

z Systems IBM Dynamic Partition Manager (DPM) Guide

SB10-7168-01



z Systems IBM Dynamic Partition Manager (DPM) Guide

SB10-7168-01

#### Note

Before you use this information and the product it supports, read the information in "Safety" on page ix, Appendix B, "Notices," on page 85 and *IBM Systems Environmental Notices and User Guide*, Z125–5823.

This edition, SB10-7168-01, applies to IBM z Systems (z Systems) and IBM LinuxONE (LinuxONE) servers. This edition replaces SB10-7168-00.

There might be a newer version of this document in a **PDF** file available on **Resource Link**.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. A newer version is indicated by a lowercase, alphabetic letter following the form number suffix (for example: 00a, 00b, 01a, 01b).

#### © Copyright IBM Corporation 2016, 2017.

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

# Contents

	•
	Tables
	Safety       ix         Safety notices       ix         World trade safety information       ix         Laser safety information       ix         Laser compliance       ix
1	About this publication       xi         Intended audience       xi         Prerequisite and related information       xi         Related HMC and SE console information       xii         Related information for Linux on z Systems       xii         Related information for KVM for IBM z Systems       xii         Related information for z/VM       xii         Related information for IBM Secure Service       xii
I	Container
	Summary of changes
	Part 1. Introduction to the IBM
	Chapter 1. Dynamic Partition Manager: A
	Chapter 1. Dynamic Partition Manager: A quicker and easier way to deploy Linux servers on a mainframe
	Chapter 1. Dynamic Partition Manager
	Chapter 1. Dynamic Partition Manager: A quicker and easier way to deploy Linux servers on a mainframe
	Chapter 1. Dynamic Partition Manager: A         quicker and easier way to deploy Linux         servers on a mainframe
	Chapter 1. Dynamic Partition Manager: A quicker and easier way to deploy Linux servers on a mainframe

Part 3. Basic tasks for Linux
Chapter 6. Planning checklist for creating a partition
Chapter 7 Creating and starting a new
partition
Selecting which New Partition task mode to use
Creating a new partition
Starting a partition and its operating system or hypervisor.
Chapter 8 Dynamically modifying the
resources of an active partition 49
Chapter 9. Summary of tasks for
managing systems, adapters, and
partitions
Part 4. Topics for system planners 61
Obenter 10. Drevenuisites for using
Chapter 10. Prerequisites for using
Chapter 11. DPM migration
considerations
Chapter 12. Adapter configuration 69
Part 5. Topics for network
administrators
Chapter 13. Customizing the network
boot server environment
Deut C. Annuardina a
Part 6. Appendixes 81
Appendix A Dynamic Partition
Manager task roles
Annendix B. Notices 95
Appendix D. Notices
Class A Notices
inaex

# Figures

1.	New user interface design for the HMC
	Welcome page
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
4.	Validation window for the New Partition task 7
5.	Methods of starting a new partition
6.	Portion of the Start window for a new partition
	named ASL Test
7.	Sample Linux command and display of
	configured processor resources 8
8.	Partitions configured on a DPM-enabled
	system
9.	Partitions with NICs configured for access to
	network adapters on a DPM-enabled system . 18
10.	Partitions with HBAs configured for access to
	storage adapters on a DPM-enabled system
11.	Example: Linux servers configured in
	partitions on a DPM-enabled system
12.	Welcome page of the New Partition task in
	basic mode
13.	Processors page of the New Partition task in
	basic mode
14.	Memory page of the New Partition task in
	basic mode
15.	Network page of the New Partition task in
	basic mode
16.	Storage page of the New Partition task in basic
	mode
17.	Boot page of the New Partition task in basic
	mode
18.	Summary page of the New Partition task in
	basic mode
19.	Validation window for the New Partition task 42
20.	Portion of the Start window for a new
	partition named ASL_Test
21.	Sample displays resulting from the lscpu
	command

22.	Sample displays resulting from the <b>lsmem</b>	
	command	5
23.	Sample displays resulting from the <b>lsqeth</b> and	
	lscss commands	5
24.	Sample displays resulting from the lschp	
	command	7
25.	Sample display for the <b>lsmem</b> command 49	)
26.	Sample screen of the Memory section in the	
	Partition Details window	)
27.	Sample display for the <b>chmem</b> command 50	)
28.	Sample display for the <b>lscpu</b> command	
	showing current processor resources	Ĺ
29.	Sample screen of the Processors section in the	
	Partition Details window	Ĺ
30.	Sample display for the <b>lscpu</b> command	
	showing additional processor resources 52	,
31.	Sample display for the <b>lscpu</b> command	
	showing configured processor resources	,
32.	Sample display for the <b>lscss</b> command	
	showing current devices	3
33.	Sample screen of the Network section with the	
	New NIC icon highlighted	3
34.	Sample screen of the New Network Interface	
	Card window	ł
35.	Sample screen of the Network section with the	
	newly added NIC	ł
36.	Sample display for the <b>znetconf</b> command	
	showing current devices	5
37.	Sample display for the <b>lsqeth</b> command	
	showing details for devices	5
38.	Sample screen of the Manage Adapters task 70	)
39.	Sample screen of the Adapter Details window,	
	with Name and Description fields highlighted . 71	L
40.	Sample screen of the Manage Adapters task,	
	with the expanded Actions list	)
41.	Sample screen of the New HiperSockets	
	Adapter window	,
42.	Sample screen of the new HiperSockets	
	adapter in the Adapter Details window 73	3

# Tables

1.	Planning checklist for defining your new
	partition
2.	Comparison of editable partition properties in
	the basic and advanced New Partition task
	modes
3.	Summary of key HMC/SE tasks and displays
	for working with DPM-enabled systems,
	partitions, and adapters

4.	Required feature codes for IBM Dynamic	
	Partition Manager (DPM)	63
5.	Summary of DPM releases and their associated	
	EC MCLs	64
6.	Actions for migrating from DPM R1 to R2	67
7.	Sample default names for installed adapters	69
8.	DPM task roles mapping legend	83
9.	DPM task roles mapping	83

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

## Laser compliance

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

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

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

# About this publication

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

DPM is a feature of the following IBM z Systems<sup>™</sup> (z Systems) and IBM LinuxONE (LinuxONE) servers:

- An IBM z13<sup>™</sup> (z13<sup>™</sup>) or IBM z13s<sup>™</sup> (z13s<sup>™</sup>)
- An IBM LinuxONE Emperor<sup>™</sup> (Emperor) or IBM LinuxONE Rockhopper<sup>™</sup> (Rockhopper)

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 at http://www.ibm.com/systems/z/os/linux/resources/testedplatforms.html.
- KVM for IBM z Systems, which is a virtualization hypervisor based on Open Source technology, on which you can run multiple Linux images.
- z/VM<sup>®</sup>, which is an IBM virtualization hypervisor on which you can run multiple Linux images.
- I DPM also supports IBM Secure Service Container partitions.

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

## Intended audience

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

Other IT personnel who might benefit from using this book include mainframe operators, automation engineers, and experienced mainframe system administrators who need to learn how to use DPM tasks, which are available starting with Hardware Management Console (HMC) and Support Element (SE) Version 2.13.1.

IBM service representatives who are responsible for installing IBM z Systems (z Systems) and IBM LinuxONE (LinuxONE) servers and enabling DPM use the following books:

- z Systems z13 Installation Manual, GC28-6936 or z Systems z13s Installation Manual, GC28-6952
- z Systems z13 Service Guide, GC28-6937, or z Systems z13s Service Guide, GC28-6955

## 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 specific mainframes, see the appropriate product document on the IBM Redbooks<sup>®</sup> web site at http://www.redbooks.ibm.com/
  - IBM z13 Technical Guide, SG24-8251
  - IBM z13s Technical Guide, SG24-8294

 For information about the DPM APIs, see *z Systems Hardware Management Console Web Services API*, SC27-2634, which is available through the Publications link on IBM Resource Link<sup>®</sup> at http://www.ibm.com/servers/resourcelink

## **Related HMC and SE console information**

Hardware Management Console (HMC) and Support Element (SE) information can be found on the console help system, or on the IBM Knowledge Center at http://www.ibm.com/support/knowledgecenter/. Click **Select a product**, and then select your server.

## **Related information for Linux on z Systems**

For information about installing and running a Linux distribution on a z Systems or LinuxONE server, see the Linux on z Systems topics in IBM Knowledge Center, at this URL: https://www.ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz\_r\_lib.html

Another useful source of information is the IBM developerWorks<sup>®</sup> site for Linux at http://www.ibm.com/developerworks/linux/linux390/documentation\_dev.html

## Related information for KVM for IBM z Systems

For information about installing and running KVM for IBM z as a hypervisor for Linux guests on a z Systems or LinuxONE server, see KVM for IBM z Systems topics in IBM Knowledge Center, at this URL: https://www.ibm.com/support/knowledgecenter/SSNW54

## Related information for z/VM

For information about installing and running z/VM as a hypervisor for Linux guests on a z Systems or LinuxONE server, see the z/VM topics in IBM Knowledge Center, at this URL: https://www.ibm.com/support/knowledgecenter/SSB27U

## Related information for IBM Secure Service Container

For information about working with Secure Service Container partitions and the appliances they support,
 see *IBM z Systems<sup>™</sup> Secure Service Container User's Guide*, SC28-6971, which is available on
 http://www.ibm.com/servers/resourcelink.

## How to use this publication

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

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

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

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

#### Part 2, "Basic concepts and terms for Linux administrators," on page 9

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

#### Part 3, "Basic tasks for Linux administrators," on page 25

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 61

Topics in this part provide the prerequisites for enabling DPM on a mainframe system, information about supported functions, and the engineering changes (EC) or microcode control levels (MCL) for upgrading to the latest DPM version. These topics are appropriate for experienced system planners and other administrators who are familiar with mainframe systems.

#### Part 5, "Topics for network administrators," on page 75

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

#### Appendixes

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

## Accessibility

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

Accessible publications for this product are offered in HTML 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 Systems<sup>®</sup> (z Systems<sup>®</sup>) and IBM LinuxONE (LinuxONE) information, go to Resource Link at http://www.ibm.com/servers/resourcelink and click **Feedback** from the navigation bar on the left. In the **Comments** input area, state your question or comment, the publication title and number, choose **General comment** as the category and click **Submit**. You can also send an email to reslink@us.ibm.com providing the same information.

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

## Accessibility features

The following list includes the major accessibility features in z Systems<sup>®</sup> and IBM LinuxONE documentation:

- 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.

## 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 vertical bar (1) to the left of the change.

### Summary of changes for SB10-7168-01

This edition contains the following updates. If a specific update is available through an engineering change (EC) microcode control level (MCL) for the host system, the MCL number is provided in Chapter 10, "Prerequisites for using Dynamic Partition Manager," on page 63.

- The following topics have been updated with information about DPM support for partition types, including the IBM Secure Service Container.
  - Chapter 3, "Partitions: Virtual images of a mainframe system," on page 11
  - Chapter 5, "Operating systems and hypervisors," on page 21
  - Chapter 6, "Planning checklist for creating a partition," on page 27
  - Chapter 7, "Creating and starting a new partition," on page 31
  - Chapter 9, "Summary of tasks for managing systems, adapters, and partitions," on page 57
  - Chapter 10, "Prerequisites for using Dynamic Partition Manager," on page 63
  - Chapter 11, "DPM migration considerations," on page 67
- The following topics have been updated with information about displaying the number of processor threads through the **Partition Details** task, when simultaneous multithreading is enabled.
  - Chapter 3, "Partitions: Virtual images of a mainframe system," on page 11
  - Chapter 6, "Planning checklist for creating a partition," on page 27
  - Chapter 7, "Creating and starting a new partition," on page 31
  - Chapter 10, "Prerequisites for using Dynamic Partition Manager," on page 63

# Part 1. Introduction to the IBM Dynamic Partition Manager

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

Topics covered in this part are:

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

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

Linux servers and applications have been running on mainframe systems for years, but configuration and setup is fairly complicated and requires the use of several z Systems-specific tools. With IBM Dynamic Partition Manager (DPM), however, system administrators now have a quicker and easier way to deploy Linux servers, using only the Hardware Management Console (HMC).

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

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



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

When your company orders an IBM system with the requisite 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 IBM mainframe systems:

- An IBM z13 (z13) or IBM z13s<sup>TM</sup> (z13s<sup>TM</sup>)
- An IBM LinuxONE Emperor (Emperor) or IBM LinuxONE Rockhopper (Rockhopper)

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. Once 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 that will 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 have completed 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 creating the partition. For additional guidance, use the notes in Chapter 6, "Planning checklist for creating a partition," on page 27.

#### Procedure

- 1. Log in to the Hardware Management Console (HMC) using either the default SYSPROG user ID or a user ID that a system administrator has 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.

dware Manageme	ent Console		11/10-24	IBM.
	S dame to many a state		s	rSPROG   Help   Logoff
	Battions Topology Monitor			
Icome	ramana roporty monto			
tems Management		P 🔐 💌 Filter	Tasks ▼ Views ▼	
1U111 D79PSIRT	Select ^ Name ^ Status ^	Processors ^ Memory ^ Processor (GB) ^ Utilization	Network ^ OS ^ Descr     Utilization ^ Name ^ Descr	iption ^
POIUTEST	DanDPM1 Active	2 2.0 0%	0 % Dan F	artition 1
P0LXSM22 P32	DanDPM2 Active	2 2.0 0%	0 % Dan F	Partition 2
P34 E	DanDPM3 Active	2 2.0 0%	0 % Dan F	fartition 3
R32	🗆 🖏 JeffP1 🛛 😂 Stopped	2 2.0	Jeff P	artition 1
S15PSIRT	🗆 🖉 part 1 😵 Stopped	1 4.0		-
S22PSIRT	Max Page Size	Total: 95 Filtered: 95 Selected: 0		
S27PSIRT S31TRENT				
SETR180	Tasks: \$2028 @ @	Limit		0
SETR186	System Details	Change Management	Configuration	
P200001	Daily	Remote Customization	New Dartition	
S202B	E Service	Operational Customization	Transmit Vital Product	z I Data
emble Management	1 (3 - 1 - 1 - 1 - 1 - 1 - 1 - 1 - 1 - 1 -		View Frame Layout	
tom Groups			Energy Management	

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.
  - a. Open the New Partition task.

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

Welcome Name Processors Memory Network Storage Accelerators Cryptos Boot Summary	<ul> <li>Welcome to the New Partition wizard. Use this wizard to create a DPM partition.</li> <li>This wizard guides you through the following tasks: <ul> <li>Naming and describing the partition</li> <li>Assigning processors and memory</li> <li>Providing access to I/O, including networks, storage, accelerators, and cryptos</li> <li>Configuring the boot device and parameters for loading the partition</li> </ul> </li> </ul>
Related Tasks Monitor System	Show this welcome page next time

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

**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 31.

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.

Note that, in Figure 3, pages that contain required fields are denoted with an asterisk. For step-by-step instructions for using the basic mode to complete the information on these pages, see "Creating a new partition" on page 33.

**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, as shown in Figure 4 on page 7, which displays a message indicating that your partition has been created, and lists additional tasks that you can use to work with the new partition.



Figure 4. Validation window for the New Partition task

- 4. If you defined HBAs for this partition, or have specified the **Storage device (SAN)** boot option, use the **Export WWPNs** task before you start the partition. Your storage administrator needs this list of WWPNs to perform zoning and LUN masking so your partition can successfully access the storage devices.
- 5. Use the **Start** task to start the partition. You can open the **Start** task by clicking the link on the Validation window, or through two other methods shown in Figure 5.

		r incer			Tasks 🔻	Views 🔻
ame	Status     Processo	ors ^ Mer (GB	) Pro	cessor ^	Network A Utilization	OS Name
ASL_Test	Partition Details	1	4.0			
DanDPM1	Daily 🕨	Grouping			0 %	]
DanDPM2	Recovery	Operating S	Operating System Messages		0 %	]
DanDPM3	Operational Customization  Configuration	Start Stop			0 %	]
JeffP1	Stopped	2	2.0			
part 1	Stopped	1	4.0			
Part 10	Stopped	1	4.0			
fest 🕀 🖻	Q- 0-		Ţ			
etails	I Rec	overv			E Conf	iguration
	ASL_Test DanDPM1 DanDPM2 DanDPM3 JeffP1 part 1 Part 10 est <b>① ②</b>	Ask_Test®       Partition Details         DanDPM1       Daily         DanDPM2       Recovery         DanDPM3       Operational Customization         JeffP1       Stopped         part 1       Stopped         Part 10       Stopped	ame       Status       Processors       Men (GB         ASL_Test       Partition Details       1         DanDPM1       Daily       Grouping         DanDPM2       Recovery       Operational Customization       Start         DanDPM3       Configuration       Stopped       2         part 1       Stopped       1         Part 10       Stopped       1	Amme       Status       Processors       Memory (GB)       Production (GB)         ASL_Test       Partition Details       1       4.0         DanDPM1       Daily       Grouping         DanDPM2       Recovery       Operational Customization       Start         DanDPM3       Configuration       Stopped       2       2.0         JeffP1       Stopped       1       4.0         Part 10       Stopped       1       4.0	ame ^ Status ^ Processors ^ Memory (GB) Processor ^   ASL_Test® Partition Details 1 4.0   DanDPM1 Daily > Grouping   DanDPM2 Recovery Operating System Messages   DanDPM3 Configuration >   JeffP1 Stopped 2   part 1 Stopped 1   Part 10 Stopped 1	ame Astaus Processors Memory Processor Network Astaus   Ast_Test Partition Details 1 4.0   DanDPM1 Daily Grouping 0 %   DanDPM2 Recovery Operating System Messages 0 %   DanDPM3 Configuration Start 0 %   JeffP1 Stopped 2 2.0   part 1 Stopped 1 4.0

Figure 5. Methods of starting a new partition

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

• Another way to access the Start task is select the new partition, expand the Daily category in the Tasks area, and click **Start**.

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

Start - ASL_Test					
		25%			
	B · 🕇 · 🗃	Actions -	Search	$\rightarrow \rightarrow $	
Partition	▲ System	Progress	Details		
ASL_Test	S202B	25%	Configuring resources	system	

Figure 6. Portion of the Start window for a new partition named ASL\_Test

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

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

Linux localhost 3.10.0 EST 2016 s390x s390x [root@localhost -]# ls	-123.20.1.el7_0.zfpc5_0.33.s390x #1 SMP Mon Feb 8 12:13:30 s390x GNU/Linux cpu
Architecture:	33 bit 64 bit
CPU op-mode(s):	J2-Dit, 04-Dit Big Endion
Byte Order:	Big sholan
CPU(B):	
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
Lld cache:	128K
Lli cache:	96K
L2d cache:	2048K
L2i cache:	2048K
[root@localhost ~]# 1s	14 Shu

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

#### Results

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

# Part 2. Basic concepts and terms for Linux administrators

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

Topics covered in this part are:

- Chapter 3, "Partitions: Virtual images of a mainframe system," on page 11
- Chapter 4, "Adapters: Connections to networks, storage, and more," on page 17
- Chapter 5, "Operating systems and hypervisors," on page 21

## Chapter 3. Partitions: Virtual images of a mainframe system

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

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

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

- Various Linux distributions, which are listed on the IBM tested platforms page for Linux environments at http://www.ibm.com/systems/z/os/linux/resources/testedplatforms.html.
- KVM for IBM z Systems, which is a virtualization hypervisor based on Open Source technology, on which you can run multiple Linux images.
- z/VM, which is an IBM virtualization hypervisor on which you can run multiple Linux images.

DPM also supports IBM 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

I mechanisms, and other features that are specifically designed for simplifying the installation of

appliances, and for securely hosting them.

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



Figure 8. Partitions configured on a DPM-enabled system

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

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

## Partition properties and configuration settings

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

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

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

Т

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 alphanumerics, 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

I

|

I

Starting with DPM Release 2.1, 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.

1	Linux	or KVM
		In this type of partition, you can install and run one of the following: a Linux on z Systems distribution or KVM for IBM z Systems.
 	z/VM	In this type of partition, you can install and run $z/VM$ as a hypervisor for multiple Linux guests.
   	Secure	<b>Service Container</b> This type of partition is an IBM Secure Service Container, in which you can run only specific firmware or software appliances that the Secure Service Container supports.
	If you see Tab	have already created partitions for the z/VM hypervisor with an earlier release of DPM, le 6 on page 67 to consider the suggested migration action.

#### Processors

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

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

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

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

#### Memory

I

T

Each partition on a DPM-enabled system has exclusive use of a user-defined portion of the total amount of entitled memory that is installed on the system. Entitled memory is the amount of memory that is licensed for use, which might be less than the total amount of memory that is installed on the system. The amount of memory that a specific partition requires depends on the

storage limits of the operating system that will run on it, on the storage requirements of the applications that run on the operating system, and on the size of the I/O configuration.

When you define the amount of memory to be assigned, or allocated, to a specific partition, you specify an initial amount of memory, and a maximum amount that must be equal to or greater than the initial amount. The partition receives its initial amount when it is started. If the maximum amount of memory is greater than the initial amount, you can add memory up to this maximum to the active partition, without stopping and restarting it.

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

#### Network

1

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 and IBM HiperSockets<sup>™</sup>. For more information about network adapters, see Chapter 4, "Adapters: Connections to networks, storage, and more," on page 17.

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

#### Storage

Host bus adapters (HBAs) provide a partition with access to external storage area networks (SANs) and devices that are connected to a system. Each HBA represents a unique connection between the partition and a physical FICON channel that is configured on the system. You need to define an HBA for each storage area network that is required for the applications that run on this partition.

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

#### Accelerators

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

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

#### Cryptos

The term *cryptos* is a commonly used abbreviation for adapters that provide cryptographic processing functions. 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 KVM** or **z/VM**, you can specify the boot option through which DPM locates and installs the executables for the hypervisor or operating system to be run on the partition. You can choose one of several different options, including booting from a storage device, network server, FTP server, and Hardware Management Console removable media.

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

### 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

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

You can create as many partition definitions as you want, but only a specific number of partitions can be active at any given time. The system limit determines the maximum number of concurrently active partitions; for example, 85 partitions can be active on a z13, while 40 can be active on a z13s<sup>TM</sup>. 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 25.

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

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

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

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

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

#### Network

In a z Systems or LinuxONE system, several types of network adapters enable communication through different networking transport protocols. These network adapters are:

- Open Systems Adapter-Express (OSA-Express) features, which provide direct, industry-standard Ethernet LAN connectivity through various operational modes and protocols. OSA adapters can provide connectivity between partitions on the same system, as well as connectivity to external LANs.
- IBM HiperSockets, which provide high-speed communications between partitions within a single system, without the need for any physical cabling or external networking connections.
- The 10 Gigabit Ethernet (10 GbE) RoCE Express feature, which supports 10 Gbps Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE). This feature provides similar benefits for system-to-system communication as the HiperSockets technology provides for communication between partitions on the same system.

DPM automatically discovers OSA and RoCE adapters because they are physical cards that are installed on the system. In contrast, HiperSockets are not physical adapters; you must configure them if you want to use them on your system. To create HiperSockets on a DPM-enabled system, use the **Create HiperSockets Adapter** task, which is available through the **Actions** list on the Adapters tab of the **Manage Adapters** task.

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

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

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

#### **DPM-enabled system**



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

#### Storage

Fibre Channel connections (FICON) provide high-speed connections between systems and storage devices. In a z Systems or LinuxONE system, FICON cards operate in Fibre Channel Protocol (FCP) mode through single- or multiple-channel switches to Small Computer System Interface (SCSI) devices. Fibre Channel networks consist of servers, storage controllers, and devices as end nodes, which are interconnected by Fibre Channel switches, directors, and hubs. Switches and directors are used to build Fibre Channel networks or fabrics.

DPM automatically discovers any storage adapters configured on the system. These storage adapters 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.

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

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

- The location of the code for the hypervisor or operating system that will run on the partition.
- The storage requirements of the workload that your partition will support.

Typically, you need to define an HBA for each storage area network that is required for the applications that run on this partition. Your system planner or storage administrator can advise
you on which storage connections to use for the workload that your partition supports. Figure 10 provides a conceptual illustration of partitions that are configured to use storage adapters to access a storage area network.

#### **DPM-enabled system**



Figure 10. Partitions with HBAs configured for access to storage adapters on a DPM-enabled system

#### Accelerators

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

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

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

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

#### Cryptos

The term *cryptos* is a commonly used abbreviation for adapters that provide cryptographic processing functions. Industry Public Key Cryptography Standards (PKCS) and the IBM 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 operating systems that can run on z Systems or LinuxONE systems.

The use of the IBM cryptographic architecture is enabled through Crypto Express5S features, which provide a secure hardware and programming environment for cryptographic processes. Crypto Express features are installed in the Peripheral Component Interconnect Express (PCIe) I/O drawer. 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 IBM 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.

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

Linux servers require specific device drivers to use the adapters that are defined for a partition. For more information about the device drivers through which Linux servers can use adapters, see the *Linux on z Systems 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 Systems, is available in IBM Knowledge Center at http://www.ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz\_r\_dd.html

# Chapter 5. Operating systems and hypervisors

Partitions on a DPM-enabled system can host a single hypervisor or operating system. DPM supports the Linux operating system and several types of hypervisors that can host multiple Linux images.

Starting with DPM Release 2.1, 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 or KVM

I

I

I

I

In this type of partition, you can install and run one of the following: a Linux on z Systems distribution or KVM for IBM z Systems.

**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 an IBM Secure Service Container, in which you can run only specific firmware or software appliances that the Secure Service Container supports.

If you have already created partitions for the z/VM hypervisor with an earlier release of DPM, see
 Table 6 on page 67 to consider the suggested migration action.

The following types of operating systems and hypervisors can run on partitions (other than SecureService Container partitions) on a DPM-enabled system:

- Various Linux distributions, which are listed on the IBM tested platforms page for Linux environments at http://www.ibm.com/systems/z/os/linux/resources/testedplatforms.html.
- KVM for IBM z Systems, which is a virtualization hypervisor based on Open Source technology, on which you can run multiple Linux images.
- z/VM, which is an IBM virtualization hypervisor on which you can run multiple Linux images.

Figure 11 on page 22 shows a sample configuration of Linux servers and the types of workloads that they might support.



Figure 11. 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 KVM for IBM z hypervisor that hosts multiple Linux web servers.
- Partition E is configured to run a z/VM hypervisor that hosts multiple Linux images running various applications.

## 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. For partitions with a type of Linux or KVM and z/VM, you can choose one of several different options, including booting from a storage device, network server, and Hardware Management Console (HMC) removable media. (DPM Ι automatically sets the boot option for the first-time start of Secure Service Container partitions.)

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.

Regardless of the task that you use, the Boot section displays a "Boot from" list of the following options. For more detail about each option, see step 9 on page 39 in "Creating a new partition" on page 33, or use the online help for the **New Partition** task or the **Partition Details** task.

Select this option if you want to start a partition without a hypervisor or operating system. None 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 or KVM and z/VM.

T T

T

#### Secure Service Container

This boot option is the default for a Secure Service Container partition. This boot option cannot be changed unless you first change the partition type.

#### Storage device (SAN)

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

To specify this option, you need to provide the 64-bit worldwide port number (WWPN) of the boot image, and the 64-bit hexadecimal logical unit number (LUN) that corresponds to the location of the boot image.

#### Network server (PXE)

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

#### FTP server

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

#### Hardware Management Console removable media

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

#### ISO image

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

#### Working with an installed operating system or hypervisor

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

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

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

- For Linux distributions, see the Linux on z Systems topics in IBM Knowledge Center, at this URL: https://www.ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz\_r\_lib.html
- For the KVM for IBM z hypervisor, see KVM for IBM z Systems topics in IBM Knowledge Center, at this URL: https://www.ibm.com/support/knowledgecenter/SSNW54
- For the z/VM hypervisor, see the z/VM topics in IBM Knowledge Center, at this URL: https://www.ibm.com/support/knowledgecenter/SSB27U

**Note:** DPM supports z/VM with the limitations listed in "Requirements for the hypervisor or operating system" on page 64, so not all z/VM functions described in IBM Knowledge Center topics are available for use.

# 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 27
- Chapter 7, "Creating and starting a new partition," on page 31
- Chapter 8, "Dynamically modifying the resources of an active partition," on page 49
- Chapter 9, "Summary of tasks for managing systems, adapters, and partitions," on page 57

# Chapter 6. Planning checklist for creating a partition

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

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

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

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

Table 1 on page 28 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 33.

#### Table 1. Planning checklist for defining your new partition

	0	Planning item or decision	Your notes:
     		<b>Partition characteristics:</b> 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
		<b>Status:</b> 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 <b>New Partition</b> task. This section is available only in advanced task mode.
		<ul> <li>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:</li> <li>Partition access</li> <li>Counter facility authorization</li> <li>Sampling facility authorization</li> </ul>	For a complete list and explanation of available controls, see the online help for the Controls section of the <b>New Partition</b> task. This section is available only in advanced task mode.
		<b>Processors:</b> 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 <b>Partition Details</b> task after starting the partition and observing the workload performance.	Amount: Type: IFL   CP Shared: Yes   No
 		If you want to enable simultaneous multithreading for this partition, you must select the IFL processor type.	
		<b>Memory:</b> 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 <b>Partition Details</b> task to make adjustments, if necessary.	Amount:        MB   GB   TB         Maximum        MB   GB   TB
 		If you are creating a Secure Service Container partition, you must specify a minimum amount of 4096 MB (4 GB).	
1		<b>Network connections:</b> 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.	NIC name: NIC device ID: Adapter type: HiperSocket   OSA   RoCE NIC name: NIC device ID: Adapter type: HiperSocket   OSA   RoCE
   		you must specify at least one NIC for communication with the Secure Service Container web interface.	

0	Planning item or decision	Your notes:
	<b>Storage:</b> Create a host bus adapter (HBA) for each storage area network (SAN) to which the partition needs access. For each HBA, provide a name and optional device number, and select one storage adapter from the list of available adapters. For availability, select at least two storage adapters of the same type, and create an HBA for each one.	HBA name: HBA device ID: HBA name: HBA device ID:
	If you are creating a Secure Service Container partition, the storage requirements depend on the type of appliance that you plan to install in the partition. <b>Note:</b> After you create HBAs and save the partition definition, you need to use the <b>Export WWPNs</b> task before you start the partition. Your storage administrator needs this list of WWPNs to perform zoning and LUN masking so your partition can successfully access the storage devices.	
	<b>Operating system / Hypervisor:</b> For a list of supported operating systems and hypervisors, see Chapter 5, "Operating systems and hypervisors," on page 21. When you know which operating system or hypervisor that you want to run in your partition, determine the location where the executables reside, and the name of the ISO image file or .INS file. The location determines which boot option you select. Only one boot option, ISO image, requires the ISO file name.	Operating system   Hypervisor: Location: SAN   Net boot   FTP   HMC   local drive ISO file name: .INS file name:
	<b>Boot option:</b> 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 <b>None</b> , you can specify a time limit (in seconds) for the load operation. With the <b>Storage device (SAN)</b> option, you also can specify additional information, such as OS load parameters. DPM automatically sets the boot option for the first-time start of a Secure Service Container partition. For a list and descriptions of the boot options, see Chapter 5, "Operating systems and hypervisors," on page	Boot loader time-out value:

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

| | |

| |

# Chapter 7. Creating and starting a new partition

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

- "Selecting which New Partition task mode to use"
- "Creating a new partition" on page 33
- "Starting a partition and its operating system or hypervisor" on page 44

## 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

ł

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 has authorized to this task through customization controls in the **User Management** task.

## Comparing the task modes

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

- A dash (—) indicates a property that you cannot edit in the basic task mode. DPM either provides default values or automatically generates values for these properties.
- A checkmark ( ) 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	0	0
Partition short name and ID	—	0
Partition type	0	0
Reserved resources	—	0
Acceptable partition status values	—	0

Partition property	Basic mode	Advanced mode
Controls: • Partition access • Counter facility authorization • Sampling facility authorization		<ul> <li>(editing requires SYSPROG or SERVICE user ID)</li> </ul>
Shared processors	0	0
Dedicated processors	—	0
Processing weights and capping	—	0
Memory (initial allocation)	0	0
Maximum memory (dynamic allocation)	0	0
Network interface cards (NICs)	0	0
Storage (HBAs)	0	0
Accelerators (virtual functions)	<ul><li>(if installed on system)</li></ul>	<ul> <li>(if installed on system)</li> </ul>
Cryptos (security)	<ul><li>(if installed on system)</li></ul>	<ul> <li>(if installed on system)</li> </ul>
Boot options	0	0

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

## Switching between task modes

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

#### If you start in basic mode and switch to advanced mode

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

#### If you start in advanced mode and switch to basic mode

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

# Creating a new partition

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

## Before you begin

The New Partition task offers two modes through which you can create a partition: basic and advanced. This procedure provides instructions only for the basic mode of the **New Partition** task. Because the advanced mode is similar, however, you can use these instructions for the advanced mode as well. Note that some pages, or sections, of the advanced mode might have slightly different names and additional content, compared to the basic mode. If you want to review the differences between the two modes, see "Selecting which New Partition task mode to use" on page 31.

- 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 has 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 27 to gather the information that you need to select or fill in values in the New Partition task.

## About this task

L L

> 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 IBM z Systems Secure Service Container User's Guide, SC28-6971, which is available 1

on http://www.ibm.com/servers/resourcelink. L

## Procedure

- 1. Open the New Partition task. You can access this task from the main HMC page by selecting the Systems Management node, by selecting a specific DPM-enabled system, or by selecting the task in the Tasks index. For example:
  - a. Select a DPM-enabled system listed under the Systems Management node.
  - b. From the Configuration task group, click the link for the New Partition task. The New Partition window opens, with an overlay that highlights key task controls on the window.

c. Click the Okay, got it button to remove the overlay. The Welcome page is displayed.

Welcome Name Processors Memory Network Storage Accelerators Cryptos Boot	<ul> <li>Welcome to the New Partition wizard. Use this wizard to create a DPM partition.</li> <li>This wizard guides you through the following tasks: <ul> <li>Naming and describing the partition</li> <li>Assigning processors and memory</li> <li>Providing access to I/O, including networks, storage, accelerators, and cryptos</li> <li>Configuring the boot device and parameters for loading the partition</li> </ul> </li> </ul>
Advanced	Show this welcome page next time Back Next Finish Cancel Help

Figure 12. Welcome page of the New Partition task in basic mode

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

#### Linux or KVM

1

1

1

1

1

In this type of partition, you can install and run one of the following: a Linux on z Systems distribution or KVM for IBM z Systems.

**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 an IBM Secure Service Container, in which you can run only specific firmware or software appliances that the Secure Service Container supports.

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

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

📗 New Par	tition - S202B		
Welcome     Name     Processors     Memory	Define the type and number of virtual processors for the partition. The virtual processors will the selected type. Processor type:  Processor type: Contral Processor (CP) Integrated Facility for Linux (IFL)	I be allocated from the	s shared physical processors of
Network Storage Accelerators Cryptos Boot Summary	Processors:	Processors	Entitled ⑦ Dedicated ASL_Test Shared
Related Tasks Monitor System		Shared Processors Virtual/Physical:	16.67% (*) ASL_Test DanDPM1 DanDPM3 DanDPM2

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

|

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

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

New Pa	rtition - S202B	
Welcome Name	Define the memory to allocate to the partition when active. In addition, maximum memory can memory, up to that limit, without requiring a restart of the partition.	be set to allow for expansion of the partition's
Memory	Memory (GB):	Installed Memory ( GB ) :
Storage	0.5 150.4 300.3 450.1 600 4	700
Accelerators	Maximum memory (GB):	600 ASI Test
Cryptos		500ASL_Test Maximum
Boot		400 - Allocated ①
Summary		300 - Entitled (1)
		200 - Installed
		100
		0

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

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

If you are creating a partition only to familiarize yourself with the process, you do not need to create any NICs unless you want to do something more than simply start the partition when you have finished creating it. In this case, skip to step 6 on page 38 of these instructions. Otherwise, you need to create a NIC for each network connection that is required for the operating system or hypervisor that runs on this partition, or for the applications that the operating system or hypervisor supports.

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

Т

T

Т

1

Welcome     Name     Processors     Memory	Create NICs for the partition, to enable access to networks, by clicking the New action in the table.								
Network Storage	• 🗩 🖻 E	8 -	9 - 8	Actions ~				Search	*
Accelerators Cryptos	Name	•	Туре	Device Number	Adapter Name	Adapter Port	Card Type	Description	
Boot     Summary					No items to	display			

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

I

L

1

|

T

L

I

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

- **a.** From the Actions list in the NICs table, select **New** to open the New Network Interface Card window.
- b. Enter a unique, meaningful name and, optionally, a description of the new NIC.
- **c.** If you are creating a Secure Service Container partition, supply values for the additional information about the network connection that is required to access the Secure Service Container web interface. This information includes an optional, virtual local area network (VLAN) identifier, the required IP address and type, and a mask / prefix.
- If you need more detailed descriptions as you provide these configuration values, see the online help.
  - d. Review the entries in the Adapter Ports and Switches table to determine which network adapters are configured on the system.
    - 1) Check the percentages listed in the Uplink Utilization and Adapter NIC Allocation columns. If the percentage in either column is high (for example, 90%) for a specific port or switch, consider selecting a different port or switch on the same network.
    - 2) Look for a warning icon next to the name in the Adapter Name column; if the warning icon is displayed for a specific port or switch, select a different one on the same network.
    - **3)** Select one port or switch by clicking the radio button in the Select column. Note that, if you select an OSA-Express adapter port other than port 0, you need to manually specify the relative port number through a Linux qeth device driver command, before entering the Linux command to bring the device online. Instructions for specifying a port number are provided at the end of the procedure in "Starting a partition and its operating system or hypervisor" on page 44.
  - e. Click OK to create the new NIC and close the New Network Interface Card window.
  - f. Check the entry for the new NIC that is displayed in the NICs table on the Network page. Change the device number if your company uses a specific numbering convention for its networks.
  - **g**. If you are creating a Secure Service Container partition, provide the network settings that are displayed after the NICs table. These settings include a host name, default gateway, and DNS servers.

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

- h. Repeat the preceding steps, as necessary, to create a new NIC for each network connection that your new partition requires.
- i. When you have finished, click Next to navigate to the next page in the task.
- 6. Use the Storage page to define host bus adapters (HBAs) to enable the new partition to access specific storage networks and devices.

If you are creating a partition only to familiarize yourself with the process, you do not need to create any HBAs unless you want to do something more than simply start the partition when you have finished creating it. In this case, skip to step 7 on page 39 of these instructions. Otherwise, you need to create an HBA for each storage area network that is required for the applications that run on this partition.

- For availability, select at least two storage adapters of the same type, and create an HBA for each one.
- If you are creating a Secure Service Container partition, the storage requirements depend on the type of appliance that you plan to install in the partition.
  - For a software appliance, define at least one HBA to access the storage device on which the appliance installation image resides.
  - For a firmware appliance, you are not required to define an HBA because access to a storage device is not required to install a firmware appliance.

Welcome Name Processors Memory	Create HBAs for th HBAs	e partit	ion, to enabl	le access to sto	rage devices by clickin	g the New action in th	e table.		
Network Storage Accelerators	⊕ ⊡ ₫	- 1	<b>Q</b> - P	Actions -				Search	$\Rightarrow$
	Name	•	Туре	Device Number	Adapter Name	Card Type	Description		
* Summary					No items	to display			

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

When you first use the **New Partition** task, the Storage display contains an empty HBAs table. To create an HBA, complete the following steps.

- a. From the Actions list in the HBAs table, select New to open the New Host Bus Adapter window.
- b. Enter a unique, meaningful name and, optionally, a description of the new HBA.
- **c.** Review the entries in the Adapter Ports table to determine which storage adapters are configured on the system.
  - 1) Check the percentage listed in the Adapter HBA Allocation column. If the percentage is high (for example, 90%) for a specific port, consider selecting a different port.
  - 2) Look for a warning icon next to the name in the Adapter Name column; if the warning icon is displayed for a specific port, select a different one.

Т

T

Т

Т

Т

1

Т

3) Select one port by clicking the radio button in the Select column.

- d. Click OK to create the new HBA and close the New Host Bus Adapter window.
- **e**. Check the entry for the new HBA that is displayed in the HBAs table on the Storage page. If your company uses a specific numbering convention for its storage networks, you can change a system-generated device number by selecting the **Details** action and editing the HBA device number.
- f. Repeat the preceding steps, as necessary, to create a new HBA for each storage area network that your new partition requires.
- g. When you have finished, click **Next** to navigate to the next page in the task.

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

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

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

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

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

Crypto features are optional and, therefore, might not be installed on the system. If none are installed, the Cryptos page is disabled, and you can skip to step 9 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. Otherwise, use instructions in the online help to enable your partition to use cryptographic adapters.

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

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

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

I If you have selected the partition type **Secure Service Container**, note that option set in the "Boot

- from" menu is also **Secure Service Container**. This boot option cannot be changed unless you first change the partition type.
- To define a boot option for other types of partitions, complete the following steps.

Т

L

- 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.

	52020	
Welcome Define the	boot type, device, and parameters for the partition	
Processors Boot from	None	~
Memory	Storage Device(SAN)	
Network	Network Server(PXE)	
Storage	FTP server	
Accelerators	Hardware Management Console removable media	
Cryptos	ISO Image	
Boot		

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 or KVM** and **z/VM**.

#### Storage device (SAN)

T

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

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

Target WWPN: 50:0a:09:85:87:09:68:ad or 500a0985870968 (hexadecimal) Target LUN: 4021400000000000

#### Network server (PXE)

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

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

#### **FTP** server

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

#### Host name

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

#### User name

Enter the user name on the target FTP server.

#### Password

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

.INS file

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

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

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

#### Hardware Management Console removable media

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

When you select this option:

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

#### **ISO** image

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

When you select this option:

- Click Browse to find the ISO image file on your workstation file system. You cannot select an ISO image from an HMC media drive. As soon as you select an ISO image file, DPM starts to upload the file, and displays a progress indicator for the upload operation.
- 2) After the upload operation completes, click Browse to search the ISO image file for the .INS file that you want to use. Any .INS files found are displayed in a separate window. Select only one .INS file and click OK to close the Browse ISO Image window.
- c. When you have finished, click Next to navigate to the next page in the task.
- 10. Use the Summary page to review the properties for the new partition.

• Welcome								*	
* Name	Verify the informatio	n and then click P	inish to create	the partition.					
<ul> <li>Processors</li> <li>Memory</li> </ul>	Name: Description:	ASL_	Test					10	
Network Storage	Processor type: Processors:	Centra 1	al Processor (4	CP)					
Cryptos • Boot	Memory (GB): Maximum Memory (	4 GB): 4							
+ Summary	NICs								
	Name	Туре	Device Number	Adapter Name	Adapter Port	Card Type	Description		
	nic02	OSA	0001	OSD 0110 A138-06	0	OSA-Express4s 10Gb			

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

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

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

## Results

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



Figure 19. Validation window for the New Partition task

## What to do next

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

#### **Export WWPNs**

Opens the **Export WWPNs** task, through which you can obtain a list of worldwide port names of the host bus adapters that are defined for a partition.

If you defined HBAs for this partition, or have specified the **Storage device (SAN)** boot option, use the **Export WWPNs** task before you start the partition. Your storage administrator needs this list of WWPNs to perform zoning and LUN masking so your partition can successfully access the storage devices.

#### 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 new partition with a type of Linux or
 KVM or z/VM, and its operating system or hypervisor.

## Before you begin

- For partitions with a type of **Secure Service Container**, see the *IBM z Systems Secure Service Container User's Guide*, SC28-6971, 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 has 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, "Dynamic Partition Manager task roles," on page 83.

## Procedure

Т

- 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**. Click 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 on this partition, or to upload the required files to initialize the hypervisor or operating system when the partition itself is started.
  - b. Click 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, click the double-arrow icon in the Name field to display the cascading task menu.
  - b. Expand the **Daily** group, and click **Start**.

A new window opens to display the progress of the start operation. Figure 20 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.

Start - ASL_Test									
		25%							
	B - 7 -	Actions 👻	Search						
Partition	▲ System	Progress	Details						
ASL_Test	S202B	25%	Configuring system resources						

Figure 20. Portion of the Start window for a new partition named ASL\_Test

## Results

When the start process completes, the Details column contains a message that indicates the results of the start operation. If the start operation was successful, the partition status is Active, and DPM has started the operating system or hypervisor. At this point, you can log in to the operating system or hypervisor and verify that DPM has brought online the system resources (processors, memory, and adapters) that are specified in the partition definition.

## 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.

#### 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	1
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/5390
BogoMIPS:	7936.00
Hypervisor:	PR/SM
Hypervisor vendor:	IBM
Virtualization type:	full
Dispatching mode:	horizontal
Lld cache:	128K
Lli cache:	96K
L2d cache:	2048K
L2i cache:	2048K
[root@lbskvm4 ~]#	

Figure 21. Sample displays resulting from the Iscpu 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 ~]# <mark>lsmem</mark> Address Range	Size (MB)	State	Removable	Device
0x000000000000000000000000000000000000	8192	online	no	0
	999424	online	yes	1-122
	40960	online	no	123-127
	2154496	offline	-	128-390

Figure 22. Sample displays resulting from the Ismem 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 [root@lk Device r	o hostl oskvm4 ~]# name	lsqeth	: en	ccw0.0	.000	1			
	card type		: 05	0 1000					
	cdev0		: 0.	0.0001					
	cdev1		: 0.	9.0002					
	cdev2		: 0.	0.0003					
	chpid		: 13						
	online		: 1						
	portname		: DUI	MMY					
	portno		: 0						
	state		: UP	(LAN	ONLI	NE)			
	priority_q	ueueing	: al	ways q	ueue	Θ			
	buffer_cou	nt	: 12	8					
	layer2		: 1						
	isolation		: no	ne					
[root@]h	skym4 ~1#	Iscss							
Device	Subchan.	Devivoe	CU TVD	e Use	PIM	РАМ	POM	CHPIDs	
0.0.0001	0.0.0000	1732/01	1731/0	1 yes	80	80	ff	13000000	00000000
0.0.0002	0.0.0001	1732/01	1731/0	1 yes	80	80	ff	13000000	00000000
0.0.0003	0.0.0002	1732/01	1731/0	1 yes	80	80	ff	13000000	00000000
0.0.0000	0.0.0003	1732/03	1731/0	3	80	80	ff	39000000	000000000
0.0.0005	0.0.0004	1732/03	1731/0	3 yes	80	80	ff	2a000000	00000000
0.0.0004	0.0.0005	1732/03	1731/0	3 yes	80	80	ff	2e000000	00000000
0.0.0007	0.0.0006	1732/03	1731/0	3	80	80	ff	01000000	00000000
[root@lb	oskvm4 ~]#								

Figure 23. Sample displays resulting from the lsqeth and lscss commands

[root@ CHPID	lbskvm Vary	4 ~]# Cfg.	lschp 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@	lbskvm	4 ~]#				

Figure 24. Sample displays resulting from the lschp command

#### Specifying the relative port number of an OSA device

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

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

#### Finding additional information about operating system or hypervisor commands

- For more information about using Linux commands to work with partition resources and adapters, see the *Linux on z Systems 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 Systems, is available in IBM Knowledge Center at http://www.ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz\_r\_dd.html
- For information about using KVM for IBM z Systems commands to work with partition resources and adapters, see *System Administration* for the KVM for IBM z version that you are using. This information is available in IBM Knowledge Center at http://www.ibm.com/ support/knowledgecenter/SSNW54
- 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. This information is available in IBM Knowledge Center at http://www.ibm.com/support/ knowledgecenter/SSB27U

# 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.

This topic provides instructions through the following examples:

• "Example: Adding more memory resources"

- "Example: Adding more processor resources" on page 51
- "Example: Adding a new network device" on page 53

#### Example: Adding more memory resources

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

[root@lbskvm4 ~]# lsmem Address Range 	Size (MB)	State	Removable	Device
9x000000000000000000000000000000000000	8192 999424 40960 2154496	online online online offline	no yes no -	0 1-122 123-127 128-390
Hemory device size : 8192 MB Hemory block size : 8192 MB Fotal online memory : 1048576 MB Fotal offline memory: 2154496 MB [root@lbskvm4 ~]#				

Figure 25. Sample display for the Ismem 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.

3 Partition	Details - KVMCloud6	
General	— - Memory	
Status	Memory (GB):	Installed Memory ( GB ) :
Controls	2064	E 000
Processors	0.5 1024.4 2048.3 3072.1 4096 2,000	5,000
Memory	Maximum memory (GB): K3	4,000 - KVMCloud6
Notwork	0.5 1024 2 2048 3 3072 1 4096 3,128 🗘	3,000 KVMCloud6 Maximum
NEIWOIK		Allocated (?)
Storage		2,000 Entilled (?
Accelerators		1,000 - Instaned
Cryptos		o -

Figure 26. Sample screen of the Memory section in the Partition Details window

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

[root@lbskvm4 /]# time				
real 0m0.691s user 0m0.157s sys 0m0.553s [root@lbskvm4 /]# lsmem Address Range	Size (MB)	State	Removable	Device
Ox000000000000000000000000000000000000	8192 999424 40960 16384 2138112	online online online online offline	no yes no yes	0 1-122 123-127 128-129 130-390

Figure 27. Sample display for the chmem command

## 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.

<pre>[root@lbskvm4 /]# lscpu</pre>	1
Architecture:	s390x
CPU op-mode(s):	32-bit, 64-bit
Byte Order:	Big Endian
CPU(s):	8
On-line CPU(s) list:	0-7
<pre>[hread(s) per core:</pre>	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 28. Sample display for the Iscpu 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.

Partitio	n Details - KVMCloud6
General	- · Processors
Status Controls	Processor type: Central Processor (CP) Integrated Facility for Linux (IFL) Processor mode: Shared Dedicated
Processors Memory	* Processors: Processors Processors 20 Entitled ③ Dedicated ③
Network Storage	10 - KVMCloud6 Shared ③
Accelerators	0

Figure 29. 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.

<pre>[root@lbskvm4 /]# lscpu</pre>	1
Architecture:	s390x
CPU op-mode(s):	32-bit, 64-bit
Byte Order:	Big Endian
CPU(s):	12
On-line CPU(s) list:	0-7
<pre>Dff-line CPU(s) list:</pre>	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
Lli cache:	96K
L2d cache:	2048K
L2i cache:	2048K

Figure 30. Sample display for the Iscpu 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 9 enabled
CPU 10 enabled
[root@lbskvm4 /]#
```

Figure 31. Sample display for the Iscpu command showing configured processor resources

## Example: Adding a new network device

To add a new network device, you follow a similar process as for the previous examples. In this example, the Linux system currently has seven network devices configured, as shown in the following display.

[root@lbs	skvm4 ~]#	lscss							
Device	Subchan.	DevType	СU Туре	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	000000000
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	000000000
[rootalb:	skvm4 ~1#								

Figure 32. Sample display for the Iscss command showing current devices

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

General		Network					
Status	NICs						
Controls	÷	⊝ ₽ -	-	Actions 👻			
Processors			Туре	Adapter Name	Adapter	Device	Card Type
Memory		WNIC			Port	Number	
Network		MGMTOSA1	OSA	OSD 017C A13B-38	0	0001	OSA-Express5
Storage							100000036-1
Assolarators							

Figure 33. Sample screen of the Network section with the New NIC icon highlighted

2. Select the adapter for the new network connection.

#### New Network Interface Card

Name: DATAOSA	A						
Device Number				Ξ.			
	Actions 👻					Search	→ →
Adapter Name 1 🔺	Adapter 2 ▲ Port	Card Type	Uplink Utilization ⑦	Adapter NIC Allocation ⑦	Location	Description	
) OSD 017C A13B-38	0	OSA-Express5s 1000Base-T	0%	0%	A13B-D138J.01- D238J.01		
) OSD 017C A13B-38	1	OSA-Express5s 1000Base-T	0%	0%	A13B-D138J.01- D238J.01		
OSD 0188 A06B-03	0	OSA-Express5s 1000Base-T	0%	0%	A06B-D103J.01- D203J.01		
OSD 0188 A06B-03	1	OSA-Express5s 1000Base-T	0%	0%	A06B-D103J.01- D203J.01		

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

Figure 34. Sample screen of the New Network Interface Card window

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

Bartitio	n Details - I	(VMClou	ıd6							
General Status Controls	- Vetwork NICs									
Processors	Name	▲ Type	Adapter Name	Adapter Port	Device Number	Card Type				
Network	DATAOSA	OSA	OSD 0188 A06B-03	0	8000	OSA-Express5s 1000Base-T				
Accelerators		A1 OSA	OSD 017C A13B-38	0	0001	OSA-Express5s 1000Base-T				

Figure 35. Sample screen of the Network section with the newly added NIC

- 4. Click Apply to save your changes and close the Partition Details window.
- 5. Then, through the Linux **znetconf** command, define the new devices 0.0.0008 through 0.0.000a.
| [root@lbskvm4 /]# znetconf<br>Scanning for network device<br>Device IDs  | -u<br>es<br>Туре                       | Card ⊤ype   | CHPID             | Drv.         |                                |                  |
|--|--|---|-------------------|--------------|--------------------------------|------------------|
| 0.0.0008,0.0.0009.0.0.000a<br>[root@lbskvm4 /]# znetconf<br>Scanning for network device<br>Successfully configured dev | 1731/01<br>-a 0.0.0<br>es<br>/ice 0.0. | OSA (ODIO)<br>0008 -o layer2=:<br>.0008 (enccw0.0 | 14<br>1<br>.0008) | qeth         |                                |                  |
| Device IDs   | Туре                                   | Card Type   | CHPID             | Drv.         | Name                           | State            |
| 0.0.0001,0.0.0002,0.0.0003<br>0.0.0008,0.0.0009,0.0.000a<br>[root@lbskvm4 /]#  | 1731/01<br>1731/01                     | OSD_1000<br>OSD_1000                              | 13<br>14          | qeth<br>qeth | enccw0.0.0001<br>enccw0.0.0008 | online<br>online |

Figure 36. Sample display for the znetconf command showing current devices

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

Issuing an **lsqeth** command displays the successfully defined network device, which can be defined to the Linux TCP/IP stack or passed to a virtual switch. Figure 37 on page 56 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
lsolation	: none
Device name	: enccw0.0.0008
lsolation	: none
Device name	: enccw0.0.0
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 ON
priority_queueing	: always que
buffer count	: 64

Figure 37. Sample display for the Isqeth command showing details for devices

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

DPM tasks are available starting with the Hardware Management Console (HMC) / Support Element (SE) Version 2.13.1. Chapter 9, "Summary of tasks for managing systems, adapters, and partitions" provides an alphabetical summary of DPM tasks and other tasks that you might use to work with a DPM-enabled system, partitions, 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, "Dynamic Partition Manager task roles," on page 83.
- 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 *z Systems Hardware Management Console Web Services API*, SC27-2634, which is available through the Publications link on IBM Resource Link at http://www.ibm.com/servers/resourcelink

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

Task name	Icon	On HMC or SE	Description
Customize Scheduled Operations		Both	Use the <b>Customize Scheduled Operations</b> 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		НМС	Use the <b>Delete Partition</b> task to delete the definition associated with one or more selected partitions on aDPM-enabled system.
Disable Dynamic Partition Manager	$\bigcirc$	SE only	Use the <b>Disable Dynamic Partition Manager</b> task to disable DPM, remove all partitions, and unconfigure all adapters from the system.
Dump (Partition)	,		Use the <b>Dump</b> task to initiate a partition dump by booting a stand-alone dump program on a DPM-enabled system.
Enable Dynamic Partition Manager		SE only	Use the <b>Enable Dynamic Partition Manager</b> task to enable DPM for the system and identify the OSA-Express5s 1000 Base-T adapters for z System Management (OSM).
Getting Started with Dynamic Partition Manager		НМС	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.

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
Integrated 3270 Console		HMC	Use the <b>Integrated 3270 Console</b> task to log on to a z/VM hypervisor that is running in a DPM partition.
Integrated ASCII Console		HMC	Use the <b>Integrated ASCII Console</b> task to log on to a Linux operating system.
Manage Adapters	B	HMC	Use the <b>Manage Adapters</b> task to view and customize the adapters and devices of an DPM-enabled system.
Adapter Details	~@}	HMC	Use the <b>Adapter Details</b> task to view or modify the adapter settings of the selected adapter.
Create Hipersockets Adapter	ရီဝ	НМС	Use the <b>Create Hipersockets Adapter</b> task to create an adapter and switch for IBM HiperSockets, 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 <b>Delete Hipersockets Adapter</b> task to delete only one selected HiperSockets adapter.
Export WWPNs	$\downarrow$	НМС	Use the <b>Export WWPNs</b> 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 <b>Reassign Channel Path IDs</b> task to change the channel path IDs that are assigned to DPM adapters.
Reassign Devices	G	HMC	Use the <b>Reassign Devices</b> task to change the adapter, port, or switch for one or more devices.
Manage Processor Sharing		НМС	Use the <b>Manage Processor Sharing</b> 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.
<b>Monitor</b> tab under the Systems Management node	۲		Use the <b>Monitor</b> tab to view the overall system monitoring data for the various components that make up a DPM-enabled system.
Monitor System Events		НМС	Use the <b>Monitor System Events</b> 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		НМС	Use the New Partition task to create a new partition on an DPM system. The <b>New Partition</b> 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 31.
Operating System Messages	Q	НМС	Use the <b>Operating System Messages</b> 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.

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
Partition Details	ŝ	Both (view-only mode on SE)	Use the <b>Partition Details</b> 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 <b>Partition Details</b> task.
Start	►	Both	Use the <b>Start</b> task to start a single DPM-enabled system, or to start one or more partitions on a DPM-enabled system.
Stop	۲	Both	Use the <b>Stop</b> 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 <b>System Details</b> task to view and manage properties of the selected DPM-enabled system.

### Part 4. Topics for system planners

Topics in this part provide the prerequisites for enabling DPM on a mainframe system, information about supported functions, and the engineering changes (EC) or microcode control levels (MCL) for upgrading to the latest DPM version. These topics are appropriate for experienced system planners and other administrators who are familiar with mainframe systems.

Topics covered in this part are:

- Chapter 10, "Prerequisites for using Dynamic Partition Manager," on page 63
- Chapter 11, "DPM migration considerations," on page 67
- Chapter 12, "Adapter configuration," on page 69

# Chapter 10. Prerequisites for using Dynamic Partition Manager

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

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

- An IBM z13 (z13) or IBM z13s<sup>TM</sup> (z13s<sup>TM</sup>)
- An IBM LinuxONE Emperor (Emperor) or IBM LinuxONE Rockhopper (Rockhopper)

DPM tasks are available starting with the Hardware Management Console (HMC) / Support Element (SE) Version 2.13.1.

#### **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.

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

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

### Versioning

DPM is available through specific engineering changes (EC) or microcode control levels (MCL) for the mainframe systems that support it. Table 5 on page 64 lists the DPM versions and the EC MCLs that must be applied to use any new functions or updates.

#### Important:

L

- 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.

DPM release	Machine type	HMC / SE version and required MCLs	Description
Release 1 (R1)	2964 or 2965	HMC / SE Version 2.13.1	Initial release of DPM for the z13, z13s, Emperor, and Rockhopper systems.
Release 2.0 (R2.0)	2964 or 2965	<ul> <li>HMC / SE Version 2.13.1</li> <li>with the following MCLs:</li> <li>SE-SYSTEM P00339.243 or later</li> <li>SE-I390ML P08414.044 or later</li> <li>SE-ENABLE2 P08434.001 or later</li> <li>SE-MCS P08456.019 or later</li> </ul>	<ul> <li>Support for z/VM as a hypervisor, and for the use of the Integrated 3270 Console task through which users can interact with the z/VM hypervisor. For additional details, see "Requirements for the hypervisor or operating system."</li> <li>Support for universally unique identifiers (UUIDs), which enables users to define their own device numbers for Peripheral Component Interconnect Express (PCIe) devices, when using a Linux distribution that also supports the use of UUIDs. For related information, see Chapter 11, "DPM migration considerations," on page 67.</li> </ul>
Release 2.1 (R2.1)	2964 or 2965	<ul> <li>HMC / SE Version 2.13.1</li> <li>with the following MCLs:</li> <li>HMC-SYSTEM P08462.261 and SE-SYSTEM P00339.304 or later</li> <li>To use the support for IBM Secure Service Container partitions, also install SE-FWPART P08442.025 and SE-BCINST P08458.035 or later</li> </ul>	<ul> <li>Support for displaying the number of processor threads through the Partition Details task.</li> <li>Support for assigning a partition type of z/VM, for partitions in which the z/VM operating system will run. If you have already installed and run z/VM on DPM R2.0 or DPM R1, see the z/VM migration information in Chapter 11, "DPM migration considerations," on page 67.</li> <li>Support for creating, starting, and managing Secure Service Container partitions. A Secure Service Container partition contains its own embedded operating system, security mechanisms, and other features that are specifically designed for simplifying the installation of appliances, and for securely hosting them. For information about working with Secure Service Container partitions and the appliances they support, see <i>IBM z Systems Secure Service Container User's Guide</i>, SC28-6971, which is available on http://www.ibm.com/servers/resourcelink.</li> </ul>

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

### 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 Systems:
  - SUSE Linux Enterprise Server (SLES) 11 or later.
  - Red Hat Enterprise Linux (RHEL) 6 or later.

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

- KVM for IBM z Systems V1.1.1 with PTFs. For required PTFs, see the Preventative Service Planning bucket (PSP) for the appropriate system.
- z/VM V6.4 or later, with the following limitations:
  - On a DPM-enabled system, z/VM can host only Linux on z Systems guests.
  - Because DPM does not support extended count key data (ECKD<sup>™</sup>) volumes, 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).

1

For information about installing and running a Linux distribution on a z Systems or LinuxONE server, see the Linux on z Systems topics in IBM Knowledge Center, at this URL: https://www.ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz\_r\_lib.html

### 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 z13 or Emperor, the 2964DEVICE PSP bucket
  - For a z13s<sup>™</sup> or Rockhopper, the 2965DEVICE PSP bucket

### **Chapter 11. DPM migration considerations**

This topic describes the migration considerations or actions that you might need to take when you apply the latest version of IBM Dynamic Partition Manager (DPM) to your system.

DPM is available through specific engineering changes (EC) or microcode control levels (MCL) for the mainframe systems that support it. For information that correlates DPM versions and their associated EC MCLs, see "Versioning" on page 63.

Table 6 lists the migration actions that you should consider when migrating from DPM Release 1 (R1) to Release 2 (R2).

Description of function or change	Migration action
UUID support in DPM R2 Previously, through the DPM R1 New Partition or Partition Details task, users were able to specify their own device numbers for Peripheral Component Interconnect Express (PCIe) devices, but they were not	If you defined and started partitions through DPM R1 tasks, and also defined specific device numbers for the PCIe devices that those partitions use, you must stop and restart those active partitions after migrating to DPM R2 and to an operating system that also supports UUIDs.
visible to the operating system that was running on the partition. Examples of PCIe devices include Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) Express features and IBM zEnterprise Data Compression (zEDC) features.	<b>Attention:</b> This migration action is a disruptive task. After applying the appropriate EC MCLs for DPM R2 and upgrading the operating system, complete the following steps.
With DPM R2 support for universally unique identifiers (UUIDs), however, user-supplied device numbers are guaranteed to be unique, and are visible to operating systems that also support the use of UUIDs.	<ol> <li>Use the Stop task to stop any active partitions that have user-defined device numbers for PCIe devices.</li> <li>Use the Start task to restart the stopped partitions.</li> <li>Through the Integrated ASCII Console task, enter the lspci command to verify that the operating system displays the same unique device numbers that are specified in the Partition Details task.</li> </ol>
Partition type z/VM in DPM R2	Through the partition type that is introduced with DPM R2.1, DPM can optimize the partition configuration for a specific hypervisor or operating system. To benefit from this optimization, the suggested practice is to use the appropriate partition type for the hypervisor or operating system that runs in any new or existing partition. If you have already installed and run z/VM on DPM R2.0 or DPM R1, and you install the MCLs for DPM R2.1, the partition type for previously defined z/VM partitions is displayed as <b>Linux or KVM</b> in the <b>Partition Details</b> task. Because the configuration requirements for Linux, KVM, and z/VM are similar, you do not have to change the partition type for existing z/VM partitions. However, the suggested practice is to change the partition type, you must use the <b>New Partition</b> task to create a new partition definition for each existing z/VM partition. You cannot change the partition type through the <b>Partition Details</b> task.

© Copyright IBM Corp. 2016, 2017

### 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 **Manage Adapters** task to review details about the adapters, and make any adjustments that might be necessary.

Also, if HiperSockets are required on the system, you need to define them. See "Creating a HiperSockets adapter" on page 72 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 7 provides examples of the default names that DPM assigns to installed adapters. In the table, PCHID refers to the physical channel path identifier, and Location includes jack and slot information.

Туре	Crypto type	PCHID	Card type	Location	Default name
FCP	—	0171	FICON Express 8	Z22B-D211-J.01	FCP 0171 Z22B-11
OSD	—	018C	OSA-Express5S	Z15B-D104-J.01	OSD 018C Z15B-04
RoCE	—	0184	10GbE RoCE Express®	Z15B-LG02J.01- LG02J.02	RoCE 0184 Z15B-02
zEDC	—	0188	zEDC Express	Z15B-LG03	zEDC 0188 Z15B-03
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

Table 7. Sample default names for installed adapters

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 38 on page 70.

Adap	Manaç Iters De	ge Ada	apters - Cryptos	S202B						Related Tasks
Filter:	•ြာ Network	s 📑 Stora	age 🕜 Acceler	rator 🦪 Crypto						
ŵ	•	- 🖯	- Action	15 🔻					Sean	ch →
	Name 🔺	ID	Туре	Status	State	Card Type	Location	Device Allocation	Number of Partitions	Description
	C Accel 0128 A13B-13	0128	Crypto	Z Active	Online	Crypto Express5s	A13B-LG13	0.0%	8	
	C CCA 0138 A13B-19	0138	Crypto	Not active	Standby	Crypto Express5s	A13B-LG18	0.0%	0	
	C CCA E 0144 A13B-2	Delete Hiper Create Hiper Reassign Ch	Sockets Adapter Sockets Adapter nannel Path IDs	Not active	Standby	Crypto Express5s	A13B-LG21	0.0%	0	
	C EP11 0104 A13B-02	0104	Crypto	Active	Online	Crypto Express5s	A13B-LG02	0.0%	7	

Figure 38. 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 39 on page 71 shows the adapter details for a cryptographic adapter. After you modify the name or add a description, click **Apply** to save your changes.

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

Figure 39. Sample screen of the Adapter Details window, with Name and Description fields highlighted

### Creating a HiperSockets adapter

To create a HiperSockets adapter, open the **Manage Adapters** task, and click the down arrow to display the Actions list, as shown in Figure 40. Click **Create HiperSockets Adapter** in the list.

🗉 Man	age Adapters - S	6223B						
Adapters	Devices Cryptos							Related Tasks
Filter: 🖧 Ne	etwork 📕 Storage 🕜 Acceler	ator d Crypto						
\$ <del>(</del>	Actions 👻	_					Search	
Name	Adapter Details Delete HiperSockets Adapter	itatus	State	Card Type	Location	Device Allocation	Number of Partitions	Description
) 🖻 🗖 F. 0131 A13B-	Create HiperSockets Adapter Reassign Channel Path IDs Configure Options	Not active	Reserved	FICON Express16s	A13B-D216-J.01	0%	0	
) 🗷 🗖 FC 0134 A13B-	Select All Deselect All Expand All	Not active	Reserved	FICON Express16s	A13B-D117-J.01	0%	0	
)	Clear Sort	Not active	Reserved	FICON Express16s	A13B-D217-J.01	0%	0	

Figure 40. Sample screen of the Manage Adapters task, with the expanded Actions list

The New HiperSockets Adapter window opens.

	w Hiper	Sockets Adapt	ter - S223B		
HiperSockets	adapters have a	switch that provides high-spe	ed communications betw	een partitions on a	system
using in-memo systems, inclu	ory TCP/IP conn ding Linux.	ections between and among p	partitions running several	different operating	
Name:	HS10				
Description:		8/16			
		16/24			
		32/40			
		56/64			
MTU/Frame s	ize (KB): ?	56/64 🗸	ОК	Cancel	Help

Figure 41. 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 Deta	ils - HS10	
General	General	
Network Interface Cards	* Name:	HS10
	Description:	
		in the second
	Object ID:	A56C8508-F06F-11E5-9051-9ABE94228451
	System:	S223B
	Status:	Not active
	Detailed status:	Not active
	State:	Standby
Ľ	NIC allocation: 🤨	0%
	Adapter ID:	07C0
	Adapter type:	Internal queued direct communications (HiperSockets)
	Channel path ID (CHPID): 🤅	48
	Card type:	HiperSockets
	MTU/Frame size (KB): ⑦	56/64 🗸

Figure 42. Sample screen of the new HiperSockets adapter in the Adapter Details window

### Part 5. Topics for network administrators

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

Topics covered in this part are:

• Chapter 13, "Customizing the network boot server environment," on page 77

### Chapter 13. 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 (IEFT) requests for comments (RFCs) related to networking.

Before you complete the steps in this procedure, your installation must already have set up a DHCP server and a TFTP server, according to the IEFT RFCs. These servers can reside on the same or on different systems within the network.

• 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.

### Procedure

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

allow booting; allow bootp;

b. 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*.

next-server <*ipv4 addr*>

c. Complete one of the following options.

#### Option 1

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";
}
```

2) Define a pool that includes both the dynamic address range and the boot configuration to load for all members of the class defined in the previous step. For example:

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

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

**Option 2** 

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

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

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

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

The file-root directory of the TFTP server is the top-level directory, where the TFTP server loads and stores files. It is specified by the -s *file root* option to **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=geth
            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:

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

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

```
pool {
```

```
option bootfile-name "NetworkBoot.DPM";
filename "NetworkBoot.DPM";
range 192.168.9.210 192.168.9.214;
}
host 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.

Part 6. Appendixes

### **Appendix A. Dynamic Partition Manager task roles**

Tasks and resources need to be made available or excluded based on the roles to which they are assigned. You can create your own specific task and resource roles that include specific tasks and resources; however, IBM provides default roles for your convenience. Table 9 identifies the DPM tasks along with default task roles.

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 9.

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 main UI 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. Table 8 describes the abbreviations used in Table 9.

Legend	Description
AA	Access Administrator Tasks
SP	System Programmer Tasks
OP	Operator Tasks
AOP	Advanced Operator Tasks
SER	Service Representative Tasks
X	Required role to perform a task.
0	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 8. DPM task roles mapping legend

Table 9.	DPM	task	roles	mapping
----------	-----	------	-------	---------

DPM Tasks	Task Roles				
	AA	SP	OP	AOP	SER
Enable Dynamic Partition Manager*					Х
Disable Dynamic Partition Manager*					Х
Getting Started with Dynamic Partition Manager	0	0	0	0	0
New Partition		0			0
Partition Details		0			0
Partition Details - Controls		0			0
Delete Partition		0			0
Start (start a single DPM system)		0			0
Start (start one or more DPM partitions)		0	0	0	0
Stop (stop a single DPM system)		0			0

Table 9. DPM task roles mapping (continued)

DPM Tasks Task Roles					
Stop (stop one or more DPM partitions)		0	0	0	0
Manage Processor Sharing		0			0
Manage Adapters		О	0	0	0
Adapter Details		0			0
Create HiperSockets Adapter		0			0
Delete HiperSockets Adapter		0			0
Reassign Channel Path IDs		0			0
Reassign Devices		0			0
Export WWPNs		0	0	0	0
System Details: Configure z Systems Management (OSM) Adapters*					X
Dump (Partition)		0	0	0	0

### **Appendix B. Notices**

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprise is entirely coincidental.

### Trademarks

IBM, the IBM logo, and ibm.com<sup>®</sup> are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

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

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

### **Class A Notices**

The following Class A statements apply to this IBM product. The statement for other IBM products intended for use with this product will appear in their accompanying manuals.

### Federal Communications Commission (FCC) Statement

**Note:** This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

### **Industry Canada Compliance Statement**

This Class A digital apparatus complies with Canadian ICES-003.

### Avis de conformité à la réglementation d'Industrie Canada

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

### **European Community Compliance Statement**

This product is in conformity with the protection requirements of EU Council Directive 2014/30/EU on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to European Standard EN 55032. The limits for Class A equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

European Community contact: IBM Deutschland GmbH Technical Regulations, Department M372 IBM-Allee 1, 71139 Ehningen, Germany Tele: +49 (0) 800 225 5423 or +49 (0) 180 331 3233 email: halloibm@de.ibm.com

**Warning:** This is a Class A product. In a domestic environment, this product may cause radio interference, in which case the user may be required to take adequate measures.

### VCCI Statement - Japan

この装置は、クラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。 VCCI-A

The following is a summary of the Japanese VCCI statement above:

This is a Class A product based on the standard of the VCCI Council. If this equipment is used in a domestic environment, radio interference may occur, in which case the user may be required to take corrective actions.

### Japan JIS C 61000-3-2 Compliance



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、PFC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、PFC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 А. В жилых помещениях оно может создавать радиопомехи, для снижения которых необходимы дополнительные меры
## Index

### Α

accessibility xiii contact IBM xiii features xiii adapter definition 3 dynamically add 49 specifying port number 47 assistive technologies xiii

## В

boot option network server 77 PXE 77 boot options for operating systems 22 hypervisor 22 operating system 22

# С

commands KVM for IBM z Systems 47 Linux operating system 47 z/VM hypervisor 47 concepts hypervisors 21 operating systems 21 corequisite software 65 create a partition comparison of New Partition task modes 31 overview 5 selecting New Partition task mode 31 step-by-step instructions 33 switching New Partition task modes 32

## D

DPM See also Dynamic Partition Manager feature code 63 prerequisites 63 DPM-enabled system supported types 63 Dynamic Partition Manager See also DPM advantages 4 definition 3 enabling for a system 3 summary of tasks 57 user interface 3

### Ε

echo command 47

engineering change See MCL engineering change (EC) 63

### F

feature code for Dynamic Partition Manager 63 for OSA-Express 1000BASE-T Ethernet 63

## Η

Hardware Management Console list of Dynamic Partition Manager tasks 57 hardware planning 65 HMC *See also* Hardware Management Console EC MCLs for DPM version 63 how to get started 5

## Κ

keyboard navigation xiii KVM for IBM z Systems commands 47 KVM on IBM z 21 required version 64 sample configuration 11

## L

Linux operating system 21 commands 47 Linux operating system 20 lschp 46 lscpu 45 lscss 46 lsmem 45 lsqeth 46 device drivers for adapters 20, 47 lscpu command 45 lsmem command 45 sample configuration 11 specifying port number 47 verifying partition resources 45 Linux system required version 64 lschp command 46 lscss command 46 lsqeth command 46

#### Μ

memory dynamically add 49 microcode control level (MCL) 63 migration considerations UUID support 67 z/VM 67 modify partition resources 49

### Ν

navigation keyboard xiii network boot server 77 New Partition step-by-step instructions 33 New Partition task comparison of modes 31 selecting the mode 31 switching modes 32

## 0

OSA-Express 1000BASE-T Ethernet feature code 63

## Ρ

partition boot options 22 configuration settings 12 definition 3 how to start 44 lifecycle 15 modify definition 49 planning checklist 27 properties 12 sample configuration 11 verifying resources in Linux 45 partition definition creating through New Partition task 12 Partition Details task add adapter 49 add memory 49 add processors 49 planning checklist for creating a partition 27 port number specifying in Linux 47 preboot execution environment (PXE) 77 prerequisites DPM-enabled system 63 KVM on IBM z 64 Linux system required version 64 Preventative Service Planning buckets See PSP buckets processors dynamically add 49 PSP buckets corequisite software 65

PSP buckets (*continued*) hardware planning 65 where to find 65 PXE 77

### R

revisions xv

# S

SE See also Support Element EC MCLs for DPM version 63 Secure Service Container prerequisites for using 64 Secure Service Container partition definition 11 shortcut keys xiii start a partition step-by-step instructions 44 Support Element list of Dynamic Partition Manager tasks 57 system types that support DPM 63

# U

universally unique identifier (UUID) See UUID UUID migration action 67 use in network boot server 77

# V

version EC MCL list 63

# Ζ

z/VM 21 migration action 67 z/VM hypervisor commands 47



Printed in USA

SB10-7168-01

