA-Express 1000BASE-T Ethernet [77](#)
 Fibre Channel connection (FICON), *See* FICON
 Fibre Channel Protocol (FCP), *See* FCP
 FICON adapter card
 configure
 initial setup [130](#)
 in-use definition [156](#)
 modify [156](#)
 FICON connection
 modify [156](#)
 FICON connections
 configure
 initial setup [130](#)
 FICON mode)
 storage adapter [18](#)
 FICON storage [117](#)
 FICON volume
 mapping [143](#), [144](#)
 file transfer protocol [111](#)
See also FTP, FTPS, SFTP
 FTP
 compared to FTPS or SFTP [111](#)
 FTP Secure, *See* FTPS
 FTPS
 boot option [26](#)
 setting up on HMC [111](#)

H

Hardware Management Console, *See* HMC (Hardware Management Console)
 hardware planning [84](#)
 HBA conversion
 to storage group [91](#)
 HiperSocket)
 network adapter [17](#)
 HMC (Hardware Management Console)
 list of Dynamic Partition Manager tasks [71](#)
 prerequisites for DPM versions [77](#)
 hypervisor
 partition type [25](#)

I

Integrated 3270 Console task [29](#)
 Integrated ASCII Console task [29](#)

K

keyboard
 navigation [xvii](#)
 KVM on IBM z
 required version [84](#)

L

Linux operating system
 boot options [26](#)
 commands
 Linux operating system [17](#)
 lschp [52](#)
 lscpu [51](#)
 lscss [52](#), [53](#)
 lsmem [51](#)
 lspci [53](#)
 lsqeth [52](#)
 lstape [53](#)
 concepts [25](#)
 device drivers for adapters [17](#), [54](#)
 Integrated ASCII Console task [29](#)
 link to library [29](#)
 lscpu command [51](#)
 lsmem command [51](#)
 sample configuration [11](#)
 secure execution [30](#), [54](#), [101](#)
 specifying port number [53](#)
 verifying partition resources [51](#)
 Linux system
 required version [84](#)
 lschp command [52](#)
 lscss command [52](#), [53](#)
 lspci command [53](#)
 lsqeth command [52](#)
 lstape command [53](#)

M

MAC address
 specify for a NIC [113](#)
 machine change level (MCL)

- machine change level (MCL) (*continued*)
 - definition [77](#)
- media access control address [113](#)
- memory
 - dynamically add [55](#)
- MES frame upgrade
 - migration instructions [87](#)
 - supported migration options [87](#)
- migration considerations
 - from prior system to new system
 - supported migration options [87](#)
 - HBA conversion
 - to storage group [91](#)
 - system replacement [91](#)
- migration options
 - MES frame upgrade [87](#)
- miscellaneous equipment specification, *See* MES frame
- upgrade
- modify
 - FCP adapter card [156](#)
 - FICON connection [156](#)
 - partition resources [55](#)

N

- navigation
 - keyboard [xvii](#)
- network adapter
 - dynamically add [55](#)
 - summary of types [17](#)
- network adapters
 - list of supported types [82](#)
- network boot server [107](#)
- Network interface card (NIC) [17](#)
- New Partition
 - step-by-step instructions [39](#)
- New Partition task
 - attach storage
 - planning considerations [152](#)
 - comparison of modes [37](#)
 - selecting the mode [37](#)
 - switching modes [38](#)
- Non-Volatile Memory Express (NVMe), *See* NVMe storage
- NVMe storage
 - adapter
 - carrier card [18](#)
 - description [18](#)
 - sample plan [167](#)
 - solid state drive (SSD) [18](#)
 - request for partition use [137](#)

O

- Open Systems Adapter (OSA)
 - network adapter [17](#)
- operating system
 - partition type [25](#)
- OSA-Express 1000BASE-T Ethernet
 - feature code [77](#)

P

- partition

- partition (*continued*)
 - boot options [26](#)
 - configuration settings [13](#)
 - definition [3](#)
 - how to start [49](#)
 - lifecycle [15](#)
 - modify definition [55](#)
 - planning checklist [33](#)
 - properties [13](#)
 - sample configuration [11](#)
 - status [15](#)
 - type
 - supported operating system or hypervisor [25](#)
 - verifying resources in Linux [51](#)
- partition definition
 - creating through New Partition task [12](#)
- Partition Details task
 - add adapter [55](#)
 - add memory [55](#)
 - add processors [55](#)
 - attach storage
 - planning considerations [152](#)
 - attach storage group [55](#)
 - reattach storage group [55](#)
- planning checklist
 - for creating a partition [33](#)
- port number
 - specifying in Linux [53](#)
- preboot execution environment (PXE) [107](#)
- prerequisites
 - DPM-enabled system [77](#)
 - KVM on IBM z [84](#)
 - Linux system
 - required version [84](#)
- Preventative Service Planning buckets, *See* PSP buckets
- processors
 - dynamically add [55](#)
- PSP buckets
 - corequisite software [84](#)
 - hardware planning [84](#)
 - where to find [84](#)
- PXE [107](#)

R

- reattach a storage group [55](#)
- Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE)
 - network adapter [17](#)
- revisions [xix](#)

S

- SE (Support Element)
 - list for DPM versions [77](#)
 - list of Dynamic Partition Manager tasks [71](#)
 - prerequisites for DPM versions [77](#)
- Secure Execution for Linux [30](#)
- Secure File Transfer Protocol, *See* SFTP
- Secure Service Container
 - boot option [26](#)
 - concepts [25](#)
 - definition [11](#)

- SFTP
 - boot option [26](#)
 - setting up on HMC [111](#)
- shortcut keys [xvii](#)
- solid state drive (SSD)
 - NVMe storage [18](#)
- start a partition
 - step-by-step instructions [49](#)
- status
 - of a partition [15](#)
- storage
 - attach to partition [152](#)
 - request for partition use
 - FCP tape storage drives [169](#)
 - optimize for 2nd level virtualization [136](#)
- storage adapter
 - FICON and FCP overview [83](#)
 - NVMe overview [84](#)
 - protocols [18](#)
 - summary of types [18](#)
- storage group
 - attach [55](#)
 - attach to partition [152](#)
 - attributes [134](#)
 - change FCP adapter assignment [143](#), [144](#)
 - created from HBAs during migration [91](#)
 - definition [18](#), [134](#)
 - detach [55](#)
 - FCP tape storage drives [169](#)
 - modify [143](#), [144](#)
 - NVMe [137](#)
 - optimize for 2nd level virtualization [136](#)
 - reattach [55](#)
 - request [134](#), [136](#), [169](#)
 - view [143](#), [144](#)
 - with or without template [134](#)
- storage groups
 - template [127](#), [134](#)
- storage management
 - introduction [119](#)
- STORAGEADMIN default user ID [129](#)
- Support Element, *See* SE (Support Element)
- system
 - connect to storage
 - initial setup [130](#)
 - modify [156](#)
 - sample cabling plan [165](#)
 - sample NVMe adapter plan [167](#)
 - types that support DPM [77](#)
- system replacement
 - migration instructions [91](#)
 - required MCLs [91](#)

T

- tape library
 - manage [158](#)
- tape link
 - attach to partition [152](#)
 - attributes [139](#)
 - definition [139](#)
 - request [139](#)
- tape storage
 - request for partition use [139](#)

- tape storage drive (FCP)
 - configure and access [169](#)
- template
 - for storage group [127](#), [134](#)

U

- upgrade
 - DPM [78](#)
- UUID
 - use in network boot server [107](#)

V

- versioning [77](#)
- virtual local area network
 - specify an identifier [113](#)
- VLAN enforcement [113](#)
- VLAN identifier (ID) [113](#)

Z

- z/VM
 - boot options [26](#)
 - concepts [25](#)
 - DVDPRIME command [26](#)
 - installation [26](#)
 - Integrated 3270 Console task [29](#)
 - link to library [29](#)
 - RAMDISK [26](#)
- z/VM hypervisor
 - commands [54](#)
 - sample configuration [11](#)



SB10-7176-04

