




IBM Power Systems HMC Implementation and Usage Guide

Practical guide to using the IBM HMC in
virtualized Power Systems servers

Documents the RAID 1 feature
on IBM HMC CR6 and CR7

Updated to include HMC
V7R760 and IBM POWER7



Attha Surya Dharma
Mark Hais
Minsung Kang
Yuta Wakayama



International Technical Support Organization

**IBM Power Systems HMC Implementation and
Usage Guide**

April 2013

Note: Before using this information and the product it supports, read the information in “Notices” on page xxi.

Second Edition (April 2013)

This edition applies to Version 7, Release 7, Modification 60 of Hardware Management Console (product number 5639-N47).

© Copyright International Business Machines Corporation 2013. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	ix
Tables	xvii
Examples	xix
Notices	xxi
Trademarks	xxii
Preface	xxiii
The team who wrote this book	xxiii
Now you can become a published author, too!	xxiv
Comments welcome	xxv
Stay connected to IBM Redbooks	xxv
Summary of changesxxvii
April 2013, Second Editionxxvii
Chapter 1. Overview	1
1.1 Sets of function of the HMC	2
1.2 HMC type	3
1.2.1 Desktop HMC	4
1.2.2 Rack mounted HMC	4
1.2.3 HMC version matrix to support Power Systems servers	4
1.2.4 HMC maximums	5
1.3 Enhancements in HMC Version 7R760	5
1.3.1 Web-browser-based user interface	6
1.3.2 System Planning Tool	6
1.3.3 Customizable Data Replication	6
1.3.4 Custom groups	7
1.3.5 New System Reference Code look-up	7
1.3.6 Processor compatibility	7
1.3.7 Donating dedicated processors	9
1.3.8 Processor recovery and partition availability priority	10
1.3.9 Capacity on Demand enhancement	10
1.3.10 Remote DST connection for IBM i	10
1.3.11 RAID 1 protection for the HMC	11
1.3.12 Blade Power Systems management	11
1.3.13 Updated browser support for remote operation	11

1.3.14 PowerVM management enhancements	11
Chapter 2. Planning	13
2.1 System planning tools	14
2.1.1 Work with system plans in the HMC	15
2.1.2 Importing a system plan to the HMC	17
2.1.3 Exporting a system plan from the HMC	18
2.1.4 Creating a system plan on the HMC	21
2.1.5 Viewing a system plan on the HMC	25
2.1.6 Removing a system plan on the HMC	28
2.2 System plans deployment	28
2.2.1 Deployment validation process	28
2.2.2 Deploy a system plan by using the graphical wizard	31
2.3 PowerVM Introduction	40
2.3.1 Processor Virtualization	40
2.3.2 Memory Virtualization	41
2.3.3 I/O Virtualization	41
2.4 Reliability, Availability, Serviceability on the HMC	43
2.4.1 Dual HMC and Redundancy	43
2.4.2 RAID 1 Protection	47
2.4.3 RAID 1 conversion	47
Chapter 3. Installing	59
3.1 Installation of HMC	60
3.1.1 Cabling the HMC	60
3.1.2 HMC Guided Setup wizard	64
3.1.3 Connecting managed systems to the HMC	89
3.2 HMC connectivity scenario	92
Chapter 4. Configuring	97
4.1 Network configuration	98
4.1.1 Types of HMC network configurations	98
4.1.2 Configuring the HMC network setting	98
4.1.3 Testing network connectivity	108
4.1.4 Viewing Network Topology	110
4.2 User Management	113
4.2.1 Managing user profiles and access	114
4.2.2 Changing the user password	114
4.2.3 Customizing user task roles and managed resource roles	115
4.3 Certificate Management	118
4.3.1 Create a certificate	119
4.3.2 Modifying existing certificates	120
4.3.3 Advanced options for modifying existing certificates	121
4.4 Virtualization using HMC	121

4.4.1	Processor virtualization	121
4.4.2	Memory virtualization	122
4.4.3	Virtual I/O	122
4.4.4	Live Partition Mobility	133
4.4.5	Host Ethernet Adapter.	133
4.4.6	Shared pool usage of dedicated capacity	134
4.4.7	Multiple Shared Processor Pool	136
4.4.8	Suspend logical partition	136
4.4.9	Partition availability priority	136
4.4.10	Logical partition management	138
4.4.11	Create an AIX or Linux logical partition	139
4.4.12	Create an IBM i logical partition	158
4.5	Capacity on Demand	160
4.5.1	Advantages of Capacity on Demand.	160
4.5.2	Permanent types of Capacity on Demand	161
4.5.3	Temporary types of Capacity on Demand.	163
4.5.4	Capacity on Demand website navigation	165
4.5.5	Entering CoD codes	166
4.5.6	Stopping Trial CoD	170
Chapter 5. Operating		173
5.1	Basic operation	174
5.1.1	Using the web-based user interface	174
5.1.2	Systems Management: Servers	180
5.1.3	Systems Management: Partitions task	203
5.1.4	Systems Management: Frames task.	219
5.1.5	Using the command-line interface.	223
5.2	HMC Management task.	231
5.2.1	Lock HMC Screen task	231
5.2.2	View HMC Events task	232
5.2.3	Open 5250 Console task.	232
5.2.4	Open Restricted Shell Terminal task.	233
5.2.5	Shut Down or Restart task	233
5.2.6	Schedule Operations task	234
5.2.7	Format Media task	234
5.2.8	Back up HMC Data task	234
5.2.9	Restore HMC Data task	235
5.2.10	Save Upgrade Data task	235
5.2.11	Change Network Settings task	235
5.2.12	Test Network Connectivity task.	235
5.2.13	View Network Topology task.	235
5.2.14	Tip of the Day task	235
5.2.15	View License task	236

5.2.16	Change Default User Interface Settings task	237
5.2.17	Change User Interface Settings task	238
5.2.18	Change Date and Time task	239
5.2.19	Launch Guided Setup wizard	240
5.2.20	Launch Remote Hardware Management Console task	240
5.2.21	Change User Password task	241
5.2.22	Manage User Profiles and Access task	241
5.2.23	Manage Task and Resource Roles task	241
5.2.24	Manage Users and Tasks task	242
5.2.25	Manage Certificates task	243
5.2.26	Configure Key Distribution Center task	243
5.2.27	Configure LDAP task	244
5.2.28	Remote Command Execution task	245
5.2.29	Remote Virtual Terminal task	245
5.2.30	Remote Operation task	246
5.2.31	Change Language and Locale task	246
5.2.32	Create Welcome Text task	247
5.2.33	Manage Data Replication task	248
5.2.34	Manage Install Resources task	253
5.3	Remotely access the logical partitions from HMC	255
5.3.1	Remote AIX, Linux, and Virtual I/O Server Terminal	255
5.3.2	Remote IBM 5250 terminal	256
5.4	Managing partition data	262
5.4.1	Restore task	262
5.4.2	Initialize task	263
5.4.3	Backup task	265
5.4.4	Delete task	266
Chapter 6. Service support		267
6.1	Service Management	268
6.1.1	Management tasks	270
6.1.2	Connectivity	289
6.2	Software maintenance	309
6.2.1	HMC Data backup	309
6.2.2	Restoring HMC Data	314
6.2.3	HMC software maintenance	315
6.2.4	Which firmware or fix level is correct for your system	316
6.2.5	Managed system firmware updates	335
6.3	Advanced System Management Interface	341
6.3.1	Connecting to ASMI	341
6.3.2	Log in to ASMI	344
6.3.3	Power and restart control	346
6.3.4	System Service Aids menu	348

6.3.5 System Information menu	353
6.3.6 System Configuration menu	357
6.3.7 Network Services menu	376
6.3.8 Performance setup	380
6.3.9 On demand utilities	381
6.3.10 Login Profile	384
Appendix A. Introduction to IBM Systems Director	387
Overview of IBM Systems Director	388
IBM Systems Director components	390
Platform Agent	390
Common Agent	391
IBM Systems Director Server	392
Appendix B. Managing POWER processor-based blades	393
Managing a POWER processor-based blade with the HMC	394
Appendix C. IBM product engineering debug data collection	401
Preparing to collect the pedbg	402
Run pedbg collection command	403
Offload pedbg collection	404
Appendix D. Live Partition Mobility support log collection	407
Background information	408
HMC log collection	408
Operating system log collection	409
Virtual I/O Server Mover Server Partition snaps	409
Hypervisor HMC resource dump	410
Sending logs to IBM	411
FTP to IBM	412
IBM ECuRep	412
Related publications	413
IBM Redbooks	413
Other publications	413
Online resources	414
Help from IBM	416
Abbreviations and acronyms	417
Index	421

Figures

1-1 IBM rack mounted 7042-CR6 (left) and IBM KVM 7316-TF3 (right)	4
1-2 Processor capability mode	9
2-1 The HMC Welcome page: System Plans	15
2-2 The main system plan management page	16
2-3 The system plan management page with a system plan selected	17
2-4 Import System Plan window	18
2-5 Format media	19
2-6 Export System Plan window	20
2-7 Create System Plan window	21
2-8 Customize Network Setting: LAN Adapters tab	22
2-9 Customize Network Setting for LAN Adapters: partition communication	23
2-10 Customize Network Setting for LAN Adapters: RMC application	24
2-11 Viewing a system plan	26
2-12 Viewing a system plan	27
2-13 Confirm removal of system plan window.	28
2-14 Launch deployment.	31
2-15 Confirm the deployment startup	32
2-16 Deployment validation in progress	33
2-17 Example of successful validation	34
2-18 Request the details of a specific action.	35
2-19 List of the deployment steps	36
2-20 Partition deployment: Deploy System Plan Wizard.	37
2-21 Deployment complete	39
2-22 Dual HMC configuration on a private network.	43
2-23 Dual HMC physical connections	44
2-24 Example of redundant remote HMC	45
2-25 Create Configuration pane: Selecting the RAID level	54
2-26 Select Drives pane: Check All.	55
2-27 Creation Configuration pane: Saving the configuration.	56
2-28 Select Virtual Drive Operation pane: Start Operation	57
3-1 Rear view of a rack-mounted HMC 7042-CR7	60
3-2 Rear view of a rack-mounted HMC 7042-CR5 and CR6	61
3-3 Rear view of a stand-alone HMC 7042-C06 and 7042-C07	62
3-4 Rear view of a stand-alone HMC 7042-C08	63
3-5 Welcome panel: Launch Guided Setup Wizard	65
3-6 Change Date and Time.	66
3-7 Configure HMC Network Settings	67
3-8 Configure eth0 Private Network	68

3-9	Configure eth0 DHCP	69
3-10	Launch Guided Setup Wizard - Configure HMC Network Settings	71
3-11	Configure eth1 IP assignment	72
3-12	Launch Guided Setup Wizard - Configure HMC Firewall for eth1	73
3-13	Launch Guided Setup Wizard - Configure HMC Firewall for eth1	74
3-14	Launch Guided Setup Wizard - Configure HMC Network Settings	75
3-15	Launch Guided Setup Wizard - Configure HMC Host Name	76
3-16	Launch Guided Setup Wizard: Configure HMC gateway	76
3-17	Launch Guided Setup Wizard - Configure DNS	77
3-18	Launch Guided Setup Wizard - Configure Domain Suffix	78
3-19	Launch Guided Setup Wizard - The Next Steps panel	79
3-20	Notification of Problem Events panel	79
3-21	Status panel	80
3-22	Setting up the call home window	80
3-23	Call Home Setup Wizard: Specify Contact Information	81
3-24	Call Home Setup Wizard: Configure Connectivity	83
3-25	Configure local modem	84
3-26	Configure a pass-through system	85
3-27	Configure Call Home Server Consoles	86
3-28	Authorize users for Electronic Service Agent	87
3-29	Summary of ESA	88
3-30	Update Password - Authentication Passed	90
3-31	Power on the managed system	91
3-32	Connectivity: Multi-drawer with one HMC	92
3-33	Network Redundancy	93
3-34	To support FSP failover	93
3-35	Redundant HMC for High-End Server	95
4-1	HMC Identification tab	99
4-2	LAN Adapters tab	100
4-3	LAN Adapter configuration	101
4-4	Firewall Settings tab	104
4-5	Name Services tab	105
4-6	Routing tab	106
4-7	Network Diagnostic Information: Ping	109
4-8	Network topology	110
4-9	User Profiles window	114
4-10	Change User Password window	115
4-11	Customize User Controls window	115
4-12	Add a new managed resource role	117
4-13	Create New Certificate	119
4-14	Role of PHYP for virtual I/O	123
4-15	Virtual SCSI overview	125
4-16	Create Virtual SCSI Adapter panel	127

4-17	NPIV overview	128
4-18	VLAN example: Two VLANs	130
4-19	Virtual Ethernet connection	131
4-20	Shared Ethernet Adapter configuration	132
4-21	Shared Ethernet Adapter OSI layer	132
4-22	Processor sharing	135
4-23	Setting partition availability priority	137
4-24	Create Virtual I/O Server LPAR	140
4-25	Create Partition panel	141
4-26	Processing Settings panel	142
4-27	Create dedicated processor partition	146
4-28	Processor setting with dedicated processors	147
4-29	Partition memory setting	148
4-30	Partition physical I/O setting	149
4-31	Configuring virtual resource	150
4-32	Create virtual Ethernet adapter	151
4-33	Create virtual SCSI adapter	152
4-34	Logical Host Ethernet Adapters (LHEAs)	153
4-35	Optional settings	154
4-36	Profile Summary window	157
4-37	Restricted IO Partition check box Option	158
4-38	IBM Capacity on Demand offering	161
4-39	Power Systems Capacity on Demand website	165
4-40	Capacity on Demand activation code panel	166
4-41	Capacity on Demand activation code output	167
4-42	HMC Capacity on Demand enter code panel	168
4-43	ASMI Capacity on Demand enter code panel	169
4-44	Stopping Trial CoD panel	171
5-1	HMC Welcome window	175
5-2	HMC Log on window	175
5-3	HMC logoff or disconnect window	176
5-4	Reconnecting the previous session	177
5-5	HMC workplace window	178
5-6	Active tasks in the Task bar	179
5-7	Add Managed Systems window	180
5-8	System Management servers window	181
5-9	Column configuration	183
5-10	Views option	183
5-11	Create new custom view	184
5-12	Properties tasks of the server	185
5-13	Power On task	187
5-14	Power Off task	189
5-15	Power Management task	190

5-16	Identify LED task	191
5-17	Set up a Scheduled Operation task: Date and Time.	193
5-18	Set up a Scheduled Operation task: Repeat.	193
5-19	Customize Scheduled Operations task.	194
5-20	Partition Workload Groups task	196
5-21	Manage Custom Groups task	197
5-22	Manage System Profile task	198
5-23	Service processor status task	199
5-24	Reset or remove connections task	200
5-25	Host Ethernet list.	201
5-26	Host Ethernet Physical Port Configuration	202
5-27	Partition Properties task	204
5-28	Change Default Profile task	205
5-29	Activate Logical Partition task	206
5-30	Restart Partition task.	207
5-31	Shut Down Partitions task	209
5-32	Delete Logical Partition task	211
5-33	Managed Profiles task.	213
5-34	Save Partition Configuration task	214
5-35	Frames option in the navigation pane.	219
5-36	Confirm the HMC version	231
5-37	View HMC events task	232
5-38	Shut down or restart task	233
5-39	Sample of Tip of the Day window	236
5-40	View license task.	236
5-41	Change Default User Interface Settings task	237
5-42	Change User Interface Settings task	238
5-43	Customize Console Date and Time task.	239
5-44	NTP Configuration task.	240
5-45	User and Tasks task	242
5-46	Configure LDAP task.	244
5-47	Remote Command Execution task	245
5-48	Remote Virtual Terminal task	245
5-49	Remote Operation task	246
5-50	Change Language and Locale task	247
5-51	Create Welcome Text task	247
5-52	Sample of welcome text	248
5-53	Manage Data Replication task	250
5-54	Peer-to-Peer replication	251
5-55	Master-to-Subordinate replication	252
5-56	Manage Install Resources task	254
5-57	Virtual terminal window	255
5-58	Virtual terminal error	255

5-59	Start a new session from IBM Personal Communications menu	256
5-60	Configure IBM 5250 session terminal emulator	257
5-61	Configure User ID signon information window	258
5-62	Enter option 21 for American English language option	259
5-63	Select IBM Power Systems to connect	260
5-64	Type of connection to IBM i LPAR	261
5-65	Profile Data Restore task	262
5-66	Initialize Profile Data task	264
5-67	Profile Data Backup task	265
5-68	Backup profile data results window	265
5-69	Profile Data Delete task	266
6-1	Access Service Management from the HMC main menu	268
6-2	Service Management: main view	269
6-3	Service Management: creating a serviceable event	270
6-4	Service Management: manage serviceable events	271
6-5	Service Management: serviceable event overview	272
6-6	Service Management; event description	273
6-7	Service Management: reporting options	274
6-8	Service Management: View details about service event	275
6-9	Service Management, repair action on service event	276
6-10	Service Management: manage problem data	277
6-11	Service Management: close a service event	278
6-12	Service management: remote access options	278
6-13	Manage Remote Connections window	279
6-14	Manage Remote Support Requests window	280
6-15	Service Management: HMC data options	280
6-16	Service Management: Format Media	281
6-17	Service Management: Initiate dump from the server	282
6-18	Service Management: Initiate system dump	283
6-19	Service Management: Initiate system dump for managed system	284
6-20	Service Management: System dump complete	284
6-21	Service Management: Dump results	285
6-22	Service Management: Transmit	286
6-23	Service Management: FTP settings	288
6-24	Service Management: Enable the call home feature	290
6-25	Manage Outbound Connectivity option: Local Modem tab	291
6-26	Modem configuration	291
6-27	Manage Outbound Connectivity option: Internet access	292
6-28	Manage Outbound Connectivity option: Internet VPN	293
6-29	Manage Outbound Connectivity option: Pass-Through Systems tab	294
6-30	Add servers for pass-through access	294
6-31	Add servers for pass-through access: Results	295
6-32	Manage Inbound Connectivity window: Remote Service tab	296

6-33	Manage Inbound Connectivity task, Call Answer tab	297
6-34	Manage Customer Information task, Administrator tab	298
6-35	Manage Customer Information task, System tab	299
6-36	Manage Customer Information task, Account Information tab	300
6-37	Electronic services, registration	301
6-38	Enter email addresses to associate the HMC with eService.	302
6-39	IBM Electronic Support website; Sign in with your IBM ID	303
6-40	Electronic services My systems portal	304
6-41	Manage Serviceable Event Notification task.	305
6-42	Add Email Address panel for notification	306
6-43	Manage Connection Monitoring task.	307
6-44	Manage POWER4 Service Agent task	308
6-45	Back up HMC Data menu	310
6-46	Customize Scheduled Operations window	311
6-47	Set up a Scheduled Operation window.	312
6-48	Scheduled backup HMC Data Repeat Option.	313
6-49	Scheduled HMC Data backup storage Options tab	313
6-50	Restore HMC Data menu from HMC management	315
6-51	Shows the version number of HMC and managed systems	316
6-52	Fix Level Recommendation Tool website	317
6-53	Fix Level Recommendation Tool product options window	318
6-54	Fix Level Recommendation Tool recommendation window	320
6-55	Displays the HMC software that is available on the website.	322
6-56	Choose HMC Firmware for select fix type option	323
6-57	HMC firmware release level menu	324
6-58	HMC software and update files available on the IBM website	325
6-59	Downloading option HMC software and fix images.	326
6-60	Install Corrective Service window	327
6-61	Save Upgrade Data wizard	331
6-62	Firmware and HMC select fix type: System firmware	338
6-63	Firmware and HMC assistance option	339
6-64	System firmware download selection	340
6-65	Launch ASMI from HMC	342
6-66	Advanced System Management main menu	345
6-67	System power on and off options	346
6-68	System Service Aids menu	348
6-69	Error and event logs	349
6-70	Capturing overall system information with system dump procedure. . .	350
6-71	Reset service processor	351
6-72	Factory configuration reset	352
6-73	System Information menu	353
6-74	Display details of VPD.	354
6-75	Displays all VPD detail	355

6-76	Progress Indicator History option	356
6-77	Real-time Progress Indicator	356
6-78	System Configuration menu	357
6-79	System Name	358
6-80	Configure I/O Enclosures window	359
6-81	Time of Day function	361
6-82	Firmware Update Policy before POWER7	362
6-83	Firmware Update Policy in certain models of POWER7 Systems	362
6-84	Hardware Management Consoles window	363
6-85	Deconfiguration Policies window	365
6-86	Processor Deconfiguration window	367
6-87	Processor Deconfiguration window	367
6-88	Memory Deconfiguration window	369
6-89	Memory deconfiguration memory bank selection	370
6-90	Program Vital Product Data panel	370
6-91	System Keyword display example	371
6-92	System Enclosures display example	372
6-93	Service Indicators menu	373
6-94	System Attention Indicator when indicator is on	374
6-95	System Attention Indicator when indicator is off	374
6-96	Enclosure Indicators window	374
6-97	Enclosure Identify window	375
6-98	Network Services menu	376
6-99	HMC Ethernet port configuration	377
6-100	Network Access window	379
6-101	Performance setup	380
6-102	CoD Order Information example	382
6-103	CoD Activation window	383
6-104	Update Installed Languages menu	385
A-1	IBM Systems Director environment	389
B-1	HMC networks diagram	395
B-2	Blade configuration management network	396
B-3	Setting FSP IP address for Power Systems blade	397
B-4	Power Systems blade ASMI login	398
B-5	Hardware Management Console	399
C-1	List HMC user command shows no hscpe user	402
C-2	List HMC user command show hscpe user	402
C-3	Sample of the process of offloading logs to a USB media device	404
C-4	Sample of process offloading logs to a network scp server	405

Tables

1-1 HMC and POWER processors interoperability	5
4-1 Possible status for each node	111
4-2 Meaning of node status.	112
4-3 Predefined HMC roles.	113
4-4 Devices associates with tagged IOAs	159
5-1 Components of the HMC workplace window	178
6-1 Several icons to indicate the type of information in the FLRT report . . .	321
6-2 Default IP address for server connectors HMC1 and HMC2	343
6-3 Default login user ID and password	344

Examples

6-1 HMC pedbg logs for LPM for more than one HMC	408
6-2 Provide AIX client SNAP report.	409
6-3 If you have redundant VIO (four VIO)	410
6-4 Collecting ctsnap.	410
6-5 Rename the file to include PMR number	410
6-6 Collect hypervisor resource dump through a terminal.	411
D-1 Example to archive log files in AIX	411

Notices

This information was developed for products and services offered in the U.S.A.

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, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: 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 states 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 supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

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 the names and addresses used by an actual business enterprise is entirely coincidental.


COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Active Memory™	IBM®	POWER®
AIX 5L™	Lotus®	PureFlex™
AIX®	Micro-Partitioning®	Rational®
BladeCenter®	POWER Hypervisor™	Redbooks®
DB2 Universal Database™	Power Systems™	Redbooks (logo)  ®
DB2®	POWER6+™	System i®
developerWorks®	POWER6®	System p®
Electronic Service Agent™	POWER7 Systems™	System x®
Focal Point™	POWER7+™	System z®
GPFS™	POWER7®	SystemMirror®
HACMP™	PowerHA®	Tivoli®
i5/OS™	PowerVM®	WebSphere®

The following terms are trademarks of other companies:

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

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

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

Other company, product, or service names may be trademarks or service marks of others.

Preface

The IBM® Hardware Management Console (HMC) provides systems administrators a tool for planning, deploying, and managing IBM Power Systems™ servers. This IBM Redbooks® publication is designed for system administrators to use as a desk-side reference when managing partition-capable IBM Power Systems servers by using the HMC.

The major functions that the HMC provides are Power Systems server hardware management and virtualization (partition) management. You can find information about virtualization management in the following documents:

- ▶ *A Practical Guide for Resource Monitoring and Control (RMC)*, SG24-6615
- ▶ *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940
- ▶ *Implementing IBM Systems Director 6.1*, SG24-7694
- ▶ *Hardware Management Console V7 Handbook*, SG24-7491
- ▶ *IBM PowerVM Live Partition Mobility*, SG24-7460
- ▶ *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590
- ▶ *Converting Hardware Management Console (HMC) 7042-CR6 or 7042-CR7 Models to RAID1*, REDP-4909

The following topics are described:

- ▶ Plan to implement the HMC
- ▶ Configure the HMC
- ▶ Operate the HMC
- ▶ Manage software levels on the HMC
- ▶ Use service functions on the HMC
- ▶ Update firmware of managed systems
- ▶ Use IBM System Planning Tool deployments

In addition, there is an explanation on how to use the new HMC graphical user interface and the new HMC commands that are available with HMC Version 7, Release 7, modification 60.

The team who wrote this book

This book was produced by a team of specialists from around the world, working at the International Technical Support Organization, Poughkeepsie Center.

Attha Surya Dharma is a Client IT Specialist in Indonesia. He has four years of experience in the IBM i field. He holds a degree in Electrical Engineering from University of Indonesia. His areas of expertise include Implementation, Backup Recovery, and System Availability consulting on the IBM i Environment.

Mark Hais is an IT Specialist in IBM Israel. He has 18 years of experience in IBM Power Systems. He has worked at IBM for 18 years. His areas of expertise include Power Systems servers and blades, IBM AIX®, IBM PowerVM®, IBM PowerHA®, and IBM PureFlex™ Systems.

Minsung Kang is a service support representative (SSR) in Korea. He has three years of experience in high-end storage and Power Systems. He holds a degree in Information and Communication Engineering from Inha University. His areas of expertise include high-end storage, Power Systems, AIX, and IBM HACMP™.

Yuta Wakayama is an IT Specialist working in Technical Sales in Osaka, IBM Japan. He has four years of experience with IBM Power Systems and AIX. He provides pre-sales technical consultation and implementation of IBM Power Systems and virtualization environments.

The project that produced this publication was managed by:
Scott Vetter, PMP

Thanks to the following people for their contributions to this project:

Dave Bennin, Richard M. Conway, Ann Lund, Al Schwab
International Technical Support Organization, Poughkeepsie Center

Karyn T. Corneli, David A. Dilling, Shawn D. Mohr, Jacobo A. Vargas
IBM US

Thanks to the authors of the previous editions of this book.

- ▶ Authors of the first edition, *Hardware Management Console V7 Handbook*, published in October 2007, were:
Narend Chand, Syamsul Hidayat, Stephen Hochstetler, Nancy Milliner,
JunHeum Min, and Matt Robbins

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance

and client satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:

<http://www.facebook.com/IBMRedbooks>

- ▶ Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- ▶ Look for us on LinkedIn:

<http://www.linkedin.com/groups?home=&gid=2130806>

- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>

Summary of changes

This section describes the technical changes made in this edition of the book and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes
for SG24-7491-01
for IBM Power Systems HMC Implementation and Usage Guide
as created or updated on May 21, 2015.

April 2013, Second Edition

This revision reflects the addition, deletion, or modification of new and changed information described below.

New information

- ▶ The new publication includes all new enhancements of HMC new version 7.7R60
- ▶ Includes new RAID 1 support for HMC rack model CR6 and CR7
- ▶ The new publication also describes the uses of HMC to support IBM Power Systems in virtualized environments

Changed information

- ▶ Updated for support of HMC Software Version 7.3 to Version 7.7. Generally, all content is updated to match the new version.
- ▶ Restructured the old publication chapters for easy uses as one lifecycle HMC implementation; from plan, installation, configuring, managing, and support of the HMC.



Overview

The *Hardware Management Console (HMC)* is a member of the IBM Systems Director platform management family. IBM Systems Director provides IT professionals with the tools that they need to better coordinate and manage all of their virtual and physical resources in the data center.

The cost of managing the IT infrastructure has become the largest and fastest-growing component of overall IT spending for many organizations. Virtualization helps to address this cost through the consolidation of physical resources. However, virtualization also adds complexity to the system by creating a sharp increase in the number of managed virtual resources.

IT professionals are seeking more advanced capabilities and tools for managing both their physical and virtual systems across multiple architectures and environments. As virtualization becomes reality in today's IT infrastructures, the IBM Systems Director family can help businesses realize their full potential by providing a unified approach to platform management. This approach is designed to lower IT operational costs and increase productivity. We describe the HMC functions that manage the IBM Power Systems, IBM System p®, and IBM System i® servers.

The HMC concepts, the types of HMC, HMC connectivity, and enhancements in HMC Version 7R760 are now described.

For more information about HMC, see the following sources:

- ▶ *Operations Guide for the Hardware Management Console and Managed Systems*, SA76-0085
- ▶ IBM Power Systems Hardware Information Center
<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp>

1.1 Sets of function of the HMC

With the HMC, a system administrator can do logical partitioning functions, service functions, and various system management functions by using either the web-browser-based user interface or the command-line interface (CLI). The HMC uses its connection to one or more systems (referred to in this book as *managed systems*) to do various functions:

- ▶ Creating and maintaining logical partitions in a managed system
The HMC controls logical partitions in managed systems. These tasks are explained in detail in section 4.4.10, “Logical partition management” on page 138.
- ▶ Displaying managed system resources and status
These tasks are explained in section 5.1.2, “Systems Management: Servers” on page 180.
- ▶ Opening a virtual terminal for each partition
The HMC provides virtual terminal emulation for AIX and Linux logical partitions and virtual 5250 console emulation for IBM i logical partitions.
- ▶ Displaying virtual operator panel values for each partition
You can see the operator panel messages for all partitions within managed systems in HMC.
- ▶ Powering managed systems on and off
We explain these tasks in “Operations” on page 186.
- ▶ Performing dynamic LPAR (DLPAR) operation
With the HMC, you can do DLPAR operations that change the resource allocation (such as processor, memory, physical I/O, and virtual I/O) dynamically for the specified partition. These tasks are explained in detail in section 4.4.10, “Logical partition management” on page 138.
- ▶ Managing Capacity on Demand operation
These tasks are explained in section 4.5, “Capacity on Demand” on page 160.

- ▶ Managing virtualization features

These tasks are explained in section 4.4, “Virtualization using HMC” on page 121.

- ▶ Managing platform firmware installation and upgrade

These tasks are explained in section 6.2.5, “Managed system firmware updates” on page 335.

- ▶ Acting as a service focal point

You can use the HMC as a service focal point for service representatives to determine an appropriate service strategy and to enable the service agent to call home to IBM. We explain these tasks in 6.1, “Service Management” on page 268.

HMC Version 7 uses a web-browser-based user interface. This interface uses a tree-style navigation model that provides hierarchical views of system resources and tasks by using drill-down and launch-in-context techniques to enable direct access to hardware resources and task management capabilities. This version provides views of system resources and provides tasks for system administration. For more information about using the web-browser-based user interface, see section 5.1.1, “Using the web-based user interface” on page 174.

The remote interface changed in this release to also use a browser interface instead of a WebSM interface.

1.2 HMC type

The HMC runs as an embedded application on an Intel based workstation that can be a desktop or rack mounted system. The embedded operating system and applications take over the entire system, and no other applications are allowed to be loaded.

Whether you opt for a desktop or rack mounted version, is a personal choice. Clients with space in their rack mounted systems would probably opt for the rack mounted version with the slide-away keyboard and display. The following options are available:

- ▶ 7042-C07 and 7042-C08 are desktop HMCs.
- ▶ 7042-CR6 and 7042-CR7 are rack mounted HMCs.

HMC Version V7R760 is the last release to be supported on models 7310-C04, and 7310-CR2. Future versions of the HMC will not support these models.

Figure 1-1 is a picture of the IBM desktop and rack mounted HMC.



Figure 1-1 IBM rack mounted 7042-CR6 (left) and IBM KVM 7316-TF3 (right)

1.2.1 Desktop HMC

The supported desktop models are the 7042-C06 and older version 7310 models. The older 7315 models are not supported by V7R760 of the HMC.

On the desktop, you can connect a keyboard and mouse to either the standard keyboard, mouse PS/2 style connectors, or USB ports. You cannot connect any other devices to the HMC. Printers are not supported off the parallel port. At the time of writing, this model is no longer available on the market.

1.2.2 Rack mounted HMC

The supported rack mounted models are the 7042-CR4 and older version 7310 models. The older 7315 models are not supported by V7R760 of the HMC. Figure 1-1 shows the HMC 7042-CR6 system unit as a standard 1 U unit and also the display and keyboard mounted in a standard 1 U pull-out tray. This model is a great choice for a dark machine room, where space is restricted.

For more information, see the Hardware Information Center:

<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp>

1.2.3 HMC version matrix to support Power Systems servers

The HMC has several versions to support the different type of IBM POWER® processors (POWER4, POWER5, POWER5+, IBM POWER6®),

IBM POWER6+™, IBM POWER7®, and IBM POWER7+™). This publication focuses on the HMC supporting POWER7. POWER5 servers must be at least GA7 SF240 firmware level to be managed by the new version of HMC.

Table 1-1 shows the interoperability between HMC and POWER processors.

Table 1-1 HMC and POWER processors interoperability

HMC Version	Supported managed system
Version 7 and later	POWER7+, POWER7, POWER6+, POWER6, POWER5+, and POWER5 ^a
Version 4 and later, 5 and later, 6 and later	POWER5+, POWER5
Version 3 and later	POWER4

a. POWER5 servers must be at least GA7 SF240 firmware level.

You can find detailed information about HMC POWER code matrix at:

<http://www.ibm.com/support/customer/care/sas/f/power5cm/home.html>

1.2.4 HMC maximums

At the time of writing, the following are general support considerations regarding the HMC:

- ▶ A maximum of 48 non-590/595/795 servers are supported
- ▶ A maximum of 32 590/595/795 servers are supported
- ▶ For all systems logical partition (LPAR) is 1000 on HMC V7R760.6 or later running on 7042-CR6 or later with minimum of 4 GB memory
- ▶ In an HMC managed enterprise, a maximum of 2 HMCs can manage a server at one time.

1.3 Enhancements in HMC Version 7R760

HMC Version 7 includes several new functions that support new POWER7 and POWER6 technology features. The HMC also configures and manages IBM System p and System i systems that are based on the POWER5 technology.

This section introduces the new functions of HMC Version 7 and POWER7. There are several new functions in this version of the HMC V7R760.

1.3.1 Web-browser-based user interface

To remotely access an HMC running HMC Version 4, 5, or 6, a special client program (WebSM) was required. WebSM functionality is replaced in HMC Version 7. Remote access to an HMC running Version 7 requires only a standard (and supported) web browser.

The new web-based user interface includes the following highlights:

- ▶ Persistent graphical user interface (GUI) session across login
- ▶ The ability to manage POWER5, POWER6, and POWER7

The following differences in the GUI are observed:

- ▶ Simplified operations
- ▶ Reorganized panels
- ▶ More descriptive GUI settings
- ▶ Redesigned panels

For more information about the web-browser based user interface, see 5.1.1, “Using the web-based user interface” on page 174.

1.3.2 System Planning Tool

The System Planning Tool (SPT) is a tool for designing logically partitioned systems and is the replacement for the LPAR Validation Tool (LVT). SPT creates a system plan that is saved as a sysplan file. That system plan can be just one system or it can contain multiple systems, each with a unique system name.

A system plan, also referred to a *sysplan*, is a representation or data model of the resources that are included in the system and how they are allocated to each partition. When you create the sysplan in SPT, the file reflects the intended LPAR configuration for a target server. This sysplan includes details about partition allocations of memory, processors, and the I/O hardware that is required for each partition.

For more information about the SPT, see 2.1, “System planning tools” on page 14.

1.3.3 Customizable Data Replication

Customizable Data Replication allows another HMC to obtain customized console data from or send data to this HMC. For more information, see 5.2.33, “Manage Data Replication task” on page 248.

1.3.4 Custom groups

Custom groups provide a mechanism for you to group system resources together in a single view or a way to organize the systems or partitions into smaller business or workload entities. For more information, see “Manage Custom Groups task” on page 196.

1.3.5 New System Reference Code look-up

There is a single repository for all system reference codes (SRCs) and progress codes for POWER6 and IBM POWER7 Systems™. The SRC is a sequence of data words (codes) that do the following functions:

- ▶ Identify a system status
- ▶ Describe a detected hardware, Licensed Internal Code (LIC), or software failure
- ▶ Describes the unit that is reporting the failure and its location

SRCs can be viewed on the control panel of a system, as a system console message, or from the following three panels on the HMC:

- ▶ Managed Serviceable Events overview display
- ▶ The Reference Code Column of the Server display
- ▶ The Reference Code History display

For POWER6 and POWER7 processor-based servers, the HMC provides active web links to the SRC repository. Clicking these links displays the additional SRC information.

For more information, see “Service events” on page 270.

1.3.6 Processor compatibility

POWER7 Systems support running in one of the following modes:

- ▶ POWER5

The POWER5 processor compatibility mode allows the running of operating system versions that use all the standard features of the POWER5 processor.

- ▶ POWER6

The POWER6 processor compatibility mode allows the running of operating system versions that use all the standard features of the POWER6 processor.

- ▶ POWER6+
The POWER6+ processor compatibility mode allows the running of operating system versions that use all the standard features of the POWER6+ processor.
- ▶ POWER6 enhanced
The POWER6 enhanced processor compatibility mode allows the running of operating system versions that use all the standard features of the POWER6 processor and also provides extra floating-point instruction to applications that use the POWER6 processor.
- ▶ POWER6+ enhanced
The POWER6+ enhanced processor compatibility mode allows the running of operating system versions that use all the standard features of the POWER6 processor and also provides extra floating-point instruction to applications that use the POWER6+ processor.
- ▶ POWER7
The POWER7 processor compatibility mode allows the running of operating system versions that use all the standard features of the POWER7 and POWER7+ processor.
- ▶ Default
The default processor compatibility mode is a preferred processor compatibility mode that enables the hypervisor to determine the current mode for the logical partition. When the preferred mode is set to default, the hypervisor sets the current mode to the most fully featured mode that is supported by the operating environment. In most cases, this mode is the processor type of the server on which the logical partition is activated. For example, assume that the preferred mode is set to default and the logical partition is running on a POWER7 processor-based server. The operating environment supports the POWER7 processor capabilities so the hypervisor sets the current processor compatibility mode to POWER7.

LPARs running on the same POWER7 system can run in different modes. All POWER7 LPARs default to POWER7 architecture mode. You can see and modify that mode in the command-line interface. The GUI allows display of the current LP compatibility mode but does not allow modification. The status is shown in the Partition Properties window as shown in Figure 1-2.

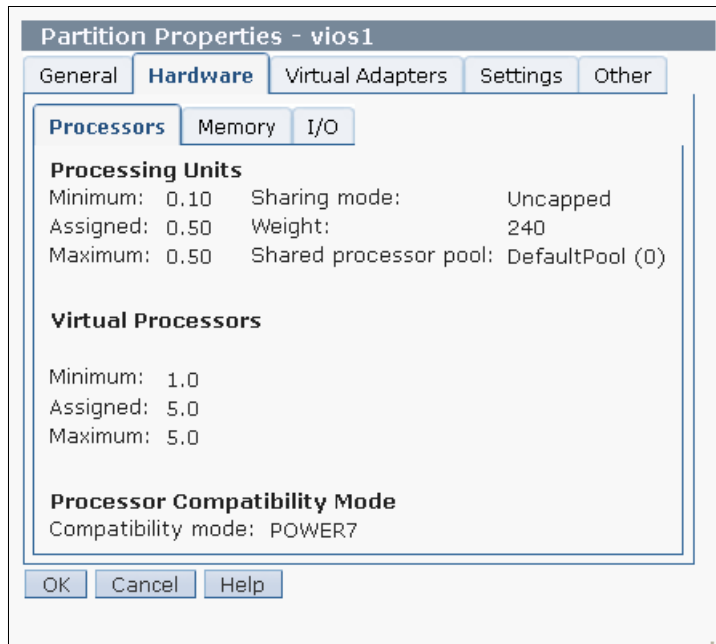


Figure 1-2 Processor capability mode

Restriction: A POWER6 processor cannot emulate all features of a POWER5 processor. Similarly, a POWER7 processor cannot emulate all features of a POWER6 or a POWER5 processor. For example, certain types of performance monitoring might not be available for a logical partition if the current processor compatibility mode of a logical partition is set to the POWER5 mode.

1.3.7 Donating dedicated processors

POWER7 and POWER6 allow dedicated processors in LPAR to donate its idle processor cycles to the shared processor pool instead of being wasted as cycles in the dedicated partition as like shared processors LPAR. You can enable this function in the HMC. We explain how you can set up this function by using HMC in 4.4.1, "Processor virtualization" on page 121.

1.3.8 Processor recovery and partition availability priority

The POWER7 and POWER6 processor supports enhanced RAS capabilities. One of enhanced RAS capacities make firmware checkpoint the state of a failed processor. The checkpoint state can be resumed on another good processor. This function is called *Processor Recovery*.

Sometimes, this process causes a loss of entitled capacity for one or more Shared Processor LPARs. Then, firmware notifies those partitions of the loss of capacity. To determine which LPARs prefer to be stolen capacity, you can set the *Partition Availability Priority* of LPAR by using HMC.

All processor recovery actions and loss of entitled capacity are logged in the system error log.

1.3.9 Capacity on Demand enhancement

Capacity on Demand (CoD) can be configured with inactive processor and memory resources that can be enabled dynamically and non-disruptively to the system. The CoD offering provides flexibility and improved granularity in processor and memory upgrades and is a useful configuration option to support future growth requirements. CoD also eliminates the need to schedule downtime to install extra resources at a specified point in time.

This CoD was introduced in POWER5 products. There are two new types of CoD and some enhancements:

- ▶ Capacity Upgrade on Demand (CUoD)
- ▶ Mobile CoD
- ▶ Trial CoD
- ▶ Utility CoD
- ▶ Capacity BackUP (CBU)
- ▶ On/Off CoD

We explain these tasks in 4.5, “Capacity on Demand” on page 160.

In POWER7 Systems HMC support, to better support IBM Pooled Capacity on the Power 795 system, you can now order an unlimited number of CoD days. The previous limit was 9999.

1.3.10 Remote DST connection for IBM i

Starting from HMC V7, we can do a remote 5250 session console dedicated service tools (DST) through an IP address connection with your 5250 Emulation

terminal from your workstation. The details of the usage of this feature you can find in 5.3.2, “Remote IBM 5250 terminal” on page 256.

1.3.11 RAID 1 protection for the HMC

Starting from the HMC 7042-CR7, RAID 1 protection is enabled by default (unless the feature is removed from the order), though the HMC 7042-CR6 optionally can be converted to have RAID 1. This feature is enabling data redundancy and high levels of performance of the HMC.

RAID 1 uses data mirroring, Two physical drives are combined into an array, and data is stripped across the array. The first of a stripe is the original data. The second half is a mirror (that is, a copy) of the data, but it is written to the other drive in the RAID 1 array. Enabling an MES upgrade for this feature on IBM 7042-CR6 or 7042-CR7 HMC Software reinstallation is required. For more information, see 2.4.3, “RAID 1 conversion” on page 47

1.3.12 Blade Power Systems management

With the release of HMC V7R760, the HMC can now manage IBM BladeCenter® Power Blade servers. This management includes support for dual Virtual I/O Servers (VIOs), live partition mobility between blades and rack servers, and management of both blades and rack servers from a single management console. You can find more information in Appendix B, “Managing POWER processor-based blades” on page 393.

1.3.13 Updated browser support for remote operation

The HMC V7R760 now supports Mozilla Firefox 7 through 10 and Microsoft™ Internet Explorer 7 through 9.

1.3.14 PowerVM management enhancements

There are several new HMC features regarding the update of the PowerVM new capabilities.

- ▶ You can specify the destination Fibre Channel port for any or all virtual Fibre Channel adapters during a live partition migration.
- ▶ You can create up to 20 virtual machines per processor core for POWER7+ processor, 10 virtual machines for POWER7, POWER6, and POWER5.
- ▶ During DLPAR, add or remove virtual I/O adapters to or from a Virtual I/O Server. The HMC now automatically attempts to run the Add/Remove

commands (**cfgdev** and **rmdev**) on the Virtual I/O Server. In earlier releases, you had to manually run these commands on the Virtual I/O Server.

- ▶ HMC now supports up to 16 concurrent *Live Partition Mobility* activities.
- ▶ A Host Ethernet Adapter (HEA) allows multiple logical partitions to share a single special physical Ethernet adapter. Unlike most other types of I/O devices, you can never assign the HEA itself to a logical partition. Instead, virtual subsets called *logical HEAs* or *LHEAs* are defined on the physical HEA. These LHEAs can then be assigned to logical partitions. This assignment allows these logical partitions to access external networks through the HEA without having to go through an Ethernet bridge on the Virtual I/O Server or another logical partition.



Planning

You can use the Hardware Management Console (HMC) to create import, export, view, create, remove, and deploy system plans. The HMC code level should be at Version 7 Release 3.1 or later, and it should include the latest service packs. The HMC provides a set of graphical user interfaces (GUIs) for these logical partition (LPAR) management functions.

2.1 System planning tools

A system plan, also referred to as *sysplan* because of the `.sysplan` file extension, is a representation of the hardware and partition configuration currently on a system. Or, the sysplan is a plan for deployment of hardware and configuration of partitions on a system, depending on how the sysplan file is generated:

- ▶ If the sysplan file is generated from an existing POWER5 or POWER6 system by using the HMC, the file reflects the actual LPAR configuration of the server then if all the partitions are active. There is less detail if one or more partitions are not active.
- ▶ If the sysplan is generated by the System Planning Tool (SPT), the file reflects the intended LPAR configuration for a target server. The sysplan includes details about partition allocations of memory, processors, and the hardware that is required for each partition.

Hardware allocations can be defined as owned by the partition, and therefore required for the partition to activate. The only other choice is for the hardware to be defined as *shared*, in which case the hardware is optional for the partition, which can be activated without the hardware. Shared hardware can be switched dynamically between two or more partitions.

The sysplan also includes general information about the system, such as system type and model, total number of processors present and the number that are activated, and the total installed and activated memory. It has detailed information about the card slots in the processor enclosure and any I/O expansion towers or I/O drawers that attach to the processor enclosure. The card slots are shown as empty, or occupied by input/output processor (IOP) or input/output adapter (IOA) feature codes, and this level of detail is used for the hardware validation during the LPAR deployment process.

Currently, the sysplan file that is created by SPT includes device-level detail (for example, what type and number of disk units are attached to a storage controller IOA). In contrast, a sysplan that is created by the HMC does not, by default, include any detail of what is attached to and controlled by an IOA.

A sysplan file is a composite object, which means that it might possibly include many files. The description of the file is embedded in the file, as is the file level and last modifying application information. When a sysplan file is created on the HMC or imported by using the HMC GUI, the file is stored on the HMC in a predefined directory. The directory path is `/opt/hsc/data/sysplan`.

2.1.1 Work with system plans in the HMC

In the HMC workplace window, *System Plans* is where you can access the graphical interfaces that you use to manage system plans on the servers directly from the HMC or remotely by using the web-browser based client connecting to the HMC (see Figure 2-1).

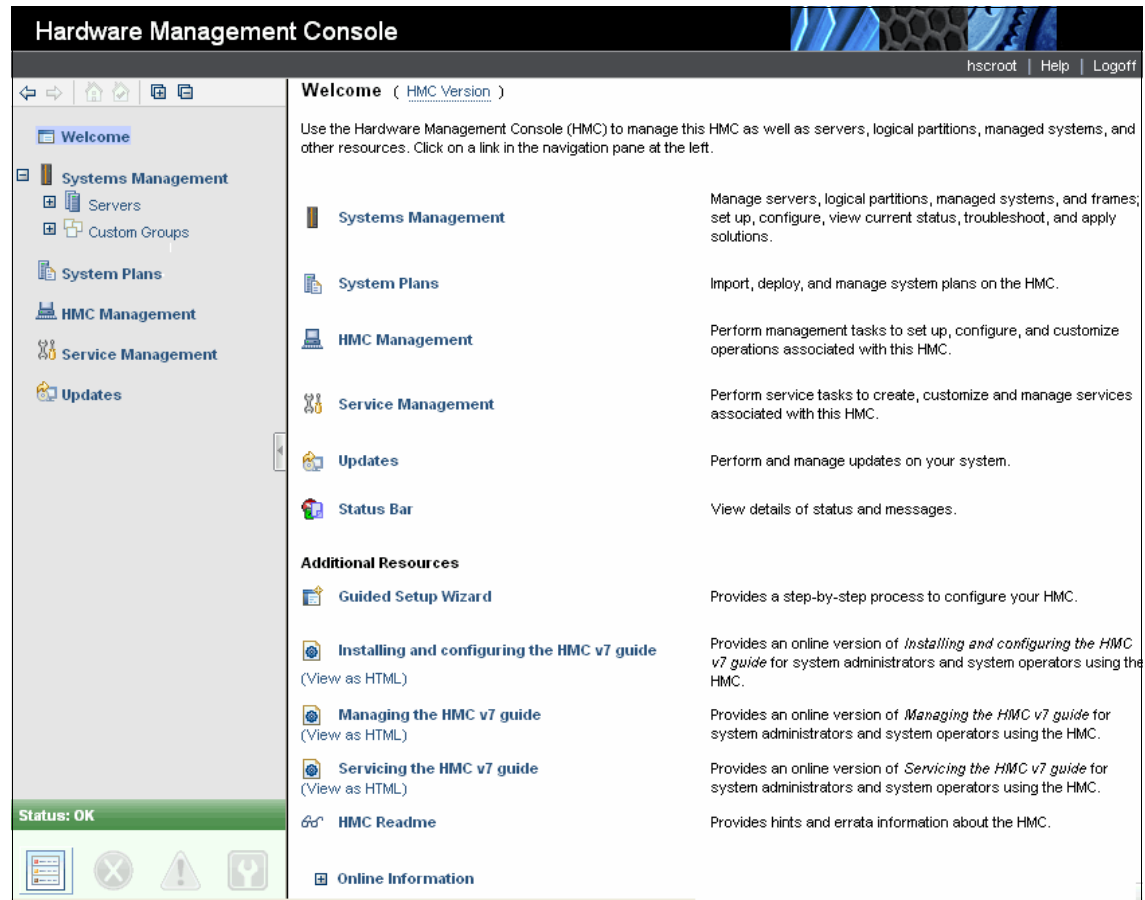


Figure 2-1 The HMC Welcome page: System Plans

To display the system plans management tasks window, click **System Plans** (Figure 2-2). The upper section of this window lists all the system plans currently on the HMC. The icons above the list let you select and clear, sort, filter, and manage the columns of the display table, and perform tasks on selected system plans. The task options are repeated in the lower *tasks* section of the main system plans management window. With no system plan selected, the only options are to import a system plan or to create a system plan.

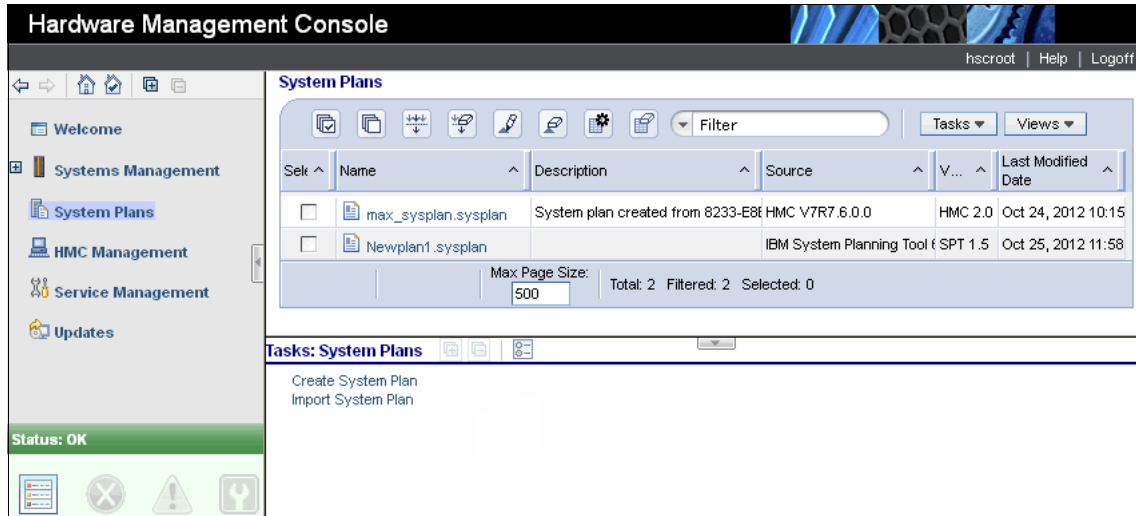


Figure 2-2 The main system plan management page

Using the HMC, you can do the following actions:

- ▶ Create a system plan
- ▶ View a system plan
- ▶ Deploy a system plan
- ▶ Export a system plan
- ▶ Import a system plan
- ▶ Remove a system plan

You can save a system plan that is created by using the HMC interface as a record of the hardware and partition configuration of the managed system at a specified time.

You can deploy an existing system plan to other systems that this HMC manages that have hardware that is identical to the hardware in the system plan.

You can export a system plan to another HMC (which imports the plan). You can then use it to deploy the system plan to other systems that the target HMC manages that have hardware that is identical to the hardware in the system plan.

You can view, create, deploy, export, import, or remove a system plan. These tasks can be selected in either the Tasks menu or the Tasks links in the lower part of the right frame. The following sections provide more details for each option.

Figure 2-3 shows a common starting point for each example. In our first example, we selected a system plan named `max_sysplan.sysplan`.

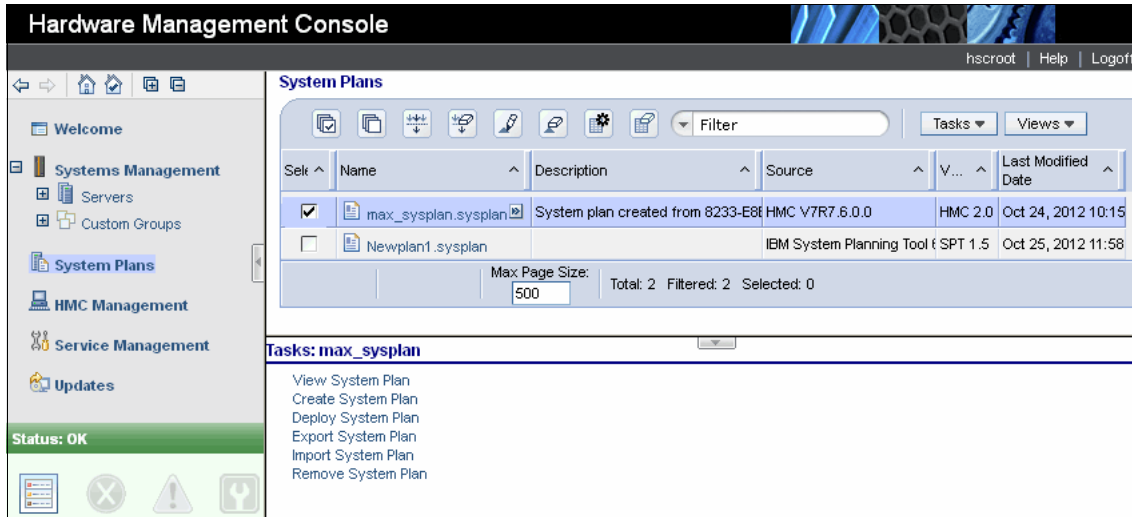


Figure 2-3 The system plan management page with a system plan selected

2.1.2 Importing a system plan to the HMC

You can load a system plan that was created by using the SPT or created on another HMC by using the import operation. You can import the system plan from one of the supported media types. Types include CD, DVD, or a USB device such as a memory card, a remote FTP site, or a PC connected to the HMC through a browser connection.

When you import a system plan, you first must prepare the media, if needed. Then, you import the sysplan file.

From the System Plans task menu, select **Import System Plan**, which opens the Import System Plan prompt window. Identify the system plan file name and whether you are importing it from media, an FTP server, or, if you are accessing the HMC through a PC-based web browser, the sysplan file can be on that PC.

In our example (Figure 2-4), the system plan file is stored on a USB flash drive. The name of the file is `newConfig.sysplan`, and the file was initially created by using SPT and saved to the flash drive. The directory path to access the file on the flash drive is `/media/sysdata`.

Import System Plan - max_sysplan.sysplan
You can import a system plan file to your HMC from the following sources.
Select the source of the system plan file

Import from this computer to the HMC

Import from media

System plan file name: * newConfig.sysplan
Sub-directory on media: media/sysdata

Import from a remote FTP site

System plan file name: *
Remote site hostname: *
User ID: *
Password: *
Remote directory:

Import Cancel Help

Figure 2-4 Import System Plan window

2.1.3 Exporting a system plan from the HMC

You can export system plans that are on the HMC to media, an FTP server, or, if you are using a PC to access the HMC through a browser, to a directory on the PC. The process is much like the importing of a sysplan file. If you are exporting to media, you must format that media for use with the HMC.

Preparing the media

To export to external media, that media has to be in a format that is available to the HMC. The easiest method is using the Format Removable Media task: Select **HMC Management** from the left navigation frame.

1. Select **HMC Management** from the left navigation frame.

2. Select **Format Media** in the right window to open the media selection box as shown in Figure 2-5.

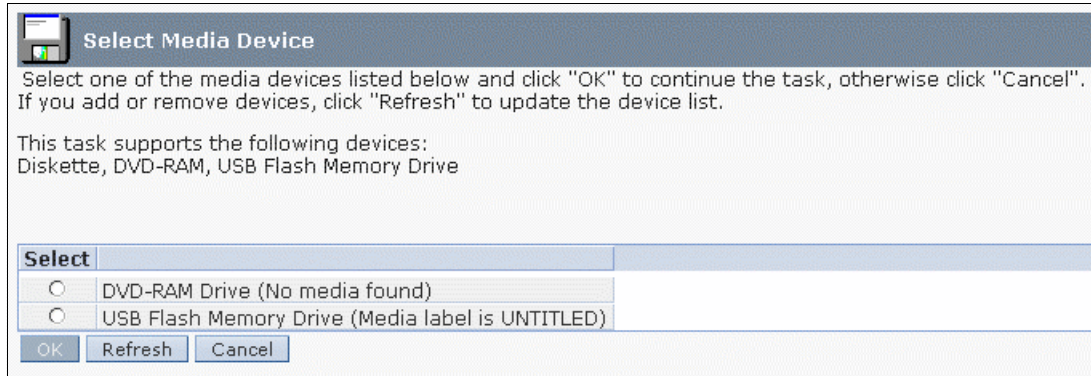


Figure 2-5 Format media

If you have a USB memory key, insert it in a USB slot on the HMC.

If you have a diskette or CD that must be formatted, insert it into the disk or CD drive.

3. Select the correct device to format and click **OK**. The memory format process starts and completes.
4. Insert the media into the PC and load the system plan file by using the save function in SPT or by browsing to the file and copying the sysplan file to the media.

Exporting the system plan

1. Select a system plan to export.
2. Click **Export System Plan** either from the Tasks menu or the content Tasks link in the lower portion of the window. A dialog box opens asking where you want to export the system plan (Figure 2-6). In our example, the name of the sysplan file is apr1119.sysplan and the target is media/USBstick directory.

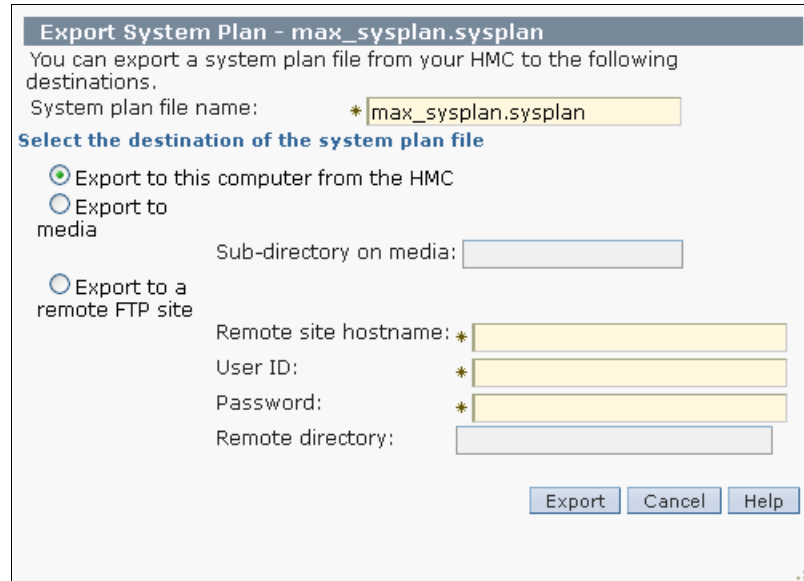


Figure 2-6 Export System Plan window

3. Click **Export** to initiate the export process. A results window with a success indication or an error message indicates the result of the export.

2.1.4 Creating a system plan on the HMC

You can create a system plan for a system that is controlled by the HMC. The system plan has information about the current partition definitions and hardware allocations. Processor, memory, and PCI cards are identified in the system plan, even if they are not owned by a partition.

Notes:

1. Hardware that is controlled through an IOA controller, such as disk units and external media devices, is not represented in the system plan unless the owning partition is running.
2. You cannot import the sysplan file that the HMC creates into the SPT to edit it. The sysplan file can be only deployed and viewed either on the HMC on which the file was created or an HMC to which the file was moved.

From the starting point in Figure 2-3, follow these steps:

1. Select **Create System Plan**.
2. The Create System Plan window prompts you for the system name, sysplan file name, a description, and a choice to view the system plan after creation, as shown in Figure 2-7.

Create System Plan - max_sysplan.sysplan
Select the managed system, and specify a name and description for the system plan.

Managed system: 8233-E8B-SN10DD51P

System plan name: Oct25.sysplan

Plan description: System plan created from 8233-E8B-SN10DD51P

Retrieve inactive and unallocated hardware resources
*Note: Using this option can add several minutes to the system plan creation process.

View system plan after creation

Create Cancel Help

Figure 2-7 Create System Plan window

3. After you enter the requested information, click **Create**.
4. Following the successful creation of a system plan, a message displays and the system plan is now in the list of plans on the HMC. Click **OK**.

Enabling hardware inventory collection from active partitions

When you use the HMC to create a system plan for a managed system, you can capture partition configuration information and a base set of associated hardware configuration information. If you have partitions already active, you can maximize the information that the HMC can obtain about the hardware.

To maximize the information that the HMC can obtain from the managed system, turn on the managed system and activate the logical partitions on the managed system, assuming that they exist, before you create the system plan.

Additionally, you must set up Resource Monitoring and Control (RMC) on the HMC before you create a system plan to capture the most detailed information. Although using the RMC can take several more minutes to finish processing, you can capture disk drive and tape drive configuration information for a managed system in the system plan. You can view this more detailed hardware information by using the View System Plan task.

To enable the HMC's internal inventory collection tool (*invscout*) to be able to do its most detailed hardware inventory retrieval operations, follow these steps:

1. In the HMC workplace window, select the HMC Management task.
2. Select **Change Network Settings**, and in the Customize Network Settings window, select the LAN Adapters tab, as shown in Figure 2-8.

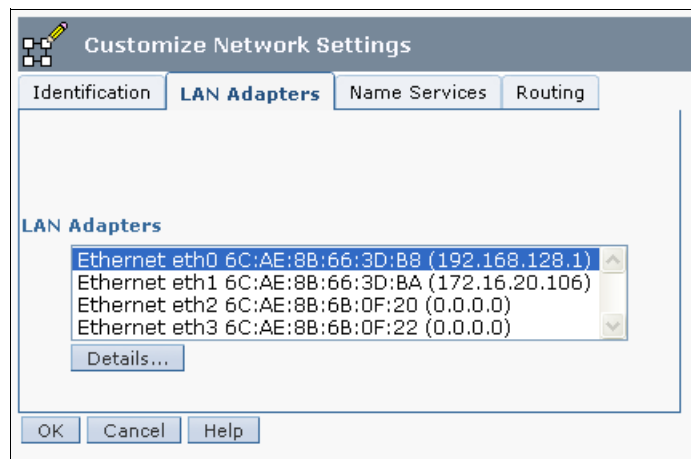


Figure 2-8 Customize Network Setting: LAN Adapters tab

3. In the LAN Adapters window, select the **eth0** LAN Adapter and click **Details**.
4. In the LAN Adapter Details window (Figure 2-9) on the LAN Adapter tab, select **Open** within the local area network information area to enable the check box for Partition Communication. Then, select **Partition communication**.

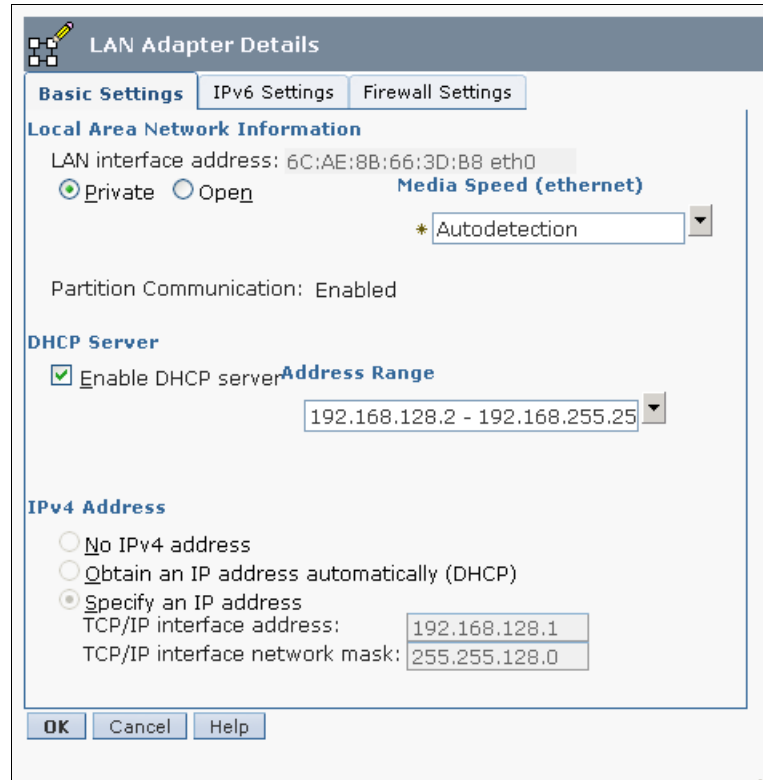


Figure 2-9 Customize Network Setting for LAN Adapters: partition communication

- Click the **Firewall Settings** tab, scroll down the Available Applications area to see whether RMC is already specified as available. In this example, we assume that RMC is not yet available. Therefore, select **RMC** in the Allowed Hosts pane and click **Allow Incoming**, as shown in Figure 2-10.

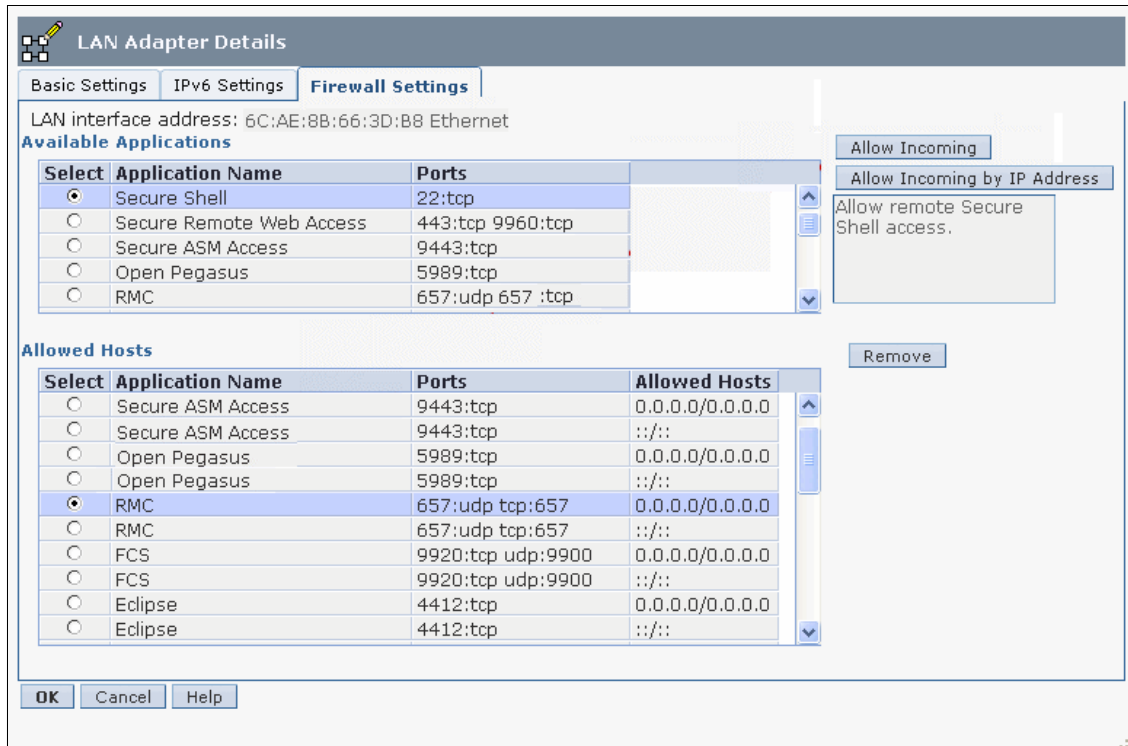


Figure 2-10 Customize Network Setting for LAN Adapters: RMC application

This action moves RMC into the Available Applications pane.

- Click **OK** twice to open a window that states that the Network Settings Changes are applied at the next HMC.
- Click **OK**. You are now back to the HMC workplace window with just the HMC Management pane on the right.

You can verify that you enabled RMC successfully by using the **lspartition** command on the HMC CLI. For more information about using the HMC CLI, see Figure 2-21 on page 39.

The list partition command:

```
lspartition -c
```

For example:

```
hmc:> lspartition -c 9117_MTM-10FZZD
```

In our example managed system, this command results in:

```
<#0> Partition:<4, partn1.business.com, 1.2.3.444>  
Active:<0>, OS<, >
```

If this command does not return any partitions, the system might not be set up for RMC. Depending on whether the system is a Power Systems server, System i, or System p, the steps for RMC are different.

IBM Systems Hardware Information Center includes more information about RMC. For background information about RMC, you can also see *A Practical Guide for Resource Monitoring and Control (RMC)*, SG24-6615. The content of this publication is based on IBM AIX 5L™, 5.1.

If the Create System plan from the GUI fails and if there is a need to create a system plan, use the underlying `mksysplan` CLI at the HMC command prompt, with the `noprobe` option. The `noprobe` option bypasses the default inventory collection of active partitions. Therefore, the resulting sysplan might not have IOA or IOP controlled disk units or media enclosures.

For example:

```
hmc:> mksysplan -m machineName -f filename.sysplan -v -o noprobe
```

When creating a sysplan, if there is a failure because of a Virtual I/O Server (VIOS) error, you can try the `noprobe` option from the CLI.

2.1.5 Viewing a system plan on the HMC

The HMC has a system plan viewer similar to the viewer in the System Planning Tool. The viewer offers a non-editable presentation of the partitions and hardware of the system. Using Figure 2-3 as a starting point, select the wanted plan in the main system plan management window. Click **View System Plan**.

When you are accessing the HMC remotely, you are presented with a View System Plan sign-on window the first time that you start the System Plan Viewer. This additional login protects unauthorized users from viewing the configuration of the system. It also prevents starting the viewer from bookmarks without providing an appropriate user name and password.

Figure 2-11 shows the system plan. The left navigation frame shows a single partition or the entire system. You can also choose just specific enclosures under the Hardware section. The file history is also viewable. The viewer also has a Print option and Show Comments / Hide Comments toggle, which is at the bottom of the viewer window.

If you are accessing the HMC from a PC browser, the print function is through the attached and network printers of the PC. If you are using the HMC terminal itself, the print function is through printers that are connected to the HMC or network printers to which the HMC has access.

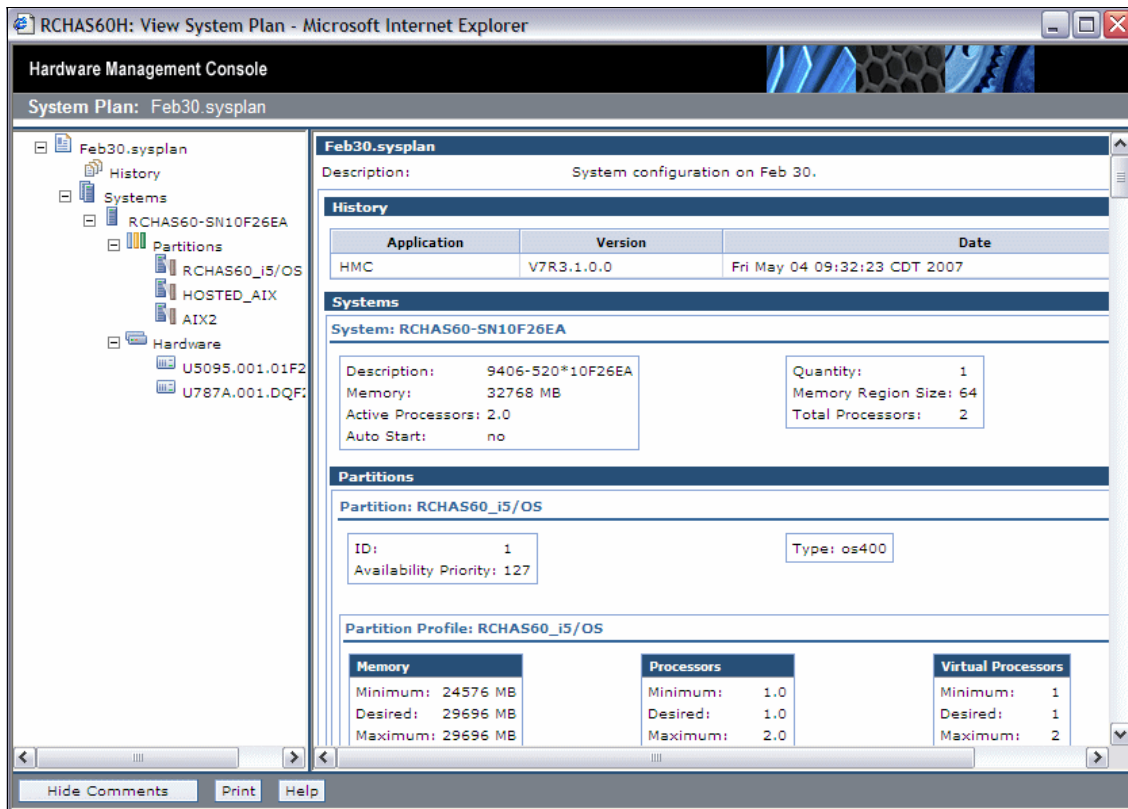


Figure 2-11 Viewing a system plan

The system plan section that is shown in Figure 2-12 shows the system's disk units. The controller for the disk unit displays in the table. This detail is obtained only if the IBM i5/OS™ operating system that is controlling the disk units is running. Linux and IBM AIX operating systems do not display disk controller information or location information.

Hardware Management Console
System Plan: RCHAS61.sysplan

Back
(rotate counterclockwise for standalone)

Backplane	Slot	Bus	Device Feature	Device Description	Device Serial #	Disk Controller	Order Status	Used by Partition / Profile
P2	D1		4326	35.16GB 15k RPM Disk Unit		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P2	D2		4326	35.16GB 15k RPM Disk Unit		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P2	D3		4326	35.16GB 15k RPM Disk Unit		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P2	D4		4326	35.16GB 15k RPM Disk Unit		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P3	D1		4326	35.16GB 15k RPM Disk Unit		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P3	D2		4326	35.16GB 15k RPM Disk Unit		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P3	D3		4326	35.16GB 15k RPM Disk Unit		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P3	D4		4326	35.16GB 15k RPM Disk Unit		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P4	D1		5754, 63A0	50GB 1/4 inch Cartridge Tape		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P4	D2		1994, 2640	IDE DVDROM		P1/T10	IBM	ITSO_i5/OS / ITSO_i5/OS
P4	D3							

Expand / Collapse System Image

P2 - D4
P2 - D3
P2 - D2
P2 - D1

P3 - D1
P3 - D2
P3 - D3
P3 - D4

P4 - D1
P4 - D2
P4 - D3

Figure 2-12 Viewing a system plan

2.1.6 Removing a system plan on the HMC

When you no longer need a system plan, you can remove the sysplan file easily from the HMC. Using Figure 2-3 on page 17 as a starting point, follow these steps:

1. Select the wanted plan in the main system plan management window.
2. Click **Remove System Plan** either from the Tasks menu or the content Tasks link in the lower portion of the window. A confirmation message displays asking if you are sure that you want to delete the file (Figure 2-13).

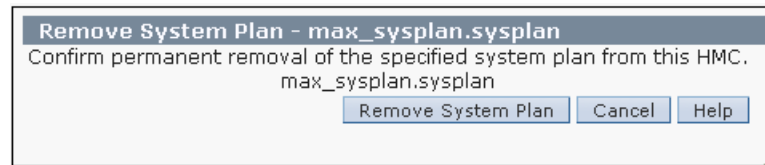


Figure 2-13 Confirm removal of system plan window

3. Click **Remove System Plan** to remove the selected sysplan file from the HMC.

2.2 System plans deployment

Since the publication of *LPAR Simplification Tools Handbook*, SG24-7231, from a general point of view, the deployment process has not changed much. Because of the updates of System Planning Tool Version 2 and HMC software, the details are not the same. The major improvements of the process are related to Virtual I/O Server implementation.

A summary of the deployment validation process is now described. We cover the new deployment wizard by using examples and provide details of the updates to the restricted shell CLI.

2.2.1 Deployment validation process

Before you deploy any system plan, it has to be validated. There are two steps in this validation process. First, the hardware is validated and then, if that validation is successful, the partition is validated.

This process is detailed in *LPAR Simplification Tools Handbook*, SG24-7231. You can refer to that book at any time. However, because it is fundamental to fully

understand how the validation process works, we now summarize the main concepts.

Hardware validation

When you run hardware validation, the HMC checks that any planned hardware exists on the managed server and that all the I/O processors and adapters are located physically in the planned slots. Hardware validation does not necessarily mean that an exact match should occur between the planned and the existing hardware. For example, you can plan on using fewer processors or memory than physically installed, or you can plan on not using all the physically installed I/O units.

Important: The HMC is not aware of the devices that are connected to the IOA. Therefore, there is no validation at a lower level than the IOA. When you use the System Plan Tool, specify devices such as disk drives, CD and DVD drives, and tape drives. The validation process *cannot* do any validation about these devices.

The validation includes all the following items:

- ▶ Server type, model, and processor feature: an exact match is required
- ▶ Number of processors: at least the planned number should exist
- ▶ Memory: at least the planned amount should exist
- ▶ Expansion units: all the expansion units in the plan should exist
- ▶ Slots: all the I/O processors and adapters in the plan should exist in a correct expansion or in the central electronics complex and should be at the same location
- ▶ Any serial number; an exact match is required

At this point, it is important to take actions to *avoid any ambiguity* about the expansion units or the processor enclosures central electronics complexes. You can have multiple central electronics complexes, for example on a 16-way model 570. In that case, you can have four central electronics complexes.

This ambiguity takes place when two or more installed expansion units or central electronics complexes have the same type and contain the same I/O processors and adapters in the same slot. You might plan a partition to use specific expansion units due, for example, to their physical location in the racks or on the floor or to specific disks drives that the HMC cannot see. The validation process allows such a system plan, but there is no guarantee for the deployment to allocate the right expansion to the partition.

The best way to eliminate expansion units or central electronics complexes ambiguity is to specify, in the system plan, their serial number.

Eliminate any hardware validation error, for the partition validation to start.

Partition validation

When you run partition validation, the HMC checks that any *existing* partition on the server exactly matches with one of the planned partitions.

The validation includes all the following items:

- ▶ Partition name
- ▶ Partition ID
- ▶ Name of the default profile
- ▶ Processing resources in the system plan
- ▶ Memory resources in the system plan
- ▶ Physical hardware in the system plan
- ▶ Virtual adapters, including slot ids and maximum adapters, in the system plan

If any of these items fail, the partition validation is unsuccessful, and the deployment fails. Some of the corrections to allow the deployment should be applied on the server. This process is the case for the name of the default profile, which cannot be changed in the System Planning Tool and is the same as the partition name. This process is also the case for some hardware features like the USB controller or the IDE CD controller that the HMC allows you to assign to an i5/OS partition (although it cannot use them), but the SPT does not.

2.2.2 Deploy a system plan by using the graphical wizard

To show the new deployment wizard that is related to V7R3 HMC software, we run three deployment examples:

- ▶ The first example fails because of hardware errors.
- ▶ The second example fails because of partitions errors.
- ▶ The third example is successful.

You can initiate deployment when you are using any right pane of the HMC by clicking **System Plans** on the left pane, as shown in Figure 2-14.

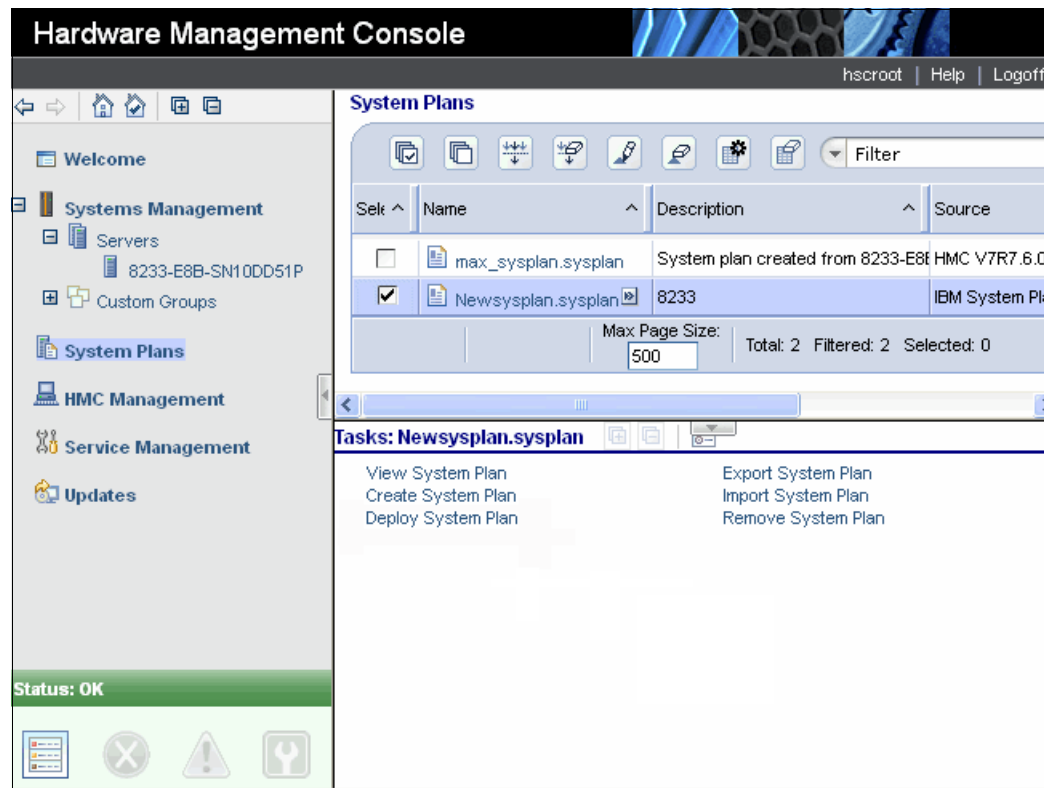


Figure 2-14 Launch deployment

To deploy a system, follow these steps:

1. On the list of the system plans, select the one that you want to deploy by clicking the check box to the left of the system plan.

There are three different ways to start the deployment of the selected system plan.

- Click the contextual menu immediately to the right of the system plan name and select **Deploy System Plan**.
 - Click **Deploy System Plan** in the bottom Tasks panel.
 - Click **Tasks** at the top of the System Plans list panel and select **Deploy System Plan**.
2. After the wizard is started, its Welcome page, which is shown in Figure 2-15, requests that you confirm the system plan to deploy and choose the managed server to be the target of the procedure. When your choices are done (in our example, we want to deploy the `marc.sysplan` file to the `RCHAS60-SN10F26EA` server), click **Next** to continue.

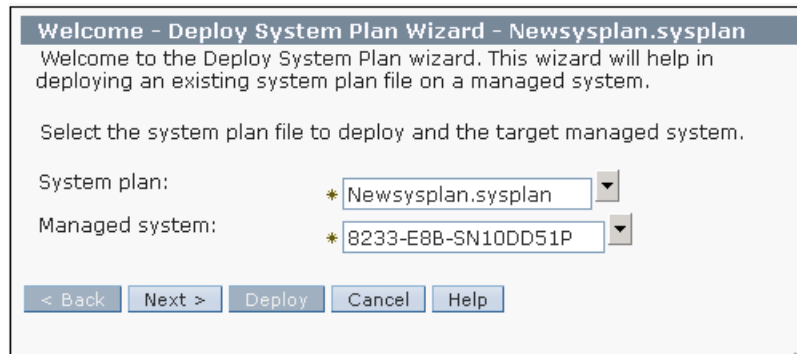


Figure 2-15 Confirm the deployment startup

3. Figure 2-16 shows you the validation progress. When the progress completes (Figure 2-17 on page 34), you can examine all the related messages. You can view the messages that are successful and the messages that are unsuccessful.




Figure 2-16 Deployment validation in progress

Figure 2-17 shows an example of a successful validation.

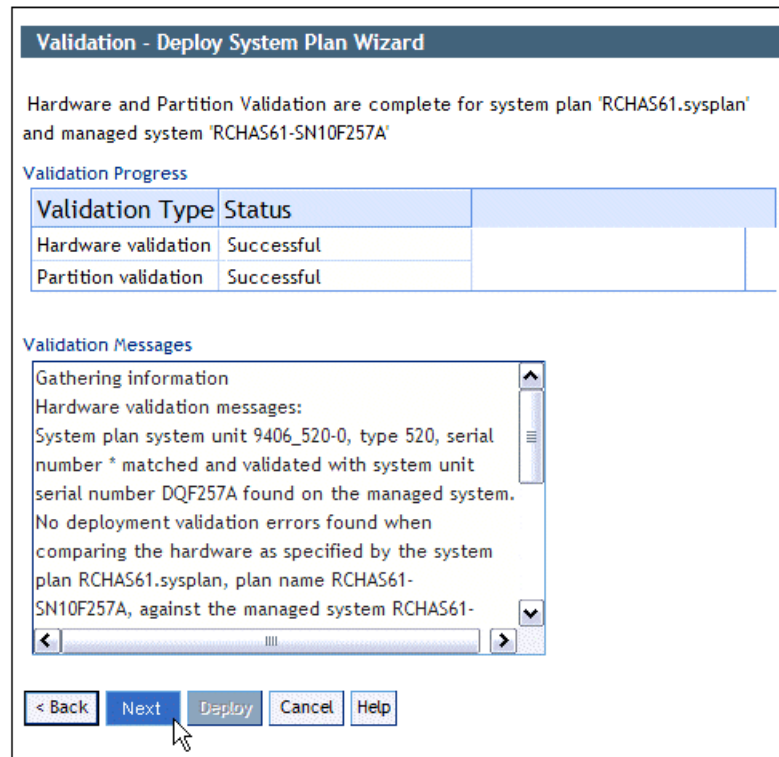


Figure 2-17 Example of successful validation

4. After the validation, the next panel is the starting point of the deployment. There are two portions in this panel:
 - a. The first contains the list of all the actions that are planned (see Figure 2-18 on page 35). In our example, the system plan is built with one i5/OS partition that hosts three Linux partitions.

Notice the *Partially* deployed status of the i5/OS partition. This status means that some items of the partition exist on the server and do not have to be deployed. Specifically, here, the partition and the profile are created (and the LPAR is running at the time of the window capture). These items are not to be deployed again.

Notice also the *Deploy* column. Each action of the plan can be cleared if you prefer to run it at a later deployment. Notice that, even if there are cleared items, the dependency is checked before running the deployment. So, you cannot, for example, deploy a hosted partition while the hosting one does not exist.

If you have to review the details of a specific action, you can select this one in the radio boxes of the Select column and click **Details**.

Partition Deployment - Deploy System Plan Wizard

Use this page to specify which partition plan actions to deploy on the managed system. Only the checked plan actions will be deployed. Select a row in the Partition Plan Actions table to view more details about the partition plan action.

Partition Plan Actions

Select	Dependency Hierarchy	Plan Action	Deploy	Status
<input type="radio"/>	1.1	Partition ITSO_i5/OS	<input checked="" type="checkbox"/>	Partially deployed
<input checked="" type="radio"/>	1.1.2	Partition vcx	<input checked="" type="checkbox"/>	
<input type="radio"/>	1.1.3	Partition LinuxVIO	<input checked="" type="checkbox"/>	
<input type="radio"/>	1.1.4	Partition IPT2	<input checked="" type="checkbox"/>	

[Details](#)

Partition Deployment Step Order

This table displays the partition deployment steps that will be performed based on the items checked in the Partition Plan Actions table.

Deployment Step
Partition vcx
Partition LinuxVIO
Partition IPT2
Partition Profile ITSO_i5/OS Virtual SCSI Adapters
Partition Profile ITSO_i5/OS Virtual Ethernet Adapters
Partition Profile ITSO_i5/OS Virtual Serial Adapters
Partition Profile vcx
Partition Profile LinuxVIO
Partition Profile IPT2
Partition Profile vcx Virtual SCSI Adapters

< Back Next > **Deploy** Cancel Help

Figure 2-18 Request the details of a specific action

- b. Click **Details**. The HMC links to the System Plan Viewer, which is restricted to the view that is associated with the selected action.

Authenticate again: You might have to authenticate again to the HMC when you access this option.

- c. The second portion contains the detailed list of all the steps that the HMC does to deploy the plan, as shown in Figure 2-19. You can scroll down and up to review all the steps.

Partition Deployment - Deploy System Plan Wizard

Use this page to specify which partition plan actions to deploy on the managed system. Only the checked plan actions will be deployed. Select a row in the Partition Plan Actions table to view more details about the partition plan action.

Partition Plan Actions

Select	Dependency Hierarchy	Plan Action	Deploy	Status
<input type="radio"/>	1.1	Partition ITSO_i5/OS	<input checked="" type="checkbox"/>	Partially deployed
<input type="radio"/>	1.1.2	Partition vcx	<input checked="" type="checkbox"/>	
<input type="radio"/>	1.1.3	Partition LinuxVIO	<input checked="" type="checkbox"/>	
<input type="radio"/>	1.1.4	Partition IPT2	<input checked="" type="checkbox"/>	

[Details](#)

Partition Deployment Step Order

This table displays the partition deployment steps that will be performed based on the items checked in the Partition Plan Actions table.

Deployment Step
Partition vcx
Partition LinuxVIO
Partition IPT2
Partition Profile ITSO_i5/OS Virtual SCSI Adapters
Partition Profile ITSO_i5/OS Virtual Ethernet Adapters
Partition Profile ITSO_i5/OS Virtual Serial Adapters
Partition Profile vcx
Partition Profile LinuxVIO
Partition Profile IPT2
Partition Profile vcx Virtual SCSI Adapters

Figure 2-19 List of the deployment steps

- d. When you are ready to deploy, after you review details of the actions, eventually deleting some of the actions and reviewing the deployment steps, as shown in the Figure 2-20, click **Deploy** to start the deployment.

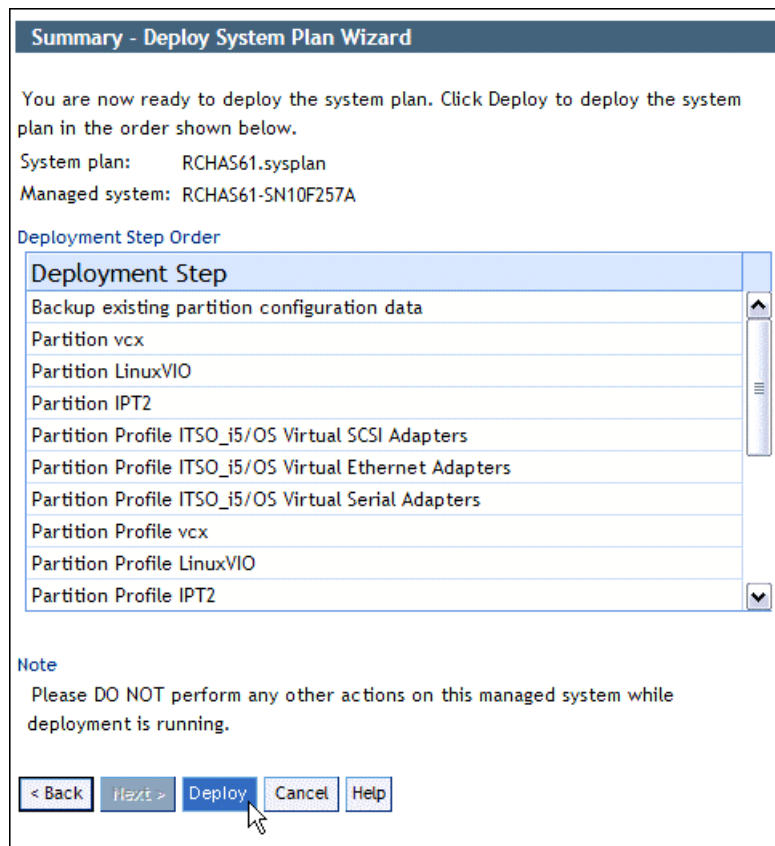


Figure 2-20 Partition deployment: Deploy System Plan Wizard

5. There is a summary of all the steps that the HMC performs. Click **Deploy** to start the operations. Notice the warning just above the icons at the bottom of the window. The deploy process time, depending on the complexity and the number of partitions, generally ranges from 5 - 20 minutes, and might be longer for specific deployments, when using Virtual I/O Server partitions for example.

While the deployment is running, each step is In progress then Successful, from the top to the bottom of the Deploy progress portion of the panel. Notice that you have to move the cursor of this list by yourself to see those steps that do not fit in the initial window.

Do a backup: The first step that you cannot disable is always to do a backup of the actual partitions configuration.

In the Messages portion of the panel, we see the detailed results of each step. In our example, when running the step “Partition Profile ITS0_i5/OS Virtual Serial Adapters“, you can see that there is no result yet. However, when running the previous step, we can see that one of the results was the following message:

```
Virtual Ethernet Adapter deployed for slot 5 on profile ITS0_i5/OS of
partition 1 on managed system RCHAS61-SN10F257A.
```

Notice that the display sort, for this list, is the opposite of the steps sort. You see the latest event first and the first one last, and therefore, there is no need to scroll down or up this list during the deployment.

When the deployment is complete, as shown in Figure 2-21, review all the steps and messages to ensure that everything is OK. Then, click **Close** to finish the session.

Deployment Progress - Deploy System Plan Wizard

Deployment status: Deployment complete

Note
Please DO NOT perform any other actions on this managed system while deployment is running.

System plan: RCHAS61.sysplan
Managed system: RCHAS61-SN10F257A

Deploy Progress

Status	Step
Successful	Partition Profile IPT2
Successful	Partition Profile vcx Virtual SCSI Adapters
Successful	Partition Profile vcx Virtual Ethernet Adapters
Successful	Partition Profile vcx Virtual Serial Adapters
Successful	Partition Profile LinuxVIO Virtual SCSI Adapters
Successful	Partition Profile LinuxVIO Virtual Ethernet Adapters
Successful	Partition Profile LinuxVIO Virtual Serial Adapters
Successful	Partition Profile IPT2 Virtual SCSI Adapters
Successful	Partition Profile IPT2 Virtual Ethernet Adapters
Successful	Partition Profile IPT2 Virtual Serial Adapters

Messages

- * Deployment complete
- * Virtual serial adapter not deployed. Virtual serial adapter already deployed for slot 0 on profile IPT2 of partition 4 on managed system RCHAS61-SN10F257A.
- * Virtual Ethernet adapter deployed for slot 2 on profile IPT2 of partition 4 on managed system RCHAS61-SN10F257A.
- * Virtual SCSI adapter deployed for slot 3 on profile IPT2 of partition 4 on managed system RCHAS61-SN10F257A.

Stop Close Help

Figure 2-21 Deployment complete

2.3 PowerVM Introduction

IBM PowerVM provides the industrial-strength virtualization solution for IBM Power Systems servers and blades. Based on more than a decade of evolution and innovation, PowerVM represents the state of the art in enterprise virtualization and is broadly deployed in production environments worldwide by most Power Systems owners. The IBM Power Systems family of servers includes proven1 workload consolidation platforms that help clients control costs while they improve overall performance, availability, and energy efficiency.

2.3.1 Processor Virtualization

PowerVM can help eliminate underutilized servers because it is designed to pool resources and optimize their use across multiple application environments and operating systems. Through advanced virtual machine (VM) capabilities, a single VM can act as a separate IBM AIX, IBM i, or Linux operating environment, by using dedicated or shared system resources. With shared resources, PowerVM can automatically adjust pooled processor, memory, or storage resources across multiple operating systems, borrowing capacity from idle VMs to handle high resource demands from other workloads.

Micro-Partitioning

PowerVM IBM Micro-Partitioning® supports multiple VMs per processor core and, depending upon the Power Systems model, can run up to 1000 VMs on a single server, each with its own processor, memory, and I/O resources. Processor resources can be assigned at a granularity of 1/100th of core. It enables up to 20 VMs per processor core.

Multiple Shared Processor Pools

Multiple Shared Processor Pools allows for the automatic nondisruptive balancing of processing power between VMs assigned to shared pools, resulting in increased throughput.

Shared Dedicated Capacity

Shared Dedicated Capacity allows for the “donation” of spare CPU cycles from dedicated processor VMs to a Shared Processor Pool. Because a dedicated VM maintains absolute priority for processor cycles, enabling this feature can increase system utilization without compromising the computing power for critical workloads.

2.3.2 Memory Virtualization

There are two kinds of technology for memory virtualization.

Active Memory Sharing

IBM Active Memory™ Sharing (AMS) is a technology that allows you to intelligently and dynamically reallocate memory from one VM to another for increased utilization, flexibility, and performance. AMS enables the sharing of a pool of physical memory among VMs on a server, helping to increase memory utilization and drive down system costs. The memory is dynamically allocated among the VMs as needed to optimize the overall physical memory usage in the pool.

For more information about Active Memory Sharing, see *PowerVM Virtualization Active Memory Sharing*, REDP-4470.

Active Memory Deduplication

Active Memory Deduplication is a powerful optimization feature that can be enabled when AMS is in use. This memory optimization intelligently detects and removes duplicate memory pages that are used between VMs and as a result reduces overall memory consumption.

2.3.3 I/O Virtualization

I/O virtualization is now described.

Virtual I/O Server

The *Virtual I/O Server (VIOS)* is a special purpose VM that can be used to virtualize I/O resources for AIX, i, and Linux client VMs. The Virtual I/O Server owns the resources that are shared with clients. A physical adapter that is assigned to the Virtual I/O Server can be shared by one or more other VMs. The Virtual I/O Server is designed to reduce costs by eliminating the need for dedicated network adapters, disk adapters and disk drives, and tape adapters and tape drives in each client VM. With Virtual I/O Server, client VMs can easily be created for test, development, or production purposes.

Shared Storage Pools

Shared Storage Pools (SSP) allow storage subsystems to be combined into a common pool of virtualized storage that can be shared by the Virtual I/O Server on multiple Power Systems servers. SSP support capabilities such as thin provisioning, whereby VM storage is dynamically allocated and released as required, to improve overall storage resource utilization.

N_Port ID Virtualization

N_Port ID Virtualization (NPIV) provides direct access to Fibre Channel Adapters from multiple VMs, simplifying the deployment and management of Fibre Channel SAN environments.

Live Partition Mobility

Live Partition Mobility supports the movement of a running AIX or Linux or IBM i VM from one Power Systems server to another without application downtime. This component helps to avoid application interruption for planned system maintenance, provisioning, and workload management. Live Partition Mobility can be used to simplify migration of operating environments to new servers temporarily or permanently.

For more information about Live Partition Mobility, see the IBM Redbooks publication, *IBM PowerVM Live Partition Mobility*, SG24-7460.

2.4 Reliability, Availability, Serviceability on the HMC

The HMC is used to do configuration, management, and maintenance activities on the system. The HMC is considered to be part of the server firmware, so it is one of the resources that contributes to the system stability and normal operation. The following topics are now described:

- ▶ Dual HMC and Redundancy
- ▶ RAID 1 protection

2.4.1 Dual HMC and Redundancy

You can configure a redundant HMC in a configuration in which dual HMC servers are connected to the service processors.

Using a redundant HMC configuration with your service processor setup requires a specific port configuration, as shown in Figure 2-22. In this configuration, each service processor connects to a network hub that is connected to each HMC. The network hubs that are connected to the service processors have to remain in the *power-on* state. Any 10/100BASE-T Ethernet switch or hub can be used to connect the server and HMC.

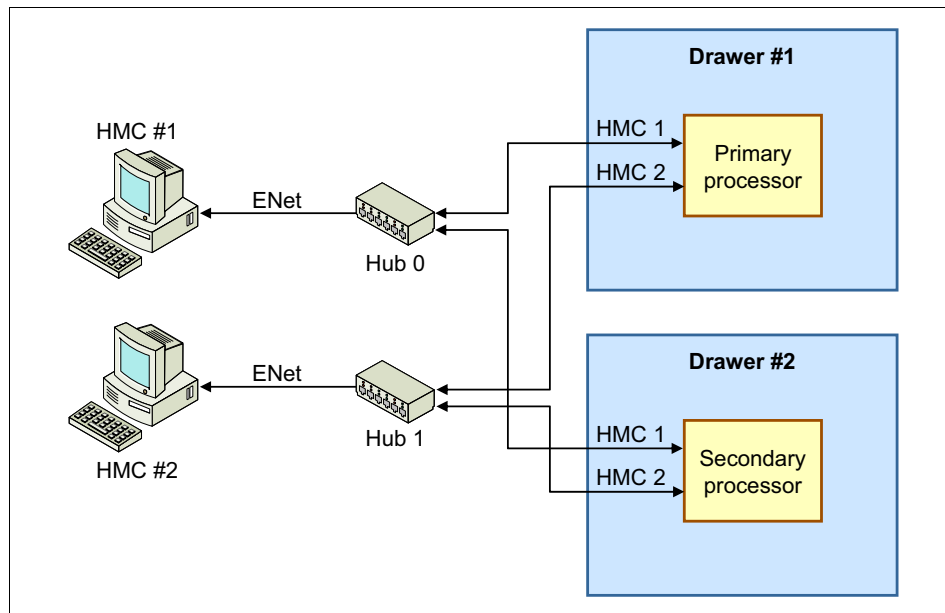


Figure 2-22 Dual HMC configuration on a private network

The HMC's first Ethernet port, *eth0*, should be configured to be a Dynamic Host Configuration Protocol (DHCP) server over a private network. By default, the flexible service processor (FSP) uses a DHCP client to request an IP address. This process occurs when power is applied to the server or the FSP is reset. The FSP has two default IP addresses: 192.168.2.147 on HMC1 port and 192.168.3.147 on HMC2 port. Always turn on the HMC first, then the server, during setup. This process is so that an IP address is available for the FSP, by which the HMC also discovers the servers on its private service network. See Figure 2-23.




Figure 2-23 Dual HMC physical connections

For information about how to configure an HMC, refer to Chapter 3, “Installing” on page 59.

Redundant remote HMC

A redundant remote HMC configuration is common. When clients have multiple sites or a disaster recovery site, they can use their second HMC in the configuration remotely over a switched network, as illustrated in Figure 2-24. The second HMC can be local, or it can be at a remote location. Use a different IP subnet for each HMC.

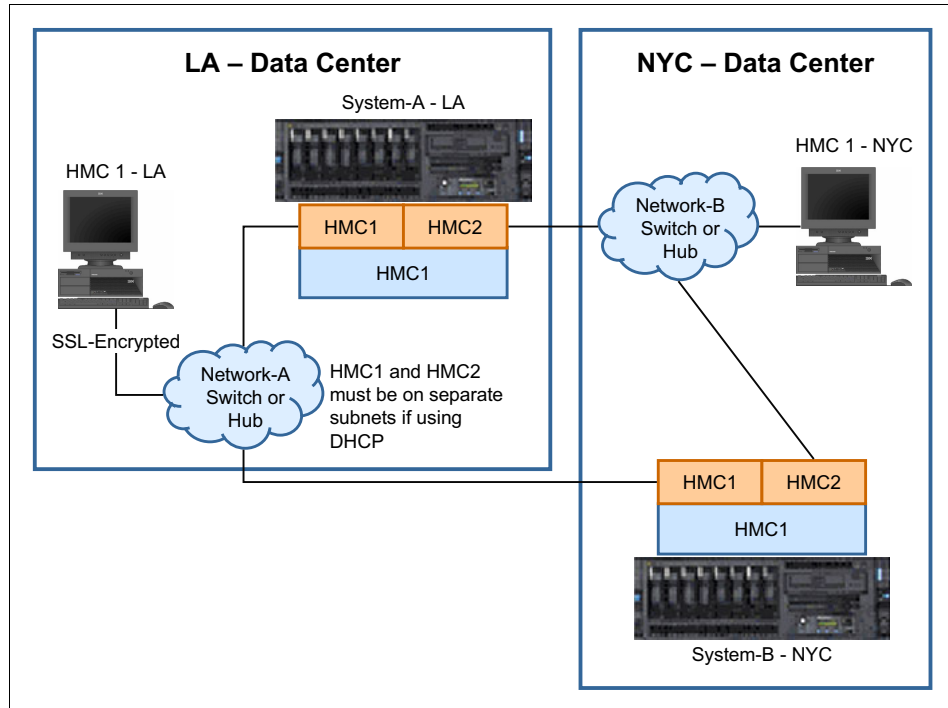


Figure 2-24 Example of redundant remote HMC

Redundant HMC configuration considerations

In a redundant HMC configuration, both HMCs are fully active and accessible always, enabling you to do management tasks from either HMC at any time. There is no primary or backup designation.

Consider the following points:

- ▶ Because authorized users can be defined independently for each HMC, determine whether the users of one HMC should be authorized on the other. If so, the user authorization should be set up separately on each HMC.
- ▶ Because both HMCs provide Service Focal Point™ and Service Agent functions, connect a modem and phone line to only one of the HMCs and

enable its Service Agent. To prevent redundant service calls, do not enable the Service Agent on both HMCs.

- ▶ Perform software maintenance separately on each HMC, at separate times, so that there is no interruption in accessing HMC function. This maintenance allows one HMC to run at the new fix level, although the other HMC can continue to run at the previous fix level. However, the preferred practice is to upgrade both HMCs to the same fix level as soon as possible.

The basic design of HMC eliminates the possible operation conflicts that are issued from two HMCs in the redundant HMC configuration. A locking mechanism that is provided by the service processor allows interoperation in a Parallel Environment. This process allows an HMC to temporarily take exclusive control of the interface, effectively locking out the other HMC. Usually, this locking is held only for the short duration of time that it takes to complete an operation, after which the interface is available for further commands.

Both HMCs are automatically notified of any changes that occur in the managed systems, so the results of commands that are issued by one HMC are visible in the other. For example, if you choose to activate a partition from one HMC, you observe the partition going to the Starting and Running states on both HMCs.

The locking between HMCs does not prevent users from running commands that might seem to be in conflict with each other. For example, if the user on one HMC activates a partition, and a short time later a user on the other HMC selects to power the system off, the system turns off. Effectively, any sequence of commands that you can do from a single HMC is also allowed when it comes from redundant HMCs.

For this reason, it is important to consider carefully how to use this redundant capability to avoid such conflicts. You might choose to use them in a primary and backup role, even though the HMCs are not restricted in that way. The interface locking between two HMCs is automatic, usually of short duration, and most console operations wait for the lock to release without requiring user intervention.

However, if one HMC experiences a problem while in the middle of an operation, it might be necessary to release the lock manually. HMC 2 can be used to disconnect HMC 1. When an HMC is disconnected, all locks that are owned by the HMC are reset. To do this process, any *hmcsuperadmin* user can run the Disconnect Another HMC GUI task on HMC 2 against HMC1. This task can be done only from the graphical interface. There is no corresponding command-line version of this task.

When you run two HMCs to the same server, also be careful with long running functions. Be careful because they might be impacted if they are not completed before an extra function is run on the second HMC.

With the previous considerations in mind, there are a number of good reasons to use the redundant HMC configuration. This list is not exhaustive:

- ▶ Redundancy of critical configuration information.
- ▶ Ability to apply maintenance to an HMC while the other is available for production management functions.
- ▶ Reduced risk of no HMC available.
- ▶ Knowing that a long running command is running against one system, being able to use the second HMC to do functions on another system without waiting.

2.4.2 RAID 1 Protection

Redundant Array of Independent Disks (RAID) is an industry-wide implementation of methods to store data of an HMC on multiple physical disks to enhance the availability of that data. The HMC V7R760 is designed to deliver support for RAID 1.

RAID 1

HMC now offers a high-availability feature. The new 7042-CR7, by default, includes two hard disk drives with RAID 1 configured. RAID 1 is also offered on both the 7042-CR6 and the 7042-CR7 as an MES upgrade option. It uses data mirroring. Two physical drives are combined into an array, and data is striped across the array. The first half of a stripe is the original data. The second half is a mirror of the data, but it is written to the other drive in the RAID 1 array. RAID 1 requires two physical drives, enabling data redundancy.

2.4.3 RAID 1 conversion

This part covers the RAID1 conversion.

Preparing for the RAID 1 conversion

In preparation for the conversion, close all serviceable events and back up your existing HMC data, as follows.

Backing up your existing HMC data

Save upgrade data to use for a restore operation, and create a profile backup as an extra precaution for saving your configuration.

Use the following procedure:

1. Expand **Systems Management** and select **Servers**.

2. Select the first managed system by selecting the check box next to the server name.
3. Under Tasks, expand **Configuration**, and then expand **Manage Partition Data**.
4. Select **Backup**.
5. In the Profile Data Backup pane, type a name for the backup file, and click **OK**. A panel opens to confirm that the profile data was backed up.
6. Click **OK**.
7. Repeat steps 1 - 6 for each additional system that is managed by this HMC. After you have a partition data backup for each managed system, proceed to step 8.
8. Prepare media to use for the Save Upgrade Data task. You need either a DVD-RAM, or a USB flash drive (available if HMC V7.3.5 or later).
9. Select **HMC Management**.
10. Select **Save Upgrade Data**.
11. Select media type, click **Next**, and then click **Finish**.
12. Wait for the Save Upgrade Data task to complete.
 - If the task fails, contact your next level of support.
 - If the task successful, use the following optional procedure to verify the data:
 - a. Open a restricted shell terminal as follows:
 - i. Select **HMC Management**.
 - ii. Select **Open restricted Shell Terminal**.
 - b. Mount the media:
 - i. Run the `lsmediadev` command to determine the mount point.

In the following example, the CD/DVD mount point is `/media/cdrom`; the USB flash drive is `/media/sda1`:

```
device=/dev/cdrom,mount_point=/media/cdrom,type=1,description=CD/DVD
device=/dev/fd0,mount_point=/media/floppy,type=2,description=internal
diskette drive
device=/dev/hda,type=6,description=internal hard disk drive
device=/dev/sda1,mount_point=/media/sda1,type=3,description=USB flash
```

memory device

- ii. Run the **mount <mount point>** command.
- c. Run the **ls -l <mount point>** command by using the same mount point as in step b.
- d. Verify that there are five *.tar* files with the following names:

ACMSaveData.tar
RSCTSaveUpgrade.tar
SaveCCFWUpgradeData.tar
SaveHSCSystemUpgradeData.tar
SaveProfileDataUpgrade.tar

If any of these *.tar* files are missing, contact your next level of support.

- e. Issue the **umount <mount point>** command by using the same mount point as in step b.

Save Upgrade Data task: The Save Upgrade Data task saves to the hard disk drive on the HMC and the DVD-RAM or USB flash drive. The Save Upgrade Data task on the media is used only to recover from an unexpected error under the direction of support personnel.

- f. Remove the DVD-RAM or USB flash drive, and insert V7.7.6 recovery media MH01326 volume 1 into the DVD-RAM drive.

You can now proceed to installing the additional drive and setting up the Integrated Mirror (RAID 1) volume, described in “Configuring RAID 1 on the 7042-CR6”.

Configuring RAID 1 on the 7042-CR6

This section applies to HMC model 7042-CR6 only. Instructions follow for installing a second physical drive, deleting an Integrated Striping (IS) volume (RAID 0), and setting up the Integrated Management (IM) volume (RAID 1).

Installing the second physical drive

Before you set up your RAID 1 volume on the HMC, do the miscellaneous equipment specification (MES) upgrade to add the second disk for mirroring. Any model 7042-CR6 that was previously used as an SDMC already has two disks that are installed. If your 7042-CR6 was an SDMC appliance that was converted to an HMC appliance, skip this task and go to “Deleting an Integrated Striping (IS) volume (RAID 0) for previous SDMC installations” on the next.

To install the physical drive, complete the following steps:

1. From the HMC GUI, click **HMC Management**, click **Shutdown or Restart**, select **Shutdown HMC**, and then click **OK**.
2. Remove the filler panel from the empty drive bay, for example, <bay number 0?>.
3. Touch the static-protective package that contains the drive to any unpainted metal surface on the server. Then, remove the drive from the package and place it on a static-protective surface.
4. Install the hard disk drive in <drive bay 1> drive bay:
 - a. Make sure that the tray handle is in the open (unlocked) position.
 - b. Align the drive assembly with the guide rails in the bay.
 - c. Gently push the drive assembly into the bay until the drive stops.

Deleting an Integrated Striping (IS) volume (RAID0) for previous SDMC installations

If your 7042-CR6 was converted from Systems Director Management Console (SDMC), it already has two hard disk drives configured as an IS or RAID0 volume. We must delete this volume before we configure the Integrated Management (IM) or RAID 1 volume.

To delete the IM volume, complete the following steps:

1. Power on the HMC appliance.
2. After the selection pane opens during the restart operation, press F1 to go to Set up.
3. Select **System Settings**.
4. Select **Adapters and UEFI Drivers**.

Compile the List of Drivers menu: If this menu is shown, press the Enter key.

5. Select the **LSI Logic Fusion MPT SAS Driver**.
6. From the LSI Logic MPT Setup Utility menu, select the **SAS1064e** driver.
7. Select **RAID Properties** in the SAS1064E Adapter Properties pane.

8. Select **Manage Array** in the View Array pane.
9. Select **Delete Array** in the Manage Array pane.
10. Select **Y** to delete the array and exit to the Adapter Properties pane.

Setting up the Integrated Management (IM) volume (RAID 1)

This task guides you through the HMC BIOS to configure the LSI Logic Fusion Adapter for RAID 1. If you just completed deleting an Integrated Striping volume for previous SDMC installations, you might be in the Adapter Properties pane of the LSI Logic MPT Setup Utility and can skip to step 7.

1. Power on the HMC appliance.
2. After the selection pane opens during the restart operation, press F1 to go to Set up.
3. Select System Settings.
4. Select **Adapters and UEFI Drivers**.

Compile the List of Drivers menu: If this menu is shown, press the Enter key.

5. Select the **LSI Logic Fusion MPT SAS Driver**.
6. From the LSI Logic MPT Setup Utility menu, select the **SAS1064e** driver.
7. Select **RAID Properties** in the SAS1064E Adapter Properties pane.
8. Select **Create IM Volume** on the Select New Array Type pane.

Ensure correct installation: This step lists the two hard disk drives (HDDs) that are installed in the HMC appliance. Check for the HDD listings and ensure that the two HDDs are correctly installed.

9. Use the arrow keys to move to the **RAID DISK** column, and press the Spacebar to toggle the value to **Yes** for each HDD that is listed. Set the RAID DISK column value to **Yes**.

Notes:

- ▶ The order in which you carry out the next step dictates the primary and secondary drives in the RAID volume.
- ▶ When you toggle the first selection, you are asked to delete or migrate existing data. The SAS1064E adapter does not support migration, so select **D – Overwrite existing data here**.

10. From the Create and Save New Array menu, select **Save Changes**.

11. Exit the menu and press Enter.

12. Press Esc to exit **LSI MPT setup utility**.

13. Press Esc to exit the previous menu.

14. Select **Exit the Configuration Utility** and **Restart**.

15. Press Enter when prompted to stop the controller: The action that is required is Stop Controller YES (Enter) or NO (Esc).

16. Press Esc three times to exit setup.

17. Press the Y key to exit setup completely.

Configuring RAID 1 on the 7042-CR7

This section applies only to HMC model 7042-CR7. Instructions follow for installing a second physical drive, deleting the pre-existing RAID 0 virtual drive, and setting up the RAID 1 virtual drive.

Installing the second physical drive

Before you set up your RAID 1 volume on the HMC, do the MES upgrade to add the second disk for mirroring.

To install the physical drive, complete the following steps:

1. From the HMC GUI, click **HMC Management**, click **Shutdown or Restart**, select the **Shutdown HMC** button, and then click **OK**.
2. Remove the filler panel from the empty drive bay <bay number 1>.

3. Touch the static-protective package that contains the drive to any unpainted metal surface on the server. Then, remove the drive from the package and place it on a static-protective surface.
4. Install the hard disk drive in <bay number 1> drive bay:
 - a. Make sure that the tray handle is in the open (unlocked) position.
 - b. Align the drive assembly with the guide rails in the bay.
 - c. Gently push the drive assembly into the bay until the drive stops.

Deleting the pre-existing RAID0 virtual drive

If your 7042-CR7 was previously installed with HMC code, it already has a defined virtual drive set up as a RAID0 volume. We must delete this volume before creating the RAID 1 volume.

To delete the original virtual drive, complete the following steps:

1. Power on the HMC appliance.
2. After the selection pane opens during the restart operation, press F1 to go to the Setup Utility.
3. Select **System Settings**.
4. Select **Storage**.
5. Select the **LSI MegaRAID <ServeRAID M5110> Configuration Utility**.
6. Select **Virtual Drive Management**.
7. Select **Select Virtual Drive Operations**.
8. Select **Delete Virtual Drive**.
9. Select **Confirm** and press the Spacebar to toggle the check box.
10. Select **Yes**.
11. Select **OK**.
12. Exit the Select Virtual Drive Operations pane (press Esc).
13. Exit the Virtual Drive Management pane (press Esc).

Setting up the RAID 1 virtual drive

This task guides you through the creation of a RAID 1 virtual drive. If you just completed deleting the pre-existing RAID 0 virtual drive, you might already be

viewing the Configuration Options pane of the LSI MegaRAID Configuration Utility and can skip to step 7.

1. Power on the HMC appliance.
2. After the selection pane opens during the restart operation, press F1 to go to Setup Utility.
3. Select **System Settings**.
4. Select **Storage**.
5. Select the **LSI MegaRAID <ServeRAID M5110> Configuration Utility**.
6. Select **Virtual Drive Management**.
7. Select **Create Configuration**.
8. Select **Select RAID Level** and select **RAID 1**, as shown in Figure 2-25.




Figure 2-25 Create Configuration pane: Selecting the RAID level

9. Select **Select Drives**.

10. Verify that both drives are listed and select **Check All**, as shown in Figure 2-26.




Figure 2-26 Select Drives pane: Check All

11. Select **Apply Changes**.

12. Select **OK** at the success pane.
13. Select **Save Configuration** as shown in Figure 2-27.




Figure 2-27 Creation Configuration pane: Saving the configuration

14. At the Warning pane, select **Confirm** and press the Spacebar to toggle the check box.

15. Select **Yes**.

16. Select **OK** at the Success pane.

A message indicates that more virtual drives cannot be created because of insufficient capacity. The reason is because all available drives are now allocated to the new virtual drive.

17. Press the Esc key to close the message.

The Virtual Drive Management pane does not refresh after the creation of the virtual drive, and, as a result, all operations are displayed but are not available (disabled). The next step addresses this issue.

18. Exit back to the Configuration Options pane by pressing the Esc key.

19. Select **Virtual Drive Management**.

20. Select **Select Virtual Drive Operations**.

21. Verify that the Virtual Drive Operation is set to Fast Initialization and select Start Operation, as shown in Figure 2-28.



Figure 2-28 Select Virtual Drive Operation pane: Start Operation

22. At the Warning pane, select **Confirm** and press the Spacebar to toggle the check box.

23. Select **Yes**.

24. Select **OK** at the Success pane.

25. The operation completes quickly, and you can exit the setup utility by pressing the Esc key until you reach the exit confirmation panel.

26. Press the Y key to exit the Setup Utility.

27. Proceed to the next section.



Installing

To set up the Hardware Management Console (HMC), complete the following groups of tasks:

- ▶ Cabling the HMC to the managed server
- ▶ Gathering configuration settings for your installation and configuring the HMC
- ▶ Connecting managed systems to the HMC

The HMC can be a stand-alone HMC or an HMC that you plan to install in a rack.

An overview of these tasks is provided.

3.1 Installation of HMC

You might learn how to cable and configure the HMC, including installing the HMC into a rack and configuring network connections and security.

3.1.1 Cabling the HMC

Attention: Do not plug the power cords into an electrical outlet until you are instructed to do so.

This section describes how to connect the HMC cables, connect the Ethernet cable, and connect the HMC to a power source. You can use these instructions to help you cable your rack-mounted or stand-alone HMC.

1. Use the specifications for the HMC to help ensure that you position the HMC in the correct location. HMC specifications provide detailed information for your HMC, including dimensions, electrical power, temperature, environment, and service clearances.

Choose from the following options:

- a. If you are installing a rack-mounted HMC, continue with step 2.
- b. If you are installing a stand-alone HMC, skip to step 3 on page 62.

2. To install a rack-mounted HMC:

a. First, identify the location of the connectors:

- A rack-mounted HMC 7042-CR7 (Figure 3-1)




Figure 3-1 Rear view of a rack-mounted HMC 7042-CR7

- A rack-mounted HMC 7042-CR5 and 7042-CR6 (Figure 3-2)




Figure 3-2 Rear view of a rack-mounted HMC 7042-CR5 and CR6

b. Next, install the HMC into a rack.

c. After the HMC is installed into a rack, connect the monitor, keyboard, and mouse.

For connection to an HMC, connect the keyboard, display, and mouse by using the USB conversion option cable.

After you complete these steps, skip to step 4.

3. If you are installing a stand-alone HMC:
 - a. First, identify the location of the connectors:
 - A stand-alone HMC 7042-C06 and 7042-C07 (Figure 3-3)
 - A stand-alone HMC 7042-C08 (Figure 3-4 on page 63)




Figure 3-3 Rear view of a stand-alone HMC 7042-C06 and 7042-C07

Port	Description
Video	The Video port is used as monitor connection.
PS/2 Keyboard or Mouse	PS/2 keyboard or mouse can attach.
Serial 1	Serial 1 is supported for an external modem.
USB	USB keyboard, mouse, USB drive, or USB diskette drive.
Ethernet	Ethernet port is used as HMC primary network connection

The HMC does not support the use of parallel, Serial 2, or Systems Management Ethernet ports.

Figure 3-4 shows a stand-alone HMC 7042-C08.




Figure 3-4 Rear view of a stand-alone HMC 7042-C08

- b. Attach the monitor cable to the monitor connector and tighten the screws. Attach the power cord to the monitor. Ensure that the voltage selection switch on the HMC is set to the voltage that is used in your area.
- c. Plug the power cord into the HMC, and then connect the keyboard and mouse. For USB connections, connect the keyboard and mouse to USB ports on the HMC. You can connect the keyboard and mouse to the USB ports on the front or back panels. For PS/2 connections, connect the keyboard and mouse to their connector on the back panel of the HMC.
4. Connect the modem. During the installation and configuration of the HMC, the modem might dial out automatically as the HMC follows routine call-out procedures.
5. Connect the Ethernet (crossover) cable from the HMC to the managed server. Your HMC should be connected to the managed server in a private service DHCP network. Your Ethernet connection to the managed server should be made by using the Ethernet port that is defined as eth0 on your HMC. If you

did not install any more Ethernet adapters in the PCI slots on your HMC, use the primary integrated Ethernet port.

6. If you use an external modem, plug the modem power supply cord into the HMC modem.
7. Plug the power cords from the monitor, HMC, and HMC external modem into electrical outlets.

Attention: Do *not* connect the managed system to a power source now.

3.1.2 HMC Guided Setup wizard

This section provides a step-by-step guide to setting up the HMC using the HMC Guided Setup wizard. Make sure that you have completed the HMC Guided Setup wizard checklist before continuing with this section.

Set up language and locale

Before starting the HMC Guided Setup wizard, set up your language in one of the 16 supported languages or the default of US English. After you select a language, you can select a locale that is associated with that language. The language and locale settings determine the language, the character set, and other specific to the country or region such as the formats for date, time, numbers, and monetary units.

Launching the HMC Guided Setup wizard

The HMC Guided Setup wizard starts automatically when you first start the HMC. You can also start the wizard from the Welcome panel, click Guided Setup wizard.

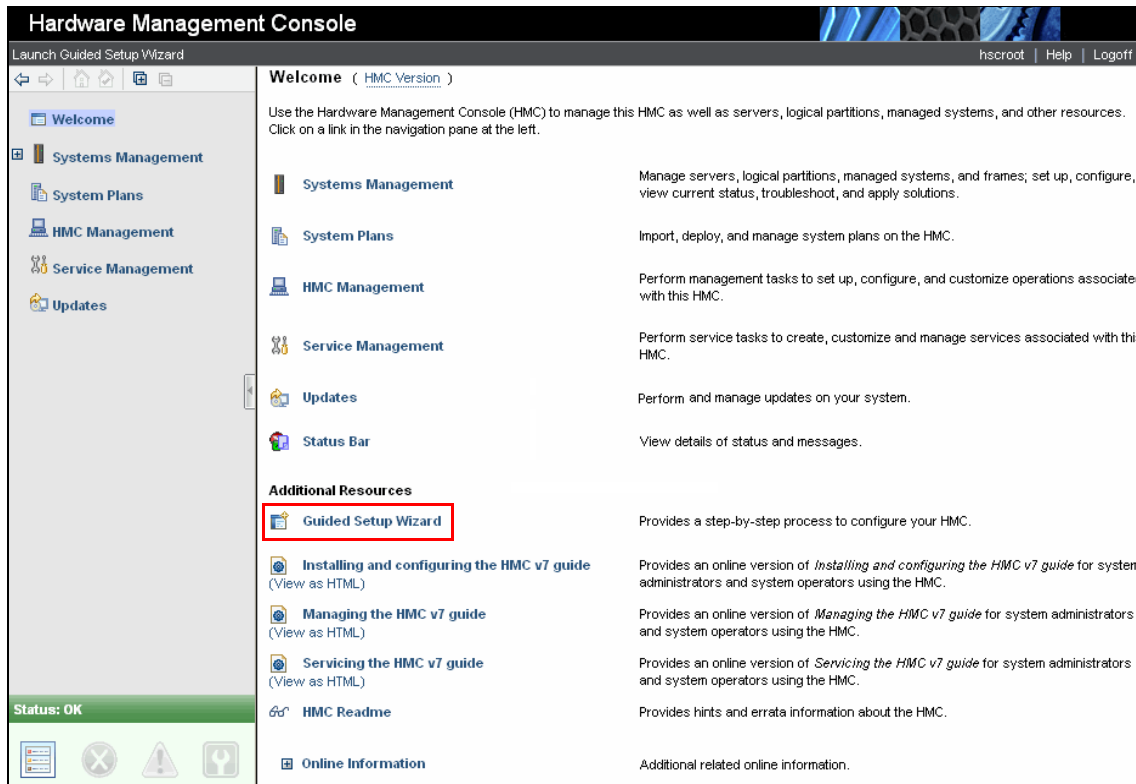


Figure 3-5 Welcome panel: Launch Guided Setup Wizard

1. The Launch Guided Setup Wizard: Welcome page opens. Click **Next** to continue with the Guided Setup wizard.
2. In the Launch Guided Setup Wizard: Change HMC Date and Time page, enter the correct date and time and time zone for your location (Figure 3-6 on page 66). Click **Next** to continue with the Guided Setup wizard (Figure 3-7 on page 67).
3. The Launch Guided Setup Wizard: Change *hscroot* Password window display. Enter the new *hscroot* password that you want (normally the default password

ID *abc123*). Use the *hscroot* user ID to access the user interface and to manage the HMC.

Use these following password rules:

- The password must contain at least seven characters.
- The characters must be standard 7-bit ASCII characters.
- Passwords can include special characters, but begin passwords with an alphanumeric character.

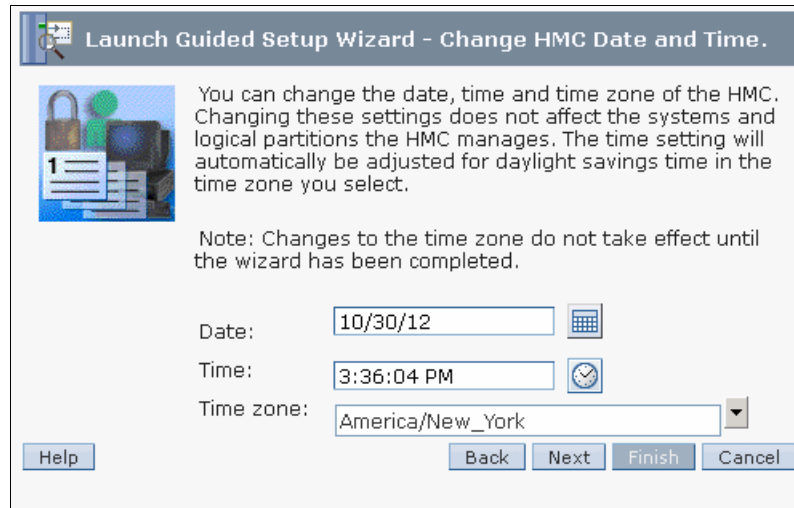


Figure 3-6 Change Date and Time

4. The Launch Guided Setup Wizard: Change root password panel display. Enter the new *root* password that you want (normally the default password ID *passwd0rd* where *0* is the number zero). Use *root* password rules same as *hscroot* password rules before. A service provider uses the *root* user ID to do maintenance on the server.

Click **Next** to continue with the Guided Setup wizard.

5. The Launch Guided Setup Wizard: Create Additional HMC Users panel display. You can optionally create new HMC users. Click **Next** to continue.
6. This completes the first part of the Guide Setup wizard.

The next step configures the HMC network settings. You might have to discuss this step with your network administrator for your HMC environment.

Click **Next** to continue with the Guided Setup wizard.

7. The Launch Guided Setup Wizard - Configure HMC Network Settings panel is displayed, as shown in Figure 3-7. In our example, we see two LAN adapters available (*eth0* and *eth1*), although you might see only one adapter in your HMC system.

We show you how to configure the HMC for both a private network and an open network. We use the first Ethernet network card (*eth0*) for a private network, then return to this panel again to configure *eth1* for an open network. We use the private network to connect the HMC with our managed systems and other HMCs. We use the second Ethernet card (*eth1*) for an open network.

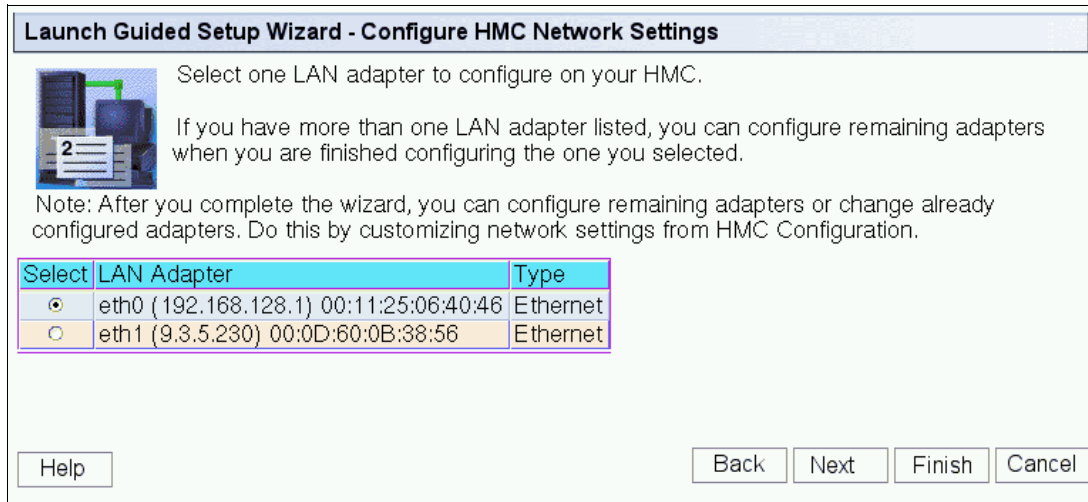


Figure 3-7 Configure HMC Network Settings

Select LAN adapter *eth0* to be configured first, then click **Next** to continue with the Guided Setup wizard.

8. Configure *eth0* Media Speed panel displays. You can choose the LAN adapter speed at **Autodetection** for initial setup, or you can set the adapter speed if you know the information.

Click **Next** to continue with the Guided Setup wizard.

9. The Launch Guided Setup Wizard - Configure eth0 Private Network panel is displayed, as shown in Figure 3-8. As mentioned previously, we set the LAN adapter eth0 as a private network to connect to our managed systems.

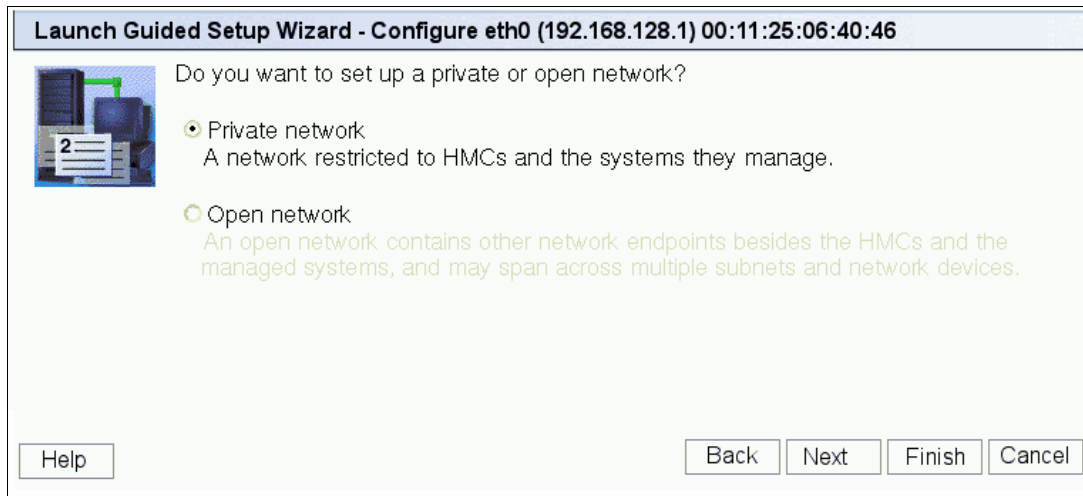


Figure 3-8 Configure eth0 Private Network

In this example, we select the Private Network. Then, click **Next** to continue.

10. The Launch Guided Setup Wizard - Configure eth0 DHCP panel is displayed (see Figure 3-9). We have to define the HMC as a DHCP server so our managed system is assigned an IP address by the DHCP server.

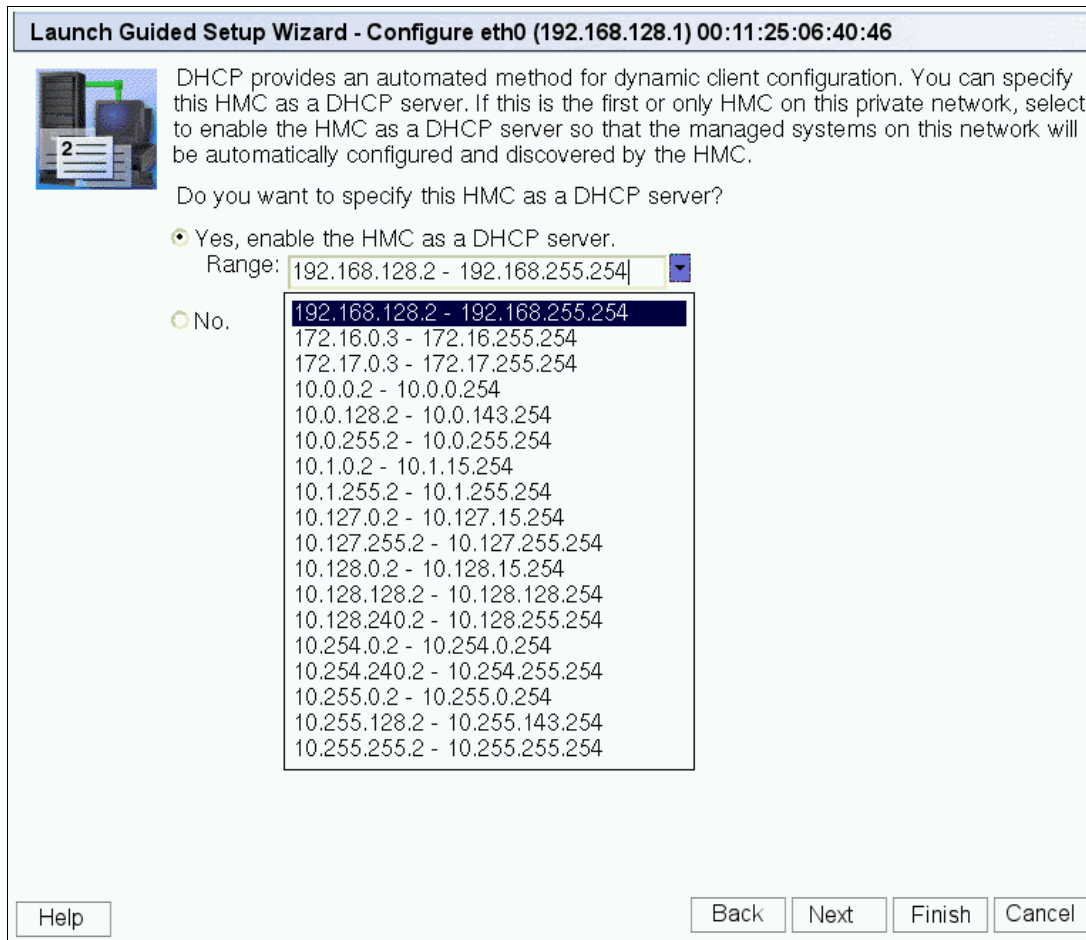


Figure 3-9 Configure eth0 DHCP

The HMC becomes the DHCP server to all clients in a private network. These clients are our managed systems and other HMCs. You can configure the HMC to select one of several different IP address ranges to use for this DHCP server. This flexibility is so that the addresses provided to the managed systems do not conflict with addresses used on other networks to which the HMC is connected.

We have some standard nonroutable IP address ranges that are assigned to its clients. The following ranges can be selected:

- 192.168.128.2 - 192.168.255.254
- 172.16.0.3 - 172.16.255.254
- 172.17.0.3 - 172.17.255.254
- 10.0.0.2 - 10.0.0.254
- 10.0.128.2 - 10.0.143.254
- 10.0.255.2 - 10.0.255.254
- 10.1.0.2 - 10.1.15.254
- 10.1.255.2 - 10.1.255.254
- 10.127.0.2 - 10.127.15.254
- 10.127.255.2 - 10.127.255.254
- 10.128.0.2 - 10.128.15.254
- 10.128.128.2 - 10.128.128.254
- 10.128.240.2 - 10.128.255.254
- 10.254.0.2 - 10.254.0.254
- 10.254.240.2 - 10.254.255.254
- 10.255.0.2 - 10.255.0.254
- 10.255.128.2 - 10.255.143.254
- 10.255.255.2 - 10.255.255.254

The HMC LAN adapter eth0 would be assigned one before the first IP address out of the range selected. In our example, we select the 192.168.0.2 to 192.168.255.254 range, so our HMC is given an IP address 192.168.0.1. Any other client (HMC or managed system) is given an address from this range.

The DHCP server in the HMC uses automatic allocation, which means that each managed system is reassigned the same IP address each time it is started. The DHCP server uses each client's built-in Media Access Control (MAC) address to ensure that it reassigns each client with the same IP address as before. When a managed system starts, it tries to contact the DHCP server to obtain its IP address. If the managed system is unable to contact the HMC DHCP server, the managed system uses its last given IP address.

We set the IP address range 192.168.0.2 to 192.168.255.254 and click **Next** to continue.

11. The Launch Guided Setup Wizard - Configure HMC Network Settings panel is displayed (Figure 3-10). This step completes the network configuration for LAN adapter eth0 as a private network. We can proceed with network configuration for LAN adapter eth1 as an open network.

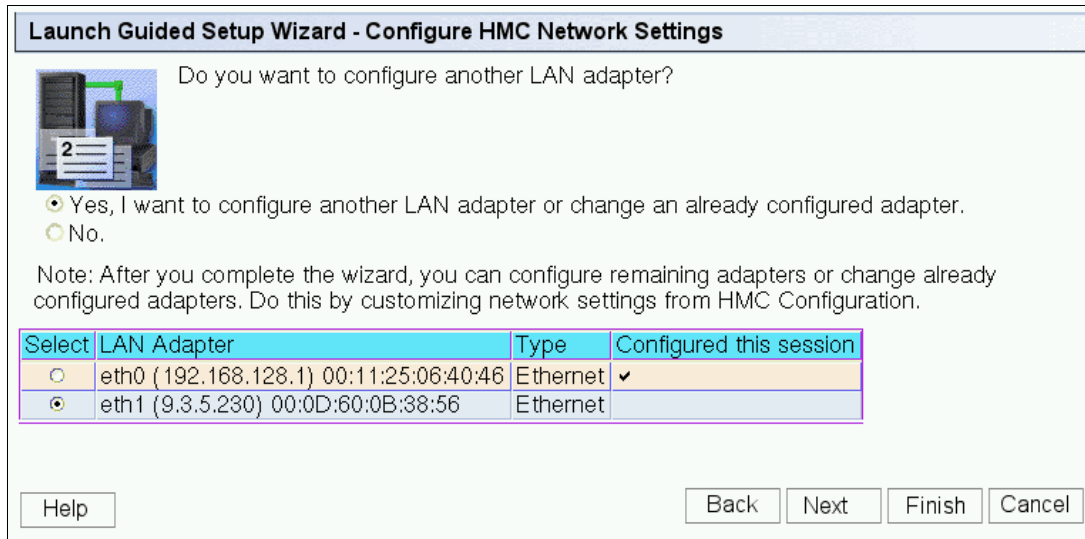


Figure 3-10 Launch Guided Setup Wizard - Configure HMC Network Settings

- Select the **Yes** option and LAN adapter eth1. Click **Next** to continue with the Guided Setup wizard.
12. The Launch Guided Setup Wizard - Configure eth1 Media Speed. As the eth0 configuration before, you can leave the LAN adapter speed at **Autodetection** for initial setup, or you can set the LAN adapter speed if you know the information.
- Click **Next** to continue with the Guided Setup wizard.
13. Configure eth1 **Open Network**. As mentioned previously, we select Open network for eth1 and Click **Next** to continue.

14. The Launch Guided Setup Wizard - Configure eth1 IP assignment panel is shown in Figure 3-11. We can configure interface eth1 to obtain an IP automatically from your DHCP server or to use a fixed IP address.

Launch Guided Setup Wizard - Configure eth1 (9.3.5.230) 00:0D:60:0B:38:56

You can have IP addresses assigned to the HMC automatically or you can specify the IP addresses to use.

Do you want to obtain an IP address automatically?

Yes, obtain an IP address automatically.

No. Use the specified address.

TCP/IP interface address:

TCP/IP interface network mask:

Figure 3-11 Configure eth1 IP assignment

In our example, we configure the interface eth1 by using fixed IP address, 9.3.5.20 with network mask 255.255.254.0. Click **Next** to continue with the Guided Setup wizard.

15. IPv6 Setting - You can optionally configure IPv6 environment.

Autoconfigure IP addresses: If this option is selected, the autoconfiguration process includes creating a link-local address and verifying its uniqueness on a link, determining what information should be autoconfigured (addresses, other information, or both). And in the case of addresses, whether they should be obtained through the stateless mechanism, the stateful mechanism, or both.

Use privacy extensions for autoconfiguration: This option causes nodes to generate global-scope addresses from interface identifiers that change over time.

IPv6 address entry: Input field for the 128 bit long IPv6 address. IPv6 addresses are normally written as eight groups of four hexadecimal digits.

For example, fe80:0:0:0:204:acff:feab:b811 is a valid IPv6 address. If one or more four-digit groups is 0000, the zeros might be omitted and replaced with two colons (::).

Prefix length entry: The prefix-length variable is a decimal value that indicates the number of contiguous, higher-order bits of the address that make up the network portion of the address.

Click **Next** to continue.

16. The Launch Guided Setup Wizard - Configure HMC Firewall for eth1 panel is now displayed, as shown in Figure 3-12. Commonly, there is a firewall that controls access from outside to your network. Because the HMC is connected to an open network, we can also restrict outside access to the HMC by using a built-in firewall in HMC. There are some applications that run on the HMC that can be secured from unauthorized outside access.

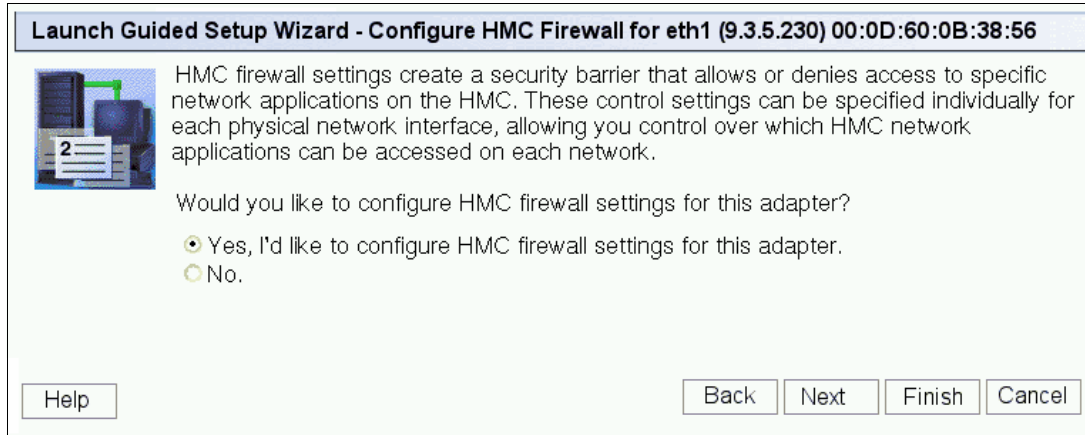


Figure 3-12 Launch Guided Setup Wizard - Configure HMC Firewall for eth1

We select **Yes** to configure HMC firewall settings and click **Next** to continue with the Guided Setup wizard.

17. The Launch Guided Setup Wizard - Configure HMC Firewall for eth1 panel displays then (see Figure 3-13). The top panel lists all of the available applications that are on the HMC. The bottom panel lists all applications that are available to the open network through the HMC firewall.

You can allow an application to pass through the firewall by selecting them from the top panel then clicking **Allow Incoming** or **Allow Incoming by IP address**. *Allow incoming* allows all remote clients access to the selected application, and *Allow Incoming by IP address* allows only specific remote clients' IP addresses to the selected application. You can select to remove an application completely from the firewall by selecting the application from the bottom panel then clicking **Remove**.

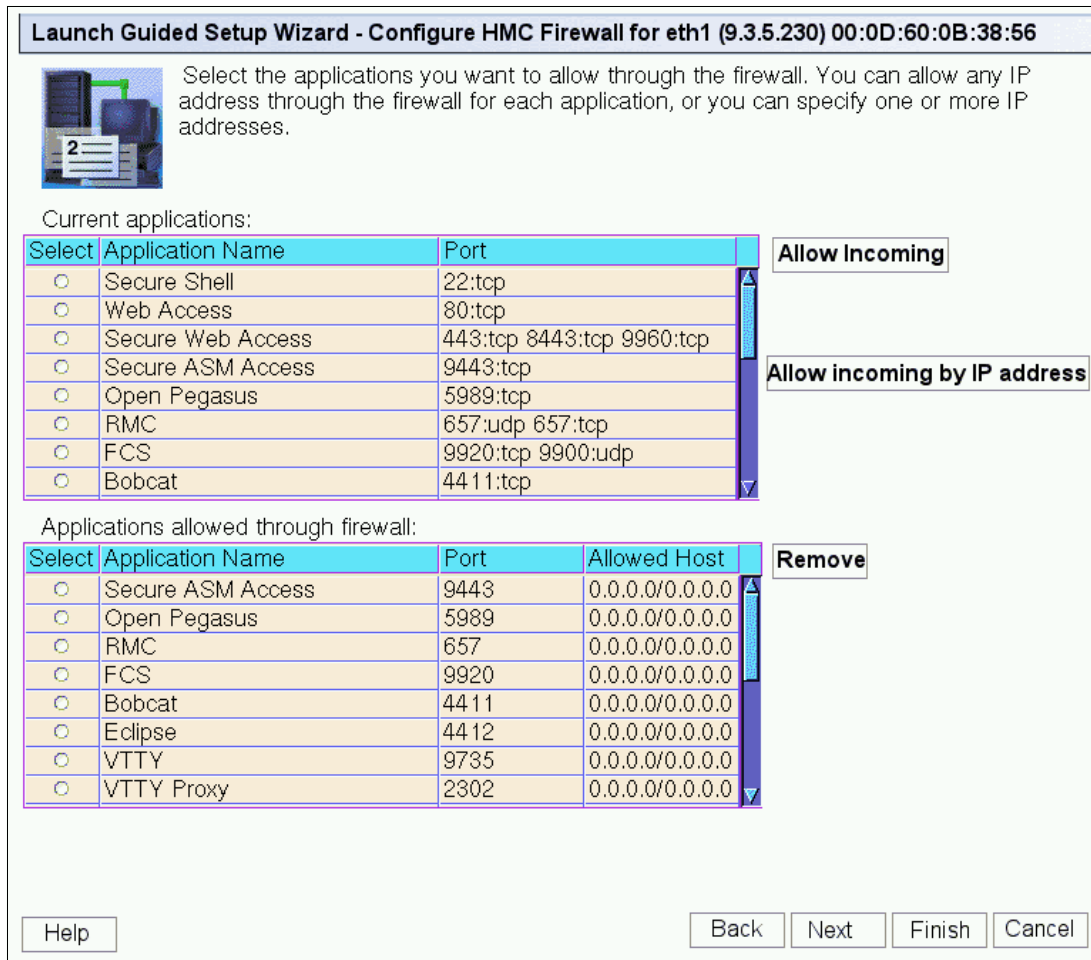


Figure 3-13 Launch Guided Setup Wizard - Configure HMC Firewall for eth1

Click **Next** to continue with the Guided Setup wizard.

18. The Launch Guided Setup Wizard - Configure HMC Network Settings panel is shown (see Figure 3-14). If you have more network adapters available, you can configure them now by selecting the adapters and **Yes**.

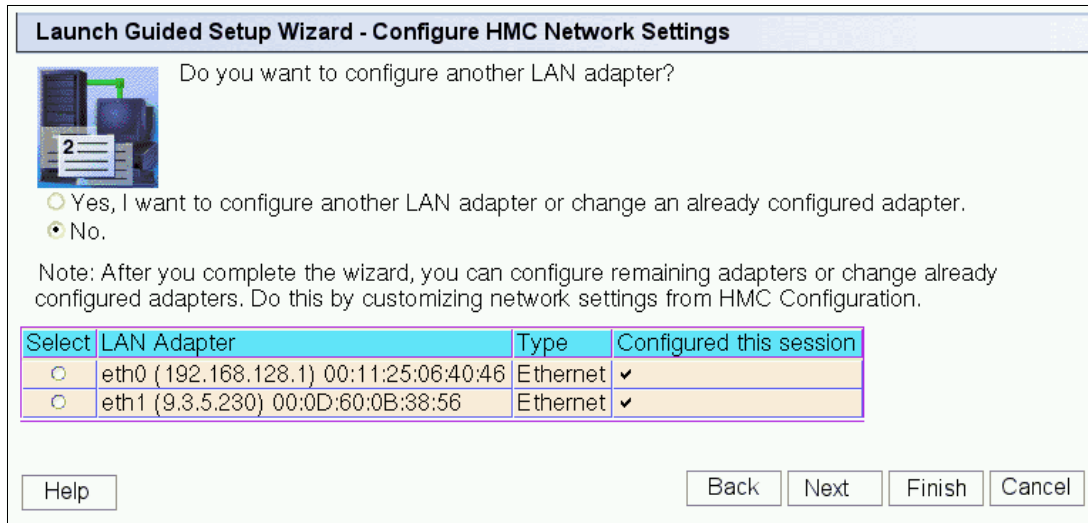
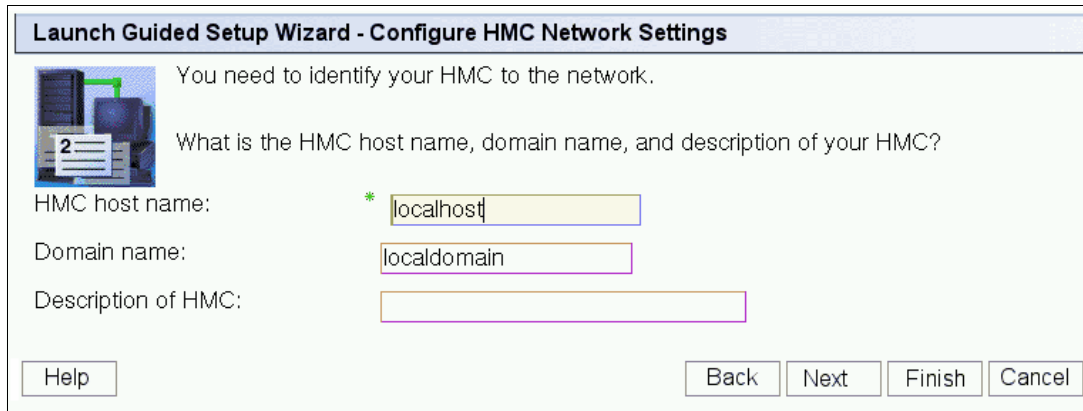


Figure 3-14 Launch Guided Setup Wizard - Configure HMC Network Settings

Because both network adapters are configured, **No** is selected. Then click **Next** to continue with the Guided Setup wizard.

19. The Launch Guided Setup Wizard - Configure HMC host name and domain panel is displayed, as shown in Figure 3-15. Enter your host name for the HMC, domain name, and description for the HMC. In our example, we enter host name *localhost* and domain name *localdomain*.



Launch Guided Setup Wizard - Configure HMC Network Settings

You need to identify your HMC to the network.

What is the HMC host name, domain name, and description of your HMC?

HMC host name: * localhost

Domain name: localdomain

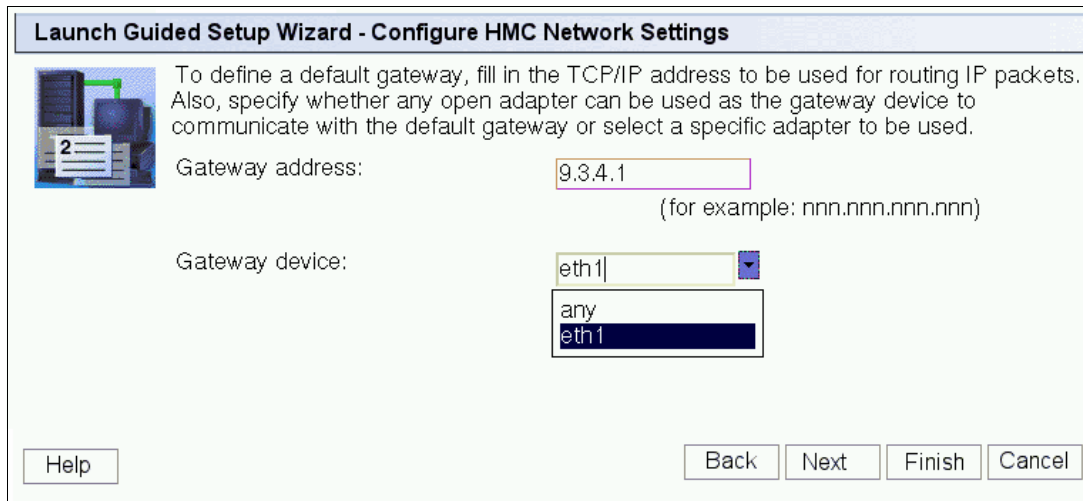
Description of HMC:

Help Back Next Finish Cancel

Figure 3-15 Launch Guided Setup Wizard - Configure HMC Host Name

Click **Next** to continue with the Guided Setup wizard.

20. The Launch Guided Setup Wizard. The configure the HMC gateway under the network settings dialog panel is shown (see Figure 3-16). If required, we can specify one of our LAN adapters as a gateway device to an open network.



Launch Guided Setup Wizard - Configure HMC Network Settings

To define a default gateway, fill in the TCP/IP address to be used for routing IP packets. Also, specify whether any open adapter can be used as the gateway device to communicate with the default gateway or select a specific adapter to be used.

Gateway address: 9.3.4.1
(for example: nnn.nnn.nnn.nnn)

Gateway device: eth1

any
eth1

Help Back Next Finish Cancel

Figure 3-16 Launch Guided Setup Wizard: Configure HMC gateway

In our example, eth1 is the gateway device to an open network. Click **Next** to continue with the Guided Setup wizard.

21. The Launch Guided Setup Wizard - Configure DNS panel displays now (Figure 3-17).

A DNS server is a distributed database for managing host names and their IP addresses. By adding a DNS server, the HMC allow us to find other hosts in our open network by their host name rather than by their IP addresses.

Enter the IP address of your DNS server or servers in the DNS server address field and click **Add** to register the IP address. You can enter multiple DNS server addresses here. The order that the addresses are entered is the order in which they are searched when trying to resolve a host name.

If you make a mistake when entering an address, you can remove it by selecting the entry and then clicking **Remove**.

Launch Guided Setup Wizard - Configure DNS

The Domain Name System (DNS) is used to provide a standard naming convention for locating IP-based computers. By defining DNS servers, you can use host names to identify servers and HMCs rather than IP addresses.

Determine if you want to enable DNS. If DNS is enabled, specify the DNS server addresses. Addresses are searched in the order they are listed.

Do you want to use DNS?

Yes, I want to use DNS by defining DNS servers.

DNS server search order:

DNS server address:	<input type="text" value="9.3.4.2"/>	<input type="button" value="Add"/>	<input type="button" value="Remove"/>
---------------------	--------------------------------------	------------------------------------	---------------------------------------

No, I do not want to use DNS.

Figure 3-17 Launch Guided Setup Wizard - Configure DNS

Click **Next** to continue with the Guided Setup wizard.

22. Launch Guided Setup Wizard. The configure the domain suffix using the Specify Domain Suffixes dialog panel is shown in Figure 3-18.

Enter a domain suffix in the Domain suffix field and click **Add** to register your entry. You can enter multiple domain suffixes for your organization if you have them. The order that the addresses are entered is the order in which they are searched when trying to map the host name to a fully qualified host name.

If you make a mistake when entering an address, you can remove it by selecting the entry and then clicking **Remove**.

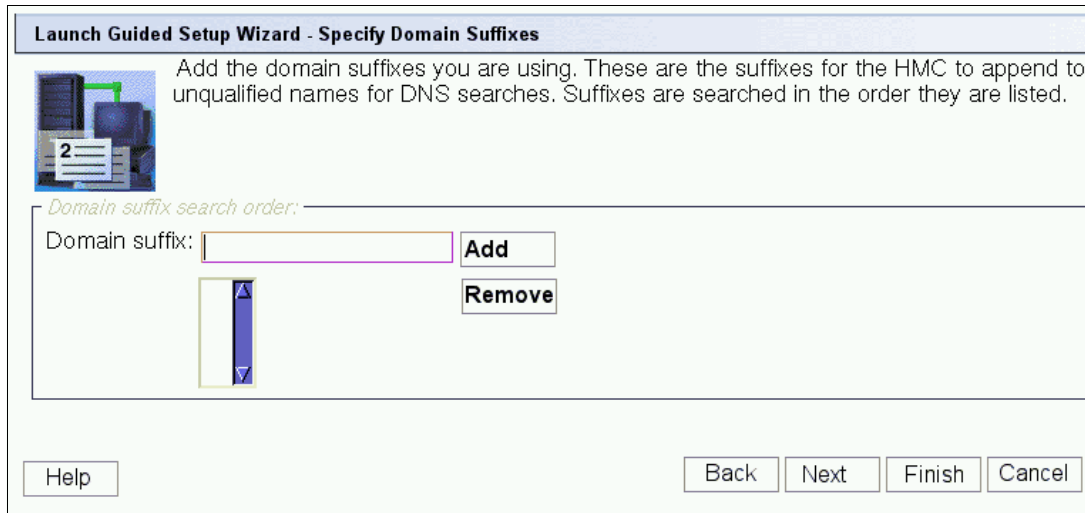


Figure 3-18 Launch Guided Setup Wizard - Configure Domain Suffix

Click **Next** to continue with the Guided Setup wizard.

23. The Launch Guided Setup Wizard - The Next Steps panel is displayed (see Figure 3-19). This step completes the network configuration section of the Guided Setup wizard. Now continue with the next part of the wizard.

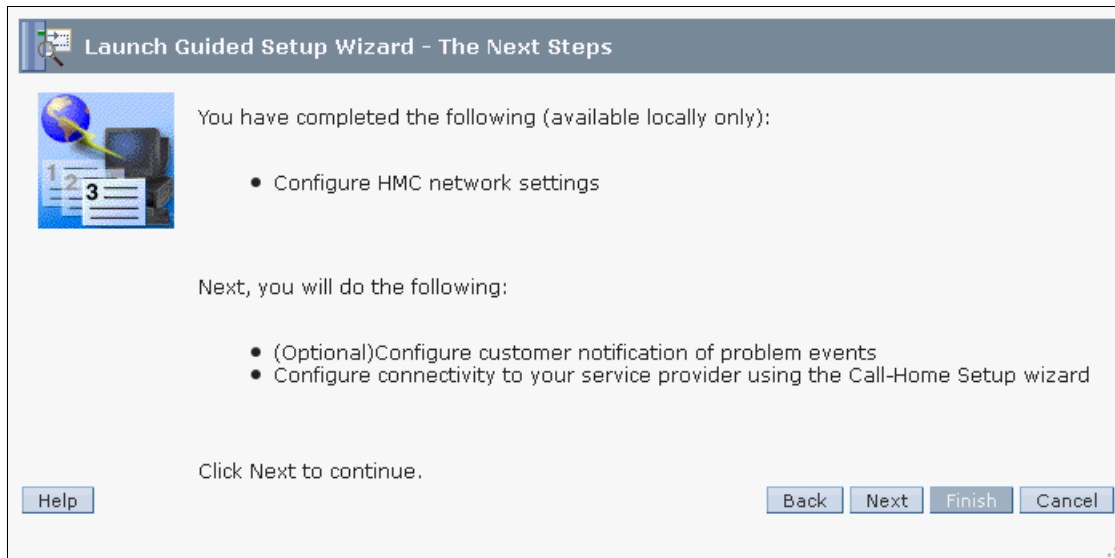


Figure 3-19 Launch Guided Setup Wizard - The Next Steps panel

Click **Next** to continue with the Guided Setup wizard.

24. Enter your SMTP server information and also add the email address for notifications, as shown Figure 3-20. Click **Next** to continue.

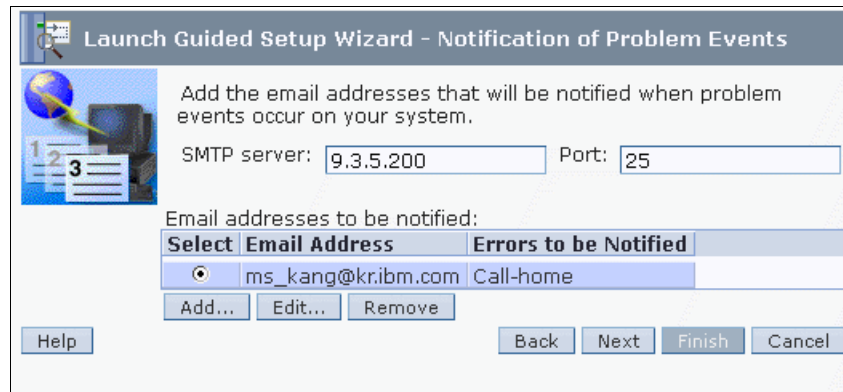


Figure 3-20 Notification of Problem Events panel

Click **Finish** to apply the changed configurations.

25. This step completes the second part of the Guide Setup wizard, as shown Figure 3-21. Click **Finish** and **Close** to continue next.

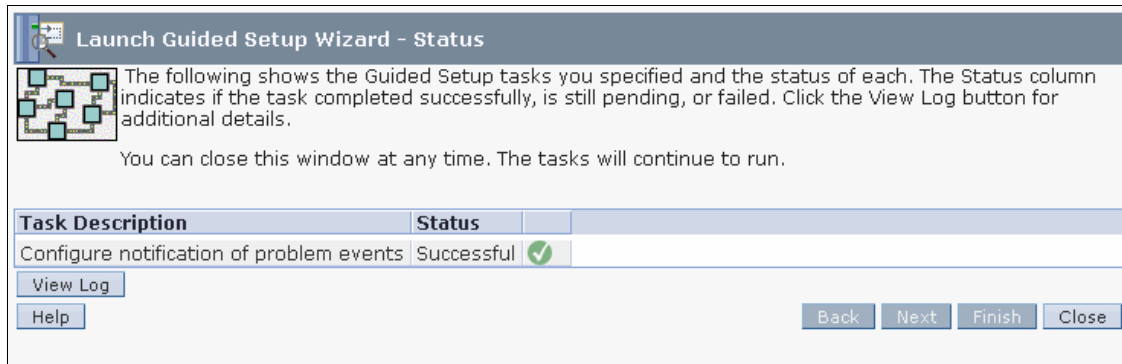


Figure 3-21 Status panel

26. The Launch Guided Setup Wizard - The Next Steps panel displays (see Figure 3-22). This step completes the network configuration section of the Guided Setup Wizard. We now continue with the Call Home Setup Wizard.

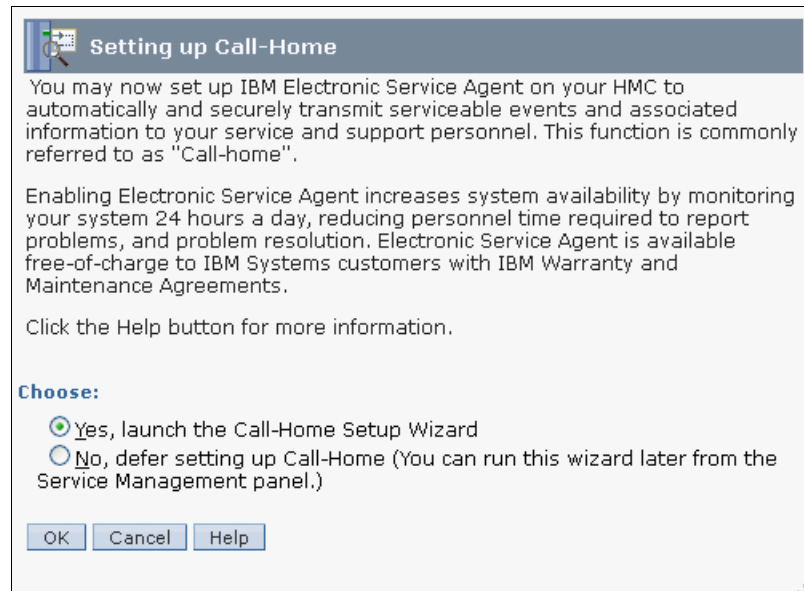


Figure 3-22 Setting up the call home window

Choose **Yes** and click **OK**.

27. You can see the Welcome page of the call home setup. Click **Next** to continue.
28. Call Home Setup Wizard - Specify Contact Information panel is shown (see Figure 3-23). This is the panel that contains the contact details for your organization. The information entered here is used by IBM when dealing with problems electronically reported (calling home), and software updates. Enter valid contact information for your own location. The fields marked with an asterisk (*) are mandatory and must be completed.

Connectivity and Call-Home Servers

✓ Welcome to Electronic Service Agent
→ Administrator
System
Account
Connectivity and Call-Home Servers
HMC Connectivity
Service Agreement
Internet
Modem Settings
Local Modem
Internet VPN
Pass-Through Systems
Call-Home Server Consoles
Authorize User
Summary of Electronic Service Agent Setup

Administrator

Contact Information

Company name: * IBM
Administrator name: * Richard Connway
Email address: * a@b.c
Phone number: * 12345
Alternate phone number:
Fax number:
Alternate fax number:

Mailing Address

Street address: * IBM
Street address 2:
City or locality: * Poughkeepsie
Country or region: * United States (of America)
State or province: * New York
Postal code: * 12601

< Back Next > Finish Cancel Help

Figure 3-23 Call Home Setup Wizard: Specify Contact Information

Click **Next** to continue.

29. Call Home Setup Wizard - Location of the HMC panel is displayed. Enter the location details of this HMC. If the location address is the same as the contact address used in the previous step, then click **Use the administrator mailing address**. Otherwise, enter the correct HMC location address details.

Click **Next** to continue.

30. Call Home Setup Wizard - Account Information. You can optionally register about it.

Click **Next** to continue.

31. Call Home Setup Wizard - Connectivity and Call Home Servers. You can see the description about call home configuration.

Click **Next** to continue.

32. Call Home Setup Wizard. The Configure Connectivity to Your Service Provider panel is now shown in Figure 3-24 on page 83.

You can configure an outbound connection between the HMC and your service provider, such as your service provider's remote support facility. You can specify how the local HMC connects to your service provider from a local modem, an Internet Secure Sockets Layer (SSL), an Internet Virtual Private Network (VPN), or a remote pass-through system.

- With an Internet SSL, you can use a high speed Internet connection on your HMC, the fastest option to send problem information to your service provider.
- With a local modem, you can use the modem on your HMC to send problem information and system data to your service provider.
- With an Internet VPN, you can use a high speed Internet connection on your HMC to send problem information to your service provider.
- With a remote pass-through system, you can use another HMC or a logical partition on your server to send problem information to your service provider.

Figure 3-24 shows the Call Home Setup Wizard: Configure Connectivity panel.

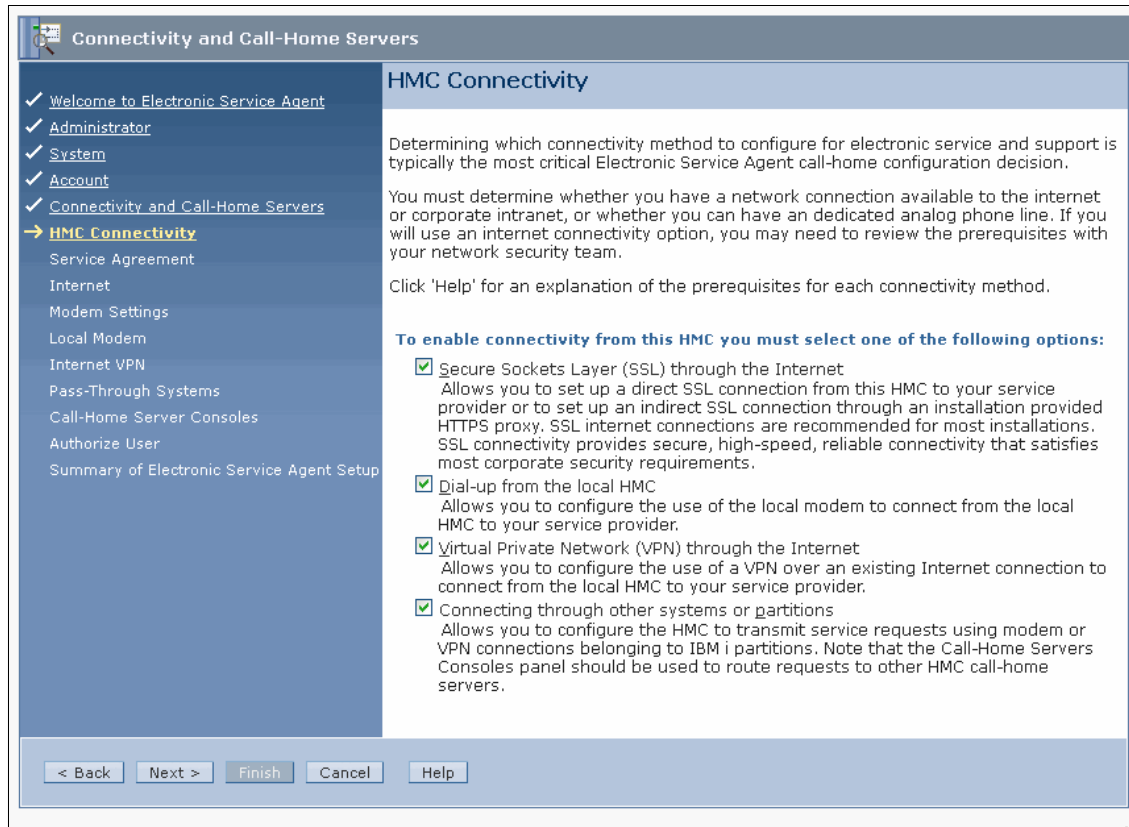


Figure 3-24 Call Home Setup Wizard: Configure Connectivity

You can select by the communications method to which you want to connect to IBM for service and support-related functions. In our example, we select all four connectivity options for demonstration purposes only. Normally, you would select only the options valid for your environment.

Click **Next** to continue with Guided Setup.

33. The Agreement for Service Programs panel is displayed. Read the agreement details carefully and click **Accept** or **Decline**.

34. Call Home Setup Wizard - Configure SSL by using an Existing Internet Connection.

Click **Next** to continue.

35. In the panel, The Modem Configuration display opens. You can set the Dial Type (Tone/Pulse), Wait for dial tone, Enable speaker, and the Dial prefix value.

Click **Next** to continue.

36. Call Home Setup Wizard: The Configure Local Modem panel is shown in Figure 3-25.

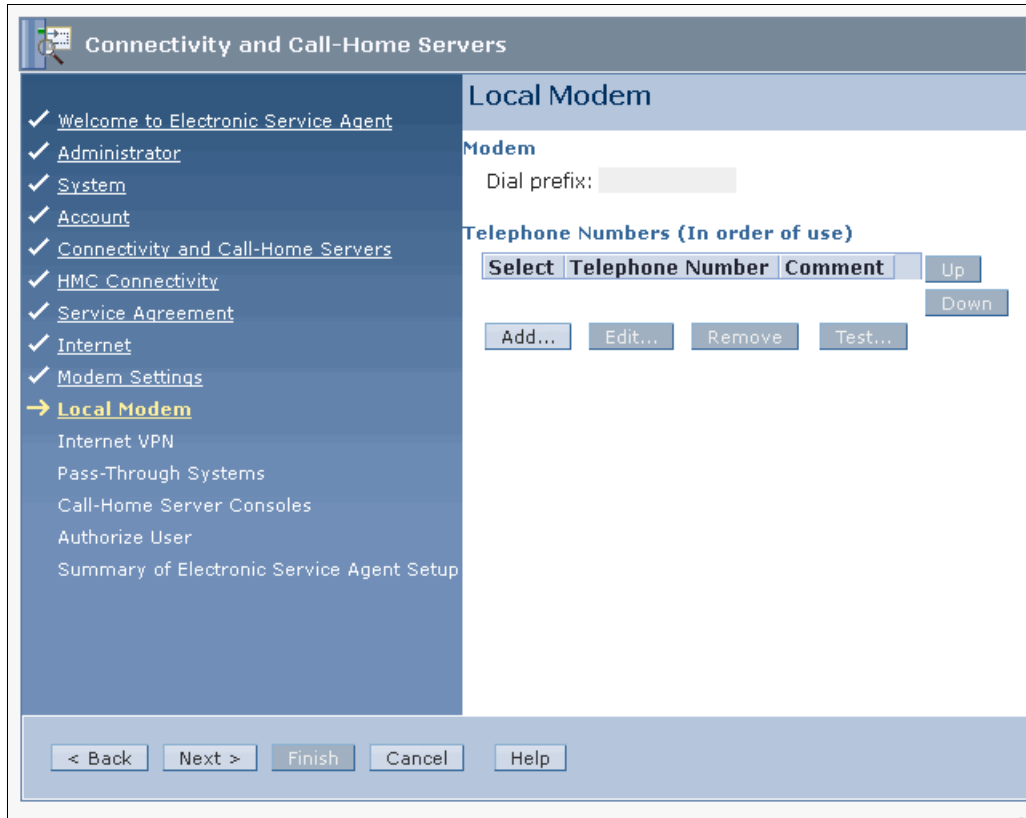


Figure 3-25 Configure local modem

Click **Add** in the Phone numbers panel to add the IBM support service phone number. The Add Phone Number - Country or Region window displays. Use the drop-down menus to select your country or region and then your state or province.

After you select your Country/region and State/province, a list of available IBM support service numbers are listed. Select the phone number nearest to your location. The phone number is populated in the Phone number field at

the bottom of the panel. You can also manually add phone numbers if you know your IBM support service number.

Click **Next** to continue with the Guided Setup wizard.

37. Call Home Setup Wizard - Use VPN by using an Existing Internet Connection panel.

Click **Next** to continue.

38. Configure Connectivity by using a Pass-Through Systems panel, which is shown in Figure 3-26. By configuring a pass-through system, send problem information to your service provider through another HMC or a logical partition.

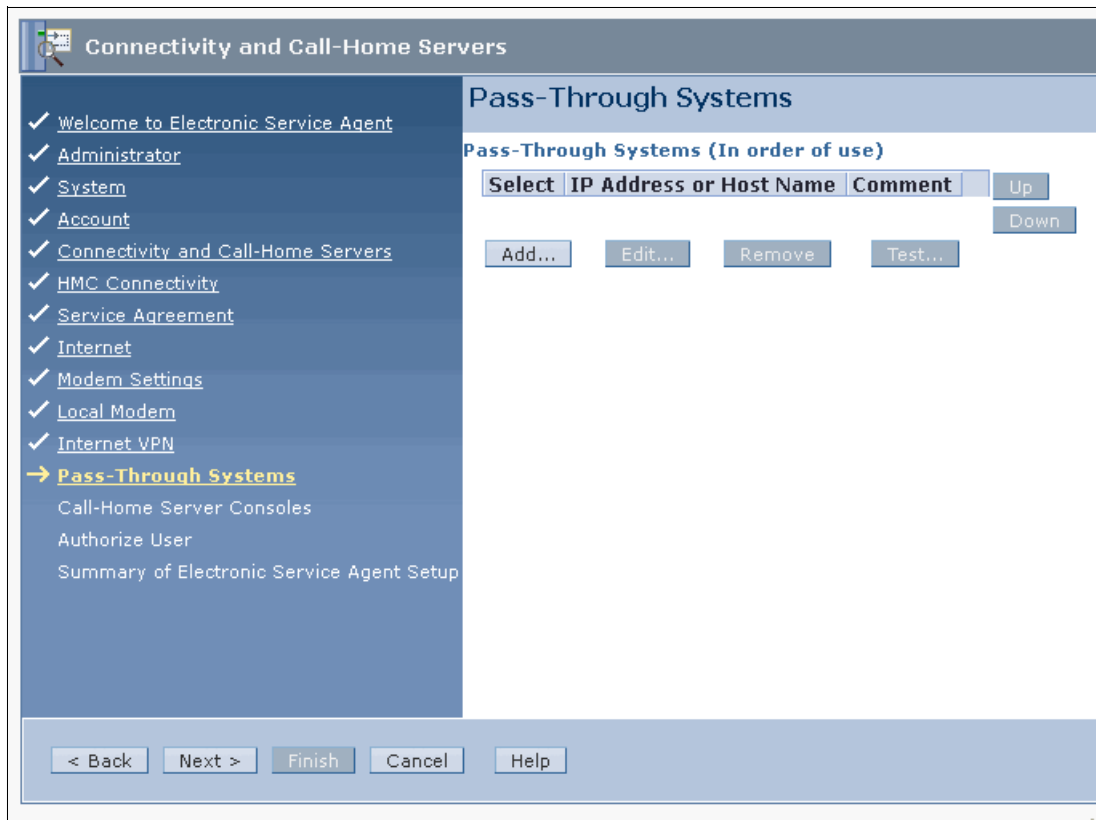


Figure 3-26 Configure a pass-through system

39. The Call Home Server Console is configured. See Figure 3-27. You can add or remove call home server consoles.

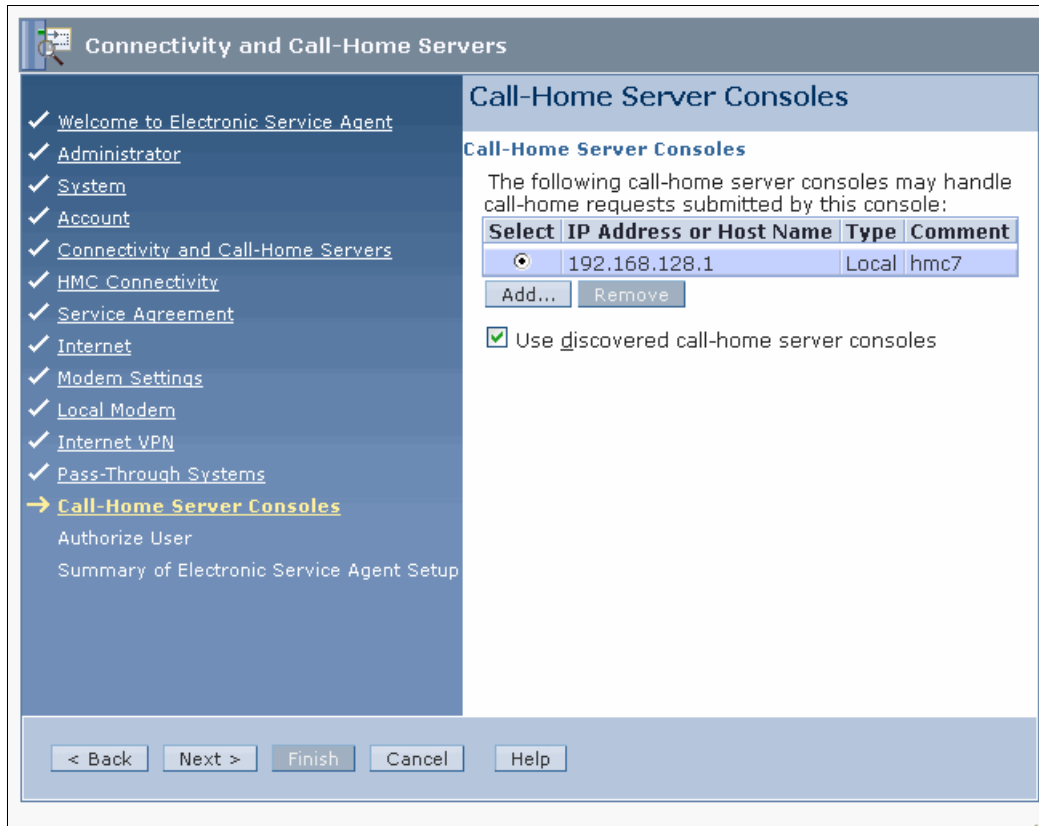


Figure 3-27 Configure Call Home Server Consoles

40. Click **Next** to continue. The authorize users for IBM Electronic Service Agent™ (ESA) panel is displayed, as shown in Figure 3-28.

By entering your IBM ID, you can access the inventory information that is collected and transmitted to IBM by ESA. If you do not have an IBM ID, you can obtain one, as explained in Figure 6-37 on page 301.

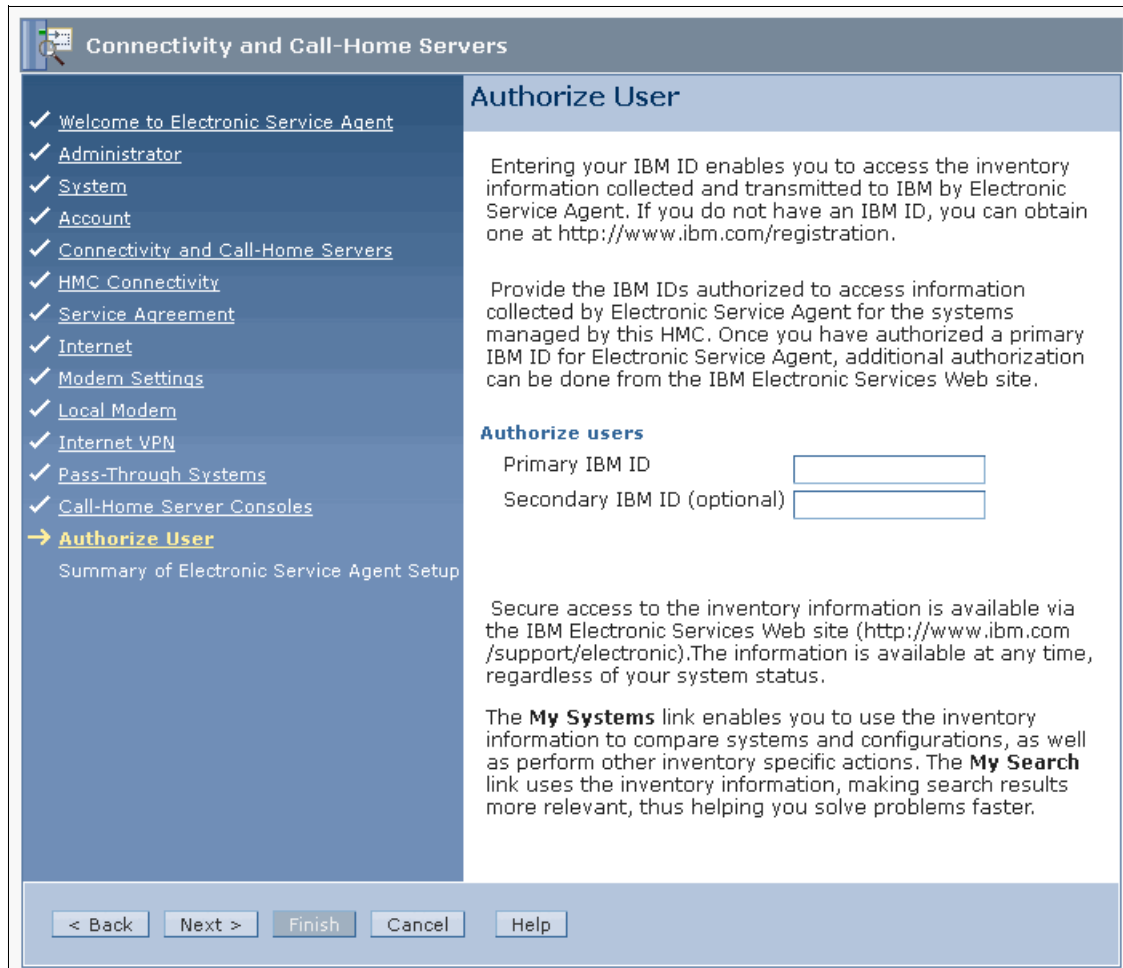


Figure 3-28 Authorize users for Electronic Service Agent

41. Click **Next** to continue with the Call Home Setup Wizard. The summary window is displayed (Figure 3-29). You can see all the changes that the Call Home Setup Wizard configures later. At this stage, nothing changed on the HMC. You can cancel the changes by clicking **Cancel**.

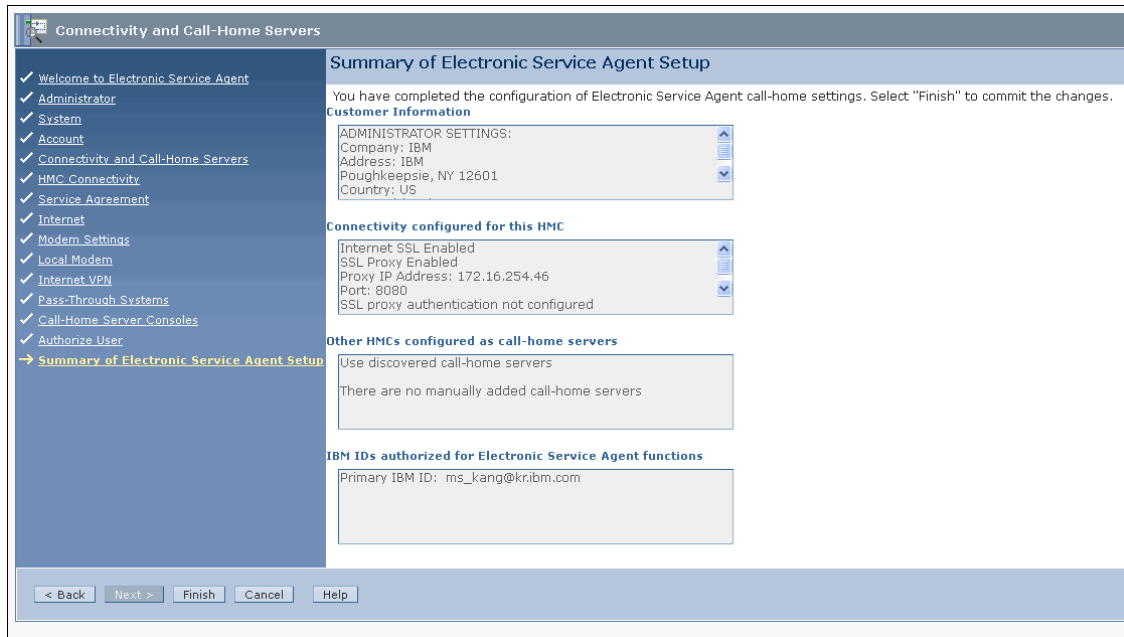


Figure 3-29 Summary of ESA

Click **Finish** to apply the changed configurations.

Changes: At this step, nothing changed on the HMC. If you click **Cancel**, all changes by the Launch Guided Setup wizard be lost.

42. Click **OK** to log off from the HMC then log on to apply some changes.

This step completes the Call Home Setup Wizard.

Post Guided Setup Tasks

If you were not able to set up all the information through the wizard, you can go back and use the standard HMC menu to complete tasks. If you are directly connected to the HMC, some tasks might be missed at first.

3.1.3 Connecting managed systems to the HMC

Attention: When installing a new Power Systems server, do not turn on the system before connecting it to an HMC. The server processor (SP) on a System p system is a DHCP client and searches for a DHCP server to obtain its IP address. If no DHCP server can be found, the SP assigns a default IP address. If this occurs, you have to use address space manager (ASM) to change the IP setting of the SP manually.

After you complete the HMC Guided Setup Wizard, you can connect your managed systems to the HMC by using following steps:

1. Connect your HMC to the HMC port of the managed system with an Ethernet cable.
2. Connect the managed system to a power source. The managed system then turns on its service processor. After the service processor is turned on, proceed to the next step. This process takes three to five minutes. You can see the following sequence of events signal that power is applied to the service processor:
 - a. Progress indicators, also referred as *checkpoints*, show on the control panel display while the system is being started. The display might appear blank for a few moments during this sequence.
 - b. When the service processor completes its power-on sequence, the green power light flashes slowly and the output on the control panel is similar to the following output:

```
01 N V=F
T
```
3. Click **Systems Management**, then **Servers** to view the status of your managed system. It can take a few minutes for the status to display.

- If the status shows *Pending Authentication*, you must set passwords for the managed system. The HMC prompts you to set passwords for the managed system. If you are not prompted by the HMC to set those passwords, click **Operations** then **Change Password**. The window for setting passwords opens, as shown in Figure 3-30. Set the password for each, as directed.

Update Password - Authentication Passed

Authentication passed for the managed system below.

Managed system name : 9117-MMA-SN10DD4AC-L10

You may change the HMC Access password at this time using the fields below. Once changed, you must update the HMC Access password for all the other Hardware Management Consoles from which you want to access this managed system.

Current HMC Access password:

New HMC Access password:

Verify HMC Access password:

Click OK to change the password or Cancel to quit the process.

Figure 3-30 Update Password - Authentication Passed

HMC configuration: If you did not configure your HMC as a Dynamic Host Configuration Protocol (DHCP) server, the HMC does not detect the managed system automatically.

- Access the ASMI to set the time of day on the system. Refer to the “Time of the Date” section to set the time of day on the system.

6. Start the managed system by clicking **Systems Management** → **Servers**, select the managed system that you want to turn on, then click **Operation**, **Power On**, as shown in Figure 3-31. There are three options for turning on the system:
 - a. Create and activate logical partitions with partition standby. When the partition standby power-on is completed, the system is in standby mode.
 - b. System profile turns on the system according to a predefined set of system profiles. Select the system profile that you want to use from the list.
 - c. Partition auto start turns on the managed system to partition standby and then activate all partitions that are marked as auto start or those partitions that were running when the system shut down.

Power On - 9117-MMA-SN10DD4AC-L10

To power on the managed system, select one of the power on options below and click OK. You must specify a system profile if you choose the System Profile power on option

Power On Options

Partition standby
 System profile
 Partition auto start

System Profiles

Select a system profile below that contains the partitions that you want to have activated when the managed system is powered on.

OK Cancel Help

Figure 3-31 Power on the managed system

- d. This step completes the instructions for installing the HMC.

3.2 HMC connectivity scenario

POWER6, POWER6+, POWER7, and POWER7+ processor technology-based servers that are managed by an HMC require Ethernet connectivity between the HMC and the Service Processor of the server. In addition, if Dynamic LPAR, Live Partition Mobility, or PowerVM Active Memory Sharing operations are required on the managed partitions, Ethernet connectivity is needed between these partitions and the HMC. A minimum of two Ethernet ports are needed on the HMC to provide such connectivity.

This section covers a variety of HMC connectivity.

Connect from single HMC to Multi-drawer

For the HMC to communicate properly with the multi-drawer server, eth0 of the HMC should connect to the port labeled HMC Port1 on the two central electronics complex drawers of each system (Figure 3-32).




Figure 3-32 Connectivity: Multi-drawer with one HMC

Although other network configurations are possible, you can attach an eth1 to each HMC Port 2 of the servers for network redundancy (Figure 3-33 on page 93). Address these by two separate subnets.




Figure 3-33 Network Redundancy

Single HMC to support flexible service processor failover

Dual service processor support is available only on the 5Core or larger P5-570 systems and requires an HMC to manage this system. The HMC and server firmware should be on minimum HMC code level, and correctly network the HMC to support dynamic flexible service processor (FSP) failover (see Figure 3-34). Note, on POWER 770, no loopback cable is required to perform FSP failover.

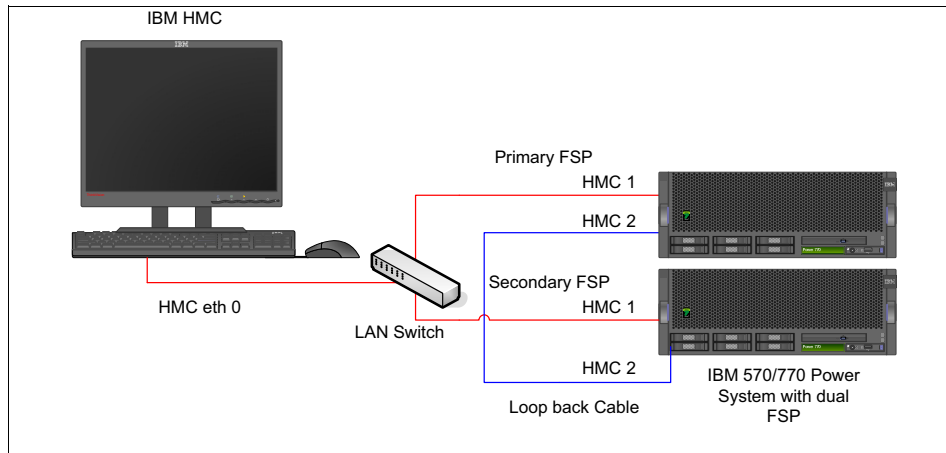


Figure 3-34 To support FSP failover

Dual HMC redundancy

A redundant HMC manages a system that is already managed by another HMC. When two HMCs manage one system, they are peers, and each can be used to control the managed system. If both HMCs are connected to the server by using private networks, each HMC should be a DHCP server setup to provide IP addresses on two unique, non-routing IP ranges. For the best redundancy, two HMCs are kept on separate subnetworks and attach to different server support network ports.

For more information, see 2.4.1, “Dual HMC and Redundancy” on page 43.

Redundant HMC for High-End POWER Server

The POWER High-End server(595,795) requires one primary HMC that can communicate to all bulk power hubs (BPHs) in the system. The primary HMC is connected to port J01 on the BPH on the front side of the system (central electronics complex) rack. For improved system availability, a redundant HMC is highly preferable. It is connected to port J01 on the BPH on the back side of the system rack.

It is common to use an Ethernet hub or switch to establish the connections between the HMC and the High-End server (see Figure 3-35).




Figure 3-35 Redundant HMC for High-End Server



Configuring

An overview of Hardware Management Console (HMC) network configuration, user management, certificate management, virtualization, and capacity on demand is provided.

4.1 Network configuration

This chapter provides a general overview of the types of network configurations for the HMC and explains how to configure HMC network settings. Also described is how to use the HMC workplace to obtain network diagnostic information.

4.1.1 Types of HMC network configurations

The HMC supports several network communications:

- ▶ *HMC to managed system connection* performs most of the hardware management functions in which HMC issues control function requests through service processor of the managed system.
- ▶ *HMC to logical partition connection* collects platform-related information, such as hardware error events or hardware inventory, from the operating system running in the logical partitions. This communication also coordinates certain platform activities, such as Dynamic Logical Partition (DLPAR) or concurrent maintenance with those operating systems.
- ▶ *HMC to remote users connection* provides remote users with access to HMC functionally. Remote users can access the HMC by using one of the following methods:
 - The remote operation to access all the HMC graphical user interface (GUI) functions remotely.
 - SSH to access the HMC command-line functions remotely.
 - A virtual terminal server for remote access to virtual logical partition consoles.
- ▶ *HMC to service and support connection* transmits data such as hardware error reports, inventory data, and microcode updates, to and from your service provider. You can use this communication path to make automatic service calls.

4.1.2 Configuring the HMC network setting

This section describes network configuration for the HMC. To open the Change Network Setting window, select **Change Network Settings** from the main menu.

HMC identification

HMC identification provides information that is needed to identify the HMC in the network. The **Identification** tab of the **Customize Network Settings** window (Figure 4-1) includes the following information:

▶ **Console name**

HMC name that identifies the console to other consoles in the network. This name is the short host name.

▶ **Domain name**

An alphabetic name that the domain name server (DNS) can translate to the Internet Protocol (IP) address.

▶ **Console Description**

Short description for the HMC.

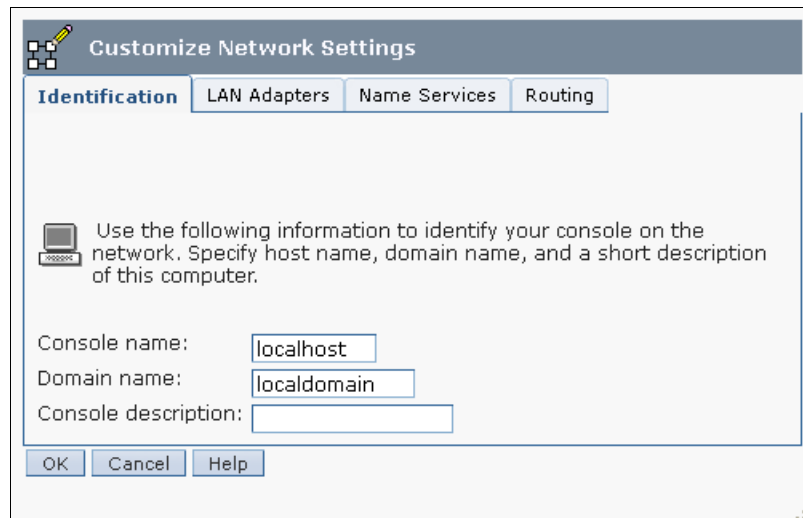


Figure 4-1 HMC Identification tab

LAN Adapters

The LAN Adapters tab (Figure 4-2) shows a summarized list of all local area network (LAN) adapters that are installed on the HMC. You can view details of each LAN adapter by clicking **Details**, which starts a window where you can change LAN adapter configuration and firewall settings.

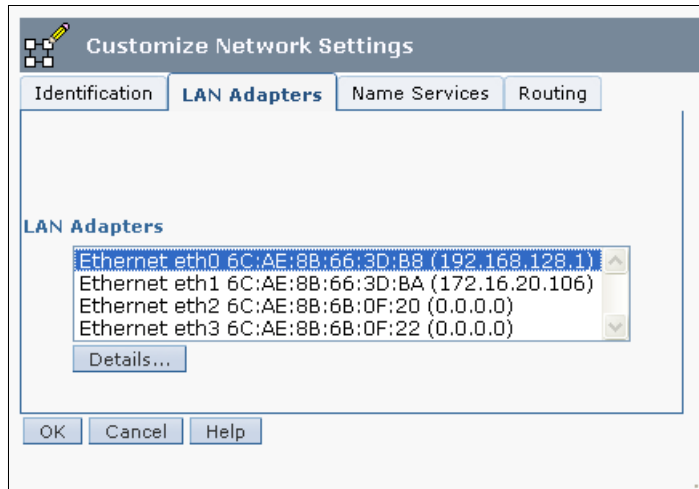


Figure 4-2 LAN Adapters tab

LAN Adapter configuration

The LAN Adapter Details window, which is shown in Figure 4-3, describes the LAN adapter configuration of Ethernet *eth0* on the LAN Adapter tab.

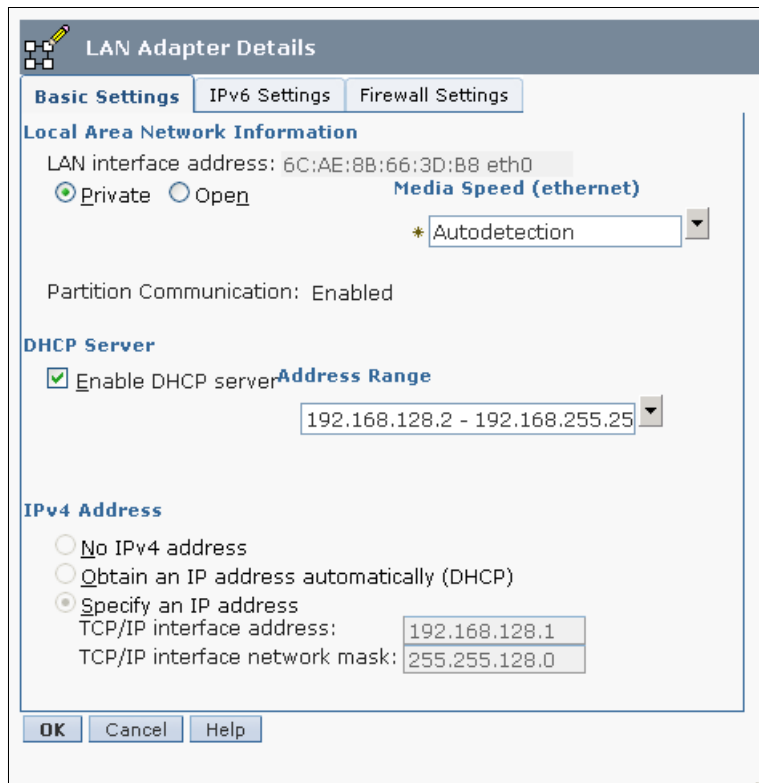


Figure 4-3 LAN Adapter configuration

The LAN Adapter tab of this window includes the following information:

- ▶ Local area network information

The LAN interface address shows Media Access Control (MAC) Address on the card and the adapter name. These values uniquely identify the LAN adapter and cannot be changed. A private network is used by the HMC to communicate by its managed system and an open network connects the HMC outside the managed system. Media speed specifies the speed in duplex mode of an Ethernet adapter. The options are Autodetection, 10 Mbps Half Duplex, 10 Mbps Full Duplex, 100 Mbps Half Duplex, 100 Mbps Full Duplex, or 1000 Mbps Full Duplex.

The connection between the HMC and its managed systems can be implemented either as a private or open network. The term *open* refers to any

general, public network that contains elements other than HMCs and service processors that are not isolated behind an HMC. The other network connections on the HMC are considered open, which means that they are configured in a way that you would expect when attaching any standard network device to an open network.

In a private service network, however, the only elements on the physical network are the HMC and the service processors of the managed systems. In addition, the HMC provides Dynamic Host Configuration Protocol (DHCP) services on that network, which allow it to automatically discover and assign IP configuration parameters to those service processors. You can configure the HMC to select one of several different address ranges to use for this DHCP service so that the addresses provided to the service processors do not conflict with addresses used on the other networks to which the HMC is connected. The DHCP services allow the elements on the private service network to be automatically configured and detected by the HMC, while at the same time preventing address conflicts in the network.

On a private network, therefore, all of the elements are controlled and managed by the HMC. The HMC also acts as a functional firewall, isolating that private network from any of the open networks to which the HMC is also attached. The HMC does not allow any IP forwarding. Clients on one network interface of the HMC cannot directly access elements on any other network interface.

To take advantage of the additional security and ease of setup, implement service network communications through a private network. However, in some environments, this configuration is not feasible because of physical wiring, floor planning, or control center considerations. In this case, the service network communications can be implemented through an open network. The same functionality is available on both types of networks, although the initial setup and configuration on an open network require more manual steps.

► DHCP Server

Choose **Enable DHCP Server** only if this adapter is defined as private network, then choose one range of addresses for the DHCP Server to distribute. If the adapter is defined as open, this setting is not available.

If you want to configure the first network interface as a private network, you can select from a range of IP addresses for the DHCP server to assign to its clients. The selectable address ranges include segments from the standard nonroutable IP address ranges.

In addition to these standard ranges, a special range of IP addresses is reserved for IP addresses. This special range can be used to avoid conflicts in cases where the HMC-attached open networks are using one of the nonroutable address ranges. Based on the range that is selected, the HMC

network interface on the private network is automatically assigned the first IP address of that range, and the service processors are then assigned addresses from the rest of the range.

The DHCP server in the HMC uses automatic allocation, which means that each unique service processor Ethernet interface is reassigned the same IP address each time it is started. Each Ethernet interface has a unique identifier that is based upon a built-in MAC address, which allows the DHCP server to reassign the same IP parameters.

► DHCP Client/IP address

There are two options:

- *Obtain an IP address automatically* allows the HMC to obtain an available IP address automatically.
- *Specify an IP address* specifies an IP address to be used, providing a TCP/IP interface address and TCP/IP interface network mask.

HMC support: HMC supports IPv6 protocol. For the instructions, link to the Systems Hardware information website:

<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7hail/configlan.htm>

Firewall settings

You use the Firewall Settings tab of the LAN Adapter Details window to view and change current firewall adapter settings for the specified LAN interface address. Select **Allow Incoming** to allow access to incoming network traffic from this address, or select **Allow Incoming by IP Address** to allow access by incoming network traffic from hosts that are specified by an IP address and network mask, as shown in Figure 4-4.

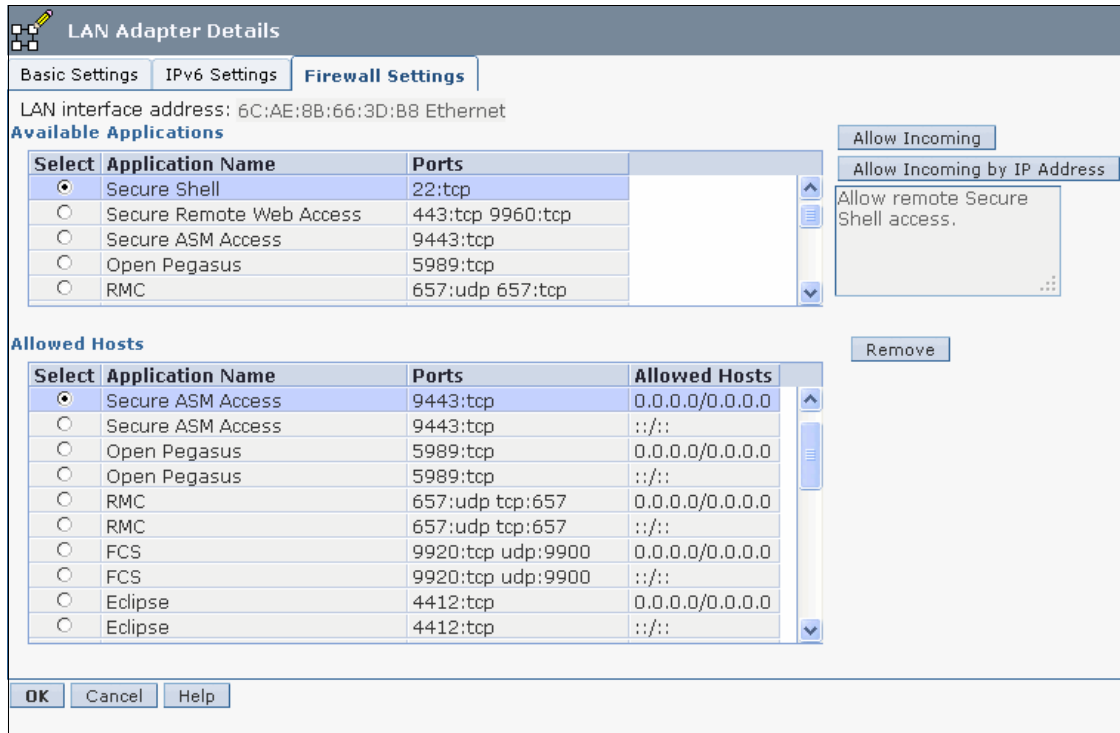


Figure 4-4 Firewall Settings tab

Name Services

You use the Name Services tab to specify DNS for configuring the console network settings (Figure 4-5). DNS is a distributed database system for managing host names and their associated IP addresses. With DNS, people can use names to locate a host, rather than using the IP address.

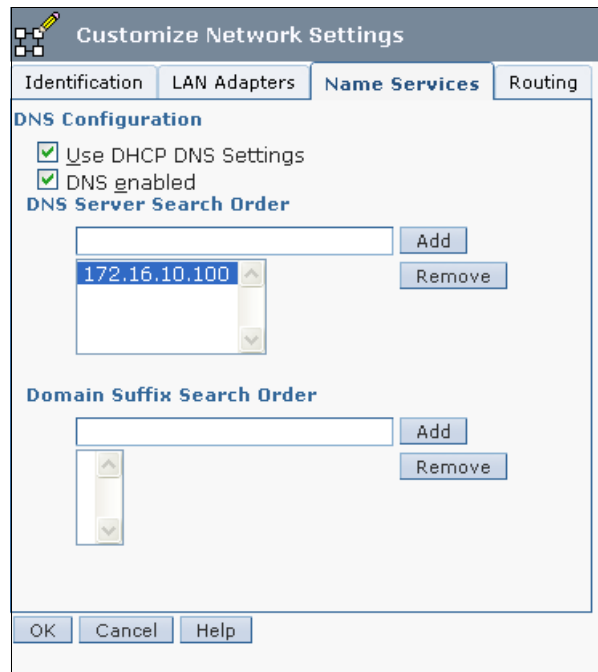


Figure 4-5 Name Services tab

Routing

In the Routing tab, you specify routing information for configuring the console network settings, such as add, delete, or change routing entries and specify routing options for the HMC, as shown in Figure 4-6.

Kernel IP routing table						
Destination	Gateway	Genmask	Flags	MSS Window	irtt	Iface
172.16.20.0	0.0.0.0	255.255.252.0	U	0 0	0	eth1
192.168.128.0	0.0.0.0	255.255.128.0	U	0 0	0	eth0
10.253.0.0	0.0.0.0	255.255.0.0	U	0 0	0	s10
169.254.0.0	0.0.0.0	255.255.0.0	U	0 0	0	eth0
127.0.0.0	0.0.0.0	255.0.0.0	U	0 0	0	lo
0.0.0.0	172.16.20.1	0.0.0.0	UG	0 0	0	eth1

Kernel IPv6 routing table						
Destination	Next Hop	Flags	Metric	Ref	Use	Iface
fe80::/64	::	U	256	0	0	eth0
fe80::/64	::	U	256	0	0	eth1
::1/128	::	U	0	466	1	lo
fe80::6eae:8bff:fe66:3db8/128	::	U	0	0	1	lo
fe80::6eae:8bff:fe66:3dba/128	::	U	0	0	1	lo
ff02::1/128	ff02::1	UC	0	1	0	eth0
ff00::/8	::	U	256	4	0	eth0
ff00::/8	::	U	256	0	0	eth1

Figure 4-6 Routing tab

Routing Information

The routing information displays the following components:

- ▶ *Type* displays the specific route, which can be one of three choices:
 - Net

Specifies a network-specific route. The destination address is the TCP/IP address of a particular network. All TCP/IP communications that are destined for that network are using the TCP/IP address of the router, unless a host route also applies for the communication to the destination host address.

Conflicts: When a conflict occurs between a host and net route, the host route is used.

- Host

Specifies a host-specific destination. The destination address is the TCP/IP address of a particular host. All TCP/IP communications that are destined for that host are routed through the router by using the router address as the TCP/IP address.

- Default

Specifies all destinations that are not defined with another routing table entry. With a default route, the destination address is all zeros. If no host or net routes apply when communicating with a destination host address, the communication is routed through the default router by using the TCP/IP address given by the router address.

- ▶ *Destination* displays the TCP/IP address of the destination host, network, or subnet.
- ▶ *Gateway* displays the TCP/IP address of the next hop in the path to the destination.
- ▶ *Subnet Mask* displays the subnet mask that is used by network interface to add routes.
- ▶ *Interface* displays the name of the network interface that is associated with the table entry.

Default gateway information

The default gateway information provides the following components:

- ▶ *Gateway address*

The default gateway is the route to all networks. The gateway informs each personal computer or other network device where to send data if the target station is not on the same subnet as the source.

- ▶ *Gateway device*

Network interface that is used as a gateway device.

The Enable “routed” option

You use the *Enable “routed”* option to enable or disable the network routing daemon, which is *routed*. If disabled, this option stops the daemon from running and prevents any routing information from being exported from this HMC.

4.1.3 Testing network connectivity

You can use the HMC workplace to obtain network diagnostic information about the HMC's network protocols. You can use the Test Network Connectivity window to access any of the following functions:

- ▶ Ping
- ▶ Interfaces
- ▶ Ethernet Settings
- ▶ Address
- ▶ Routes
- ▶ Address Resolution Protocol (ARP)
- ▶ Sockets
- ▶ Transmission Control Protocol (TCP)
- ▶ IP tables
- ▶ User Datagram Protocol (UDP)

This section explains each of these functions. To open the Test Network Connectivity window, select **Test Network Connectivity** on the main window.

Ping

Use the Ping function to send an echo request (ping) to a remote host to see whether the host is accessible and to receive information about that TCP/IP address or name. Specify any TCP/IP address or name in the TCP/IP Address or Name to Ping field, then click **Ping**, as shown in Figure 4-7.

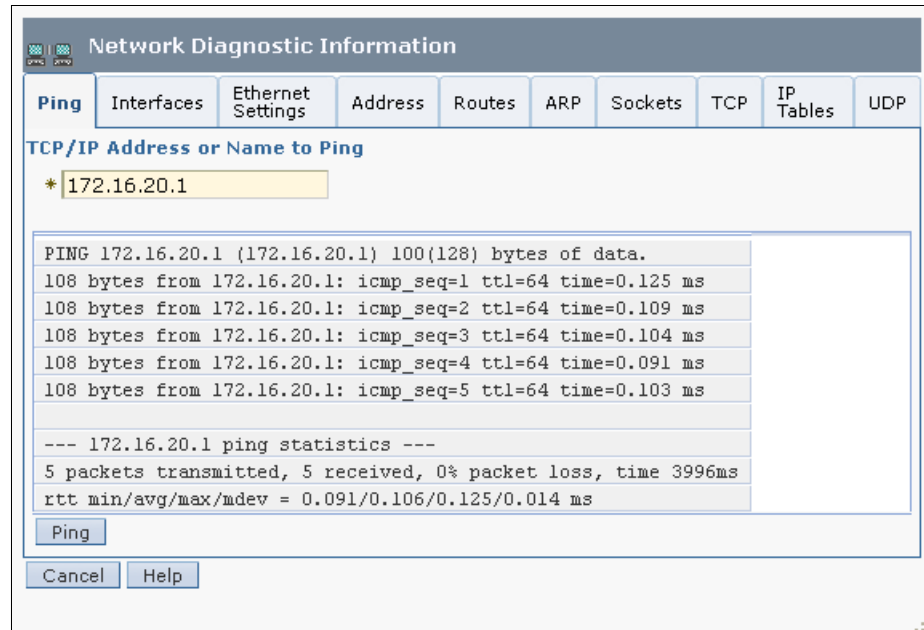


Figure 4-7 Network Diagnostic Information: Ping

In a similar way, you can run tests for *Interfaces*, *Ethernet Settings*, *Addresses*, *Routes*, *ARP*, *Sockets*, *TCP/IP Tables*, and *UDP*.

4.1.4 Viewing Network Topology

Use the Network Topology window to see a tree view of the network nodes that are known to this HMC (Figure 4-8). Examples of such nodes are managed systems, logical partitions, storage, and other HMCs. You can view attributes of a node by selecting the node in the tree view that is shown in the left pane under Current Topology. Attributes vary according to the type of node. Some examples are IP address, host name, location code, and status. Click **Refresh** to rediscover the topology and to query the nodes again for status and other attributes.

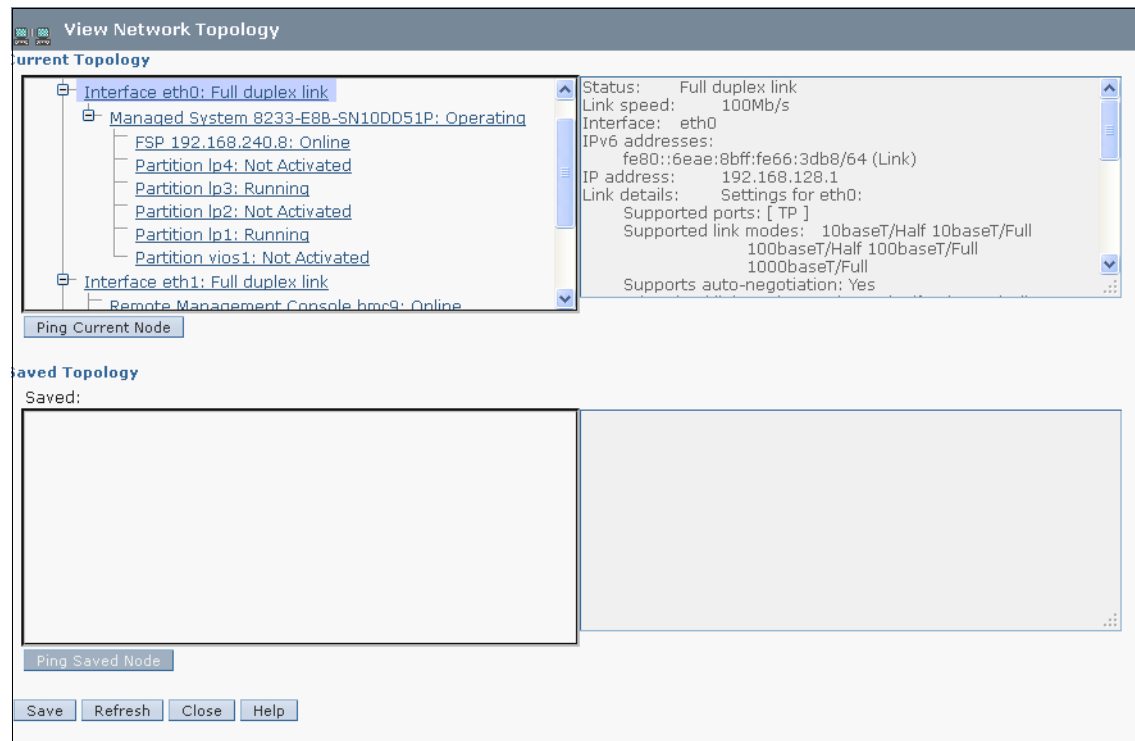


Figure 4-8 Network topology

Table 4-1 shows the possible status for each node.

Table 4-1 Possible status for each node

Node	Possible Status
Local HMC	All nodes OK; Some nodes failed; All nodes failed
Remote HMC	Online, Offline
Interface	No link; Half duplex link; Full duplex link
Storage Facility	Status not reported
Managed system	Managed system status reported by the <code>lssyscfg</code> command (Operating, Running)
FSP	Online, Offline
LPAR	LPAR status reported by the <code>lsstscfg</code> command. LPARs can also carry a connection status to report their current network status as Active, On, Off, Offline
BPA	BPA status reported by the <code>lssyscfg</code> command
BPC	Online, Offline

Each status has its meaning that is evaluated when determining cumulative status for the Local HMC node, as shown in Table 4-2.

Table 4-2 Meaning of node status

Status	Evaluation for cumulative status (OK/Fail)	Meaning
All nodes OK	OK	Child node statuses are OK
Some nodes failed	Fail	One or more child node statuses failed
All nodes failed	Fail	All child nodes statuses failed
No link	Fail	No link detected on interface
Half duplex link	OK	Half duplex link detected on interface
Full duplex link	OK	Full duplex link detected on interface
Active	OK	LPAR is pingable and known to RMC
On	Fail	LPAR is pingable but not known to RMC
Off	Fail	LPAR is not pingable nor known to RMC
Offline	Fail	For LPARs: LPARs is not pingable but is known to RMC. For Remote HMCs: Remote HMC is not pingable but is known to this HMC. For FSPs, BCPs: FSP or BPC are not pingable.
Online	OK	Remote HMC is pingable FSP is pingable BPC is pingable
Unknown	Fail	Status window to be determined
Operating, Running, or any other text from <code>lssyscfg</code>	N/A	Not evaluated when determining the cumulative status

This task also allows you to save a snapshot of the current topology and to view that saved reference topology. You can view attributes of a node in this saved topology by selecting the node in the tree view that is shown in the left pane under Saved Topology.

To test network connectivity to a node, you can select the node in either the current or the saved topology and click Ping Current Node or Ping Saved Node, available only for nodes that include an IP address or a host name.

4.2 User Management

Each HMC user can be a member of a different role. Each of these roles allows the user to access different parts of the HMC and to perform different tasks on the managed system. HMC roles are either *predefined* or *customized*. When you create an HMC user, you must assign that user a task role. Each task role allows the user varying levels of access to tasks that are available on the HMC interface.

You can assign managed systems and logical partitions to individual HMC users, allowing you to create a user that has access to managed system A but not to managed system B. Each grouping of managed resource access is called a *managed resource role*.

Table 4-3 lists the predefined HMC roles, which are the default on the HMC.

Table 4-3 Predefined HMC roles

User name	Role	Description
hmcoperator	Operator	The operator is responsible for daily system operation.
hmcsuperadmin	Super Administrator	The super administrator acts as the root user or manager of the HMC system. The super administrator has unrestricted authority to access and modify most of the HMC system.
hmcpe	Product Engineer	A product engineer helps in support situations but cannot access HMC user management functions. To provide support access for your system, you must create and administer user IDs with the product engineer role.
hmcservicerep	Service Representative	A service representative is an employee who is at your location to install, configure, or repair the system.
hmcviewer	Viewer	A viewer can view HMC information, but cannot change any configuration information.

In the Administration section of the HMC Management Task, there are options for **Managing users profiles and access** (4.2.1, “Managing user profiles and access” on page 114), **Changing the user password** (4.2.2, “Changing the user password” on page 114), and **Customizing user task roles and managed resource roles** (4.2.3, “Customizing user task roles and managed resource roles” on page 115).

4.2.1 Managing user profiles and access

This option allows you to add, copy, remove, and modify HMC system users and user profiles. The administrative functions display in a drop-down menu from the User menu, as shown in Figure 4-9.

To use this function, select **HMC Management** → **Manage User Profiles and Access**.

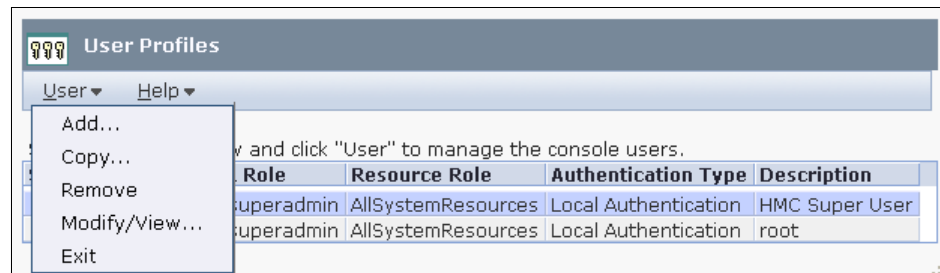


Figure 4-9 User Profiles window

4.2.2 Changing the user password

This option allows you to change the password of the current user, as shown in Figure 4-10 on page 115. The current password is needed for this option and the new password must be different from the current password.

To use this function, select **HMC Management** → **Change User Password**.

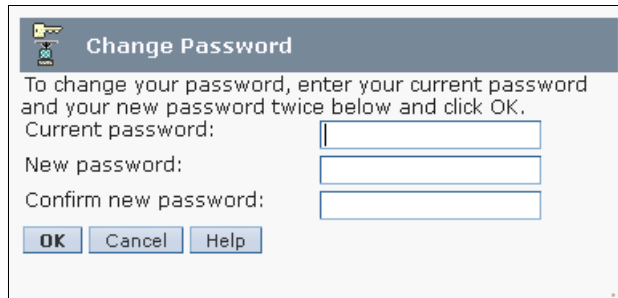


Figure 4-10 Change User Password window

4.2.3 Customizing user task roles and managed resource roles

You can customize HMC Task Roles and Managed Resource Roles through the HMC console. You can add new Task Roles and Managed Resource Roles that are based on existing roles in the HMC. System defined roles cannot be modified, but you can create a role that is based on a system defined role or existing role.

To manage access task and resource roles, select **HMC Management** → **Administration** → **Manage Access Task and Resource Roles**. The Customize User Controls window is displayed, as shown in Figure 4-11.



Figure 4-11 Customize User Controls window

Managed resource roles tasks

A managed resource role assigns permissions for a managed object or group of objects, such as a managed system or a logical partition. In a managed resource role, you can define access to specific managed systems rather than all managed systems controlled by the HMC.

You can create a managed resource role, copy an existing managed resource role, modify existing managed resource roles, or delete an existing managed resource role from the Customize User Controls window. Select **Managed Resource Roles**, then select the wanted operation from the Edit menu. By default, there is only one managed resource role: it is *AllSystemResources*.

To create a managed resource role:

1. Click **Edit** → **Add**, and the Add Role window displays.
2. Enter the name for the new managed resource role, and choose the resource role from which the new managed resource role objects will be based.
3. Select which object is available for the new managed resource role, then click **Add** to add them to the new managed resource role current objects.

Click **OK** to create a managed resource role.

Figure 4-12 shows an example of creating a managed resource role.

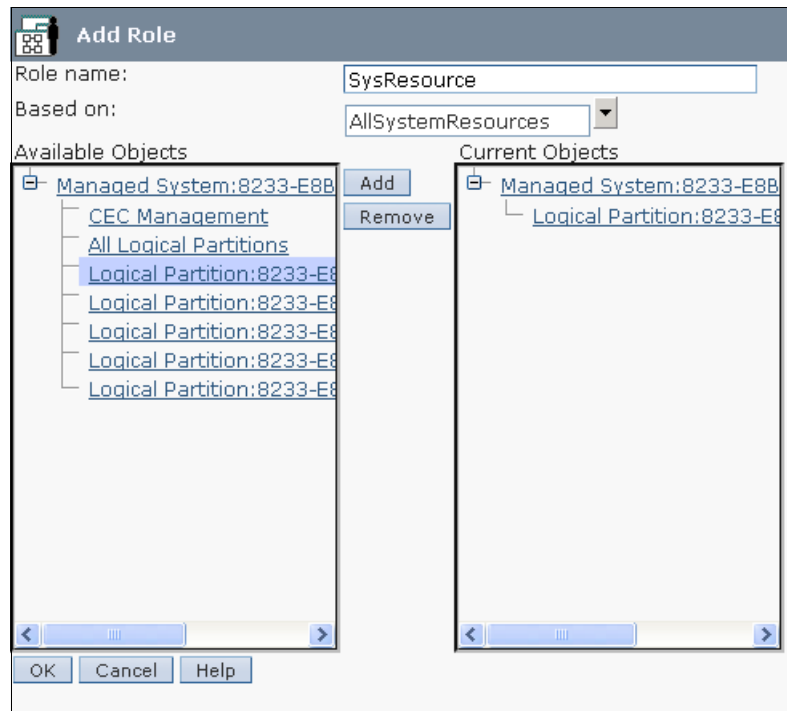


Figure 4-12 Add a new managed resource role

To copy a managed resource role, select the wanted managed resource role and select **Edit** → **Copy**. You cannot copy a user defined managed system role that is created from the **Add** menu, but you can copy system defined managed resource roles, which are AllSystemRoles. From the Copy Role window, you can also customize the object configurations for a new copy of a managed resource role.

To delete a managed resource role, select the wanted managed resource role and select **Edit** → **Remove**. A message box displays asking for *Yes/No* verification.

To modify existing managed resource roles, select a managed resource role that you want to change, and select **Edit** → **Modify**. You can change the configuration of the objects, then click **OK** to save the changes.

Creating, copying, modifying, or deleting task roles

A task role defines the access level for a user to do tasks on the managed object or group of objects, such as a managed system or logical partition. There are five system defined task roles:

- ▶ hmcshericerep
- ▶ hmcviewer
- ▶ hmcoperater
- ▶ hmcpe
- ▶ hmcshericerep

You can create a task role, copy an existing task role, modify an existing task role, or delete an existing task role from the Customize User Controls window. You cannot modify or remove system defined task roles. Select **Task Roles**, then select the wanted operation from the **Edit** menu.

To create a user task role:

1. Click **Edit** → **Add**, and the Add Role window displays.
2. Enter the name for the new managed resource role, and choose the task role from which the new task role objects will be based.
3. Select which object is available for the new task role, and then click **Add** to add them to new task role current objects.
4. Click **OK** to create a task role.

To copy a task role, select the wanted task role and select **Edit** → **Copy**. From the Copy Role window, you can also customize the object configurations for a copy of the task role.

To delete a task role, select the desired task role and select **Edit** → **Remove**. A message box displays asking for *Yes/No* verification.

4.3 Certificate Management

This chapter describes Certificate Management within the Hardware Management Console (HMC) environment.

Manage Certificates is in the Administration section of HMC Management.

Security certificates ensure that the HMC can operate securely in the client/server mode. The managed machines are servers and the managed users are clients. Servers and clients communicate over the Secure Sockets Layer (SSL) protocol, which provides server authentication, data encryption, and data integrity.

When a user wants remote access to the HMC user interface through a web browser, the user requests the secure page by `https://hmc_hostname`. The HMC then presents its certificate to the remote client (web browser) when establishing connection with the HMC. The browser verifies that the certificate was issued by a trusted party, checks that the dates are still valid, and ensures that the certificate was created for that specific HMC.

Archive default certificate: If you are new to the topic of certificate management, it is recommended that you archive the default certificate. To archive a certificate, see 4.3.3, “Advanced options for modifying existing certificates” on page 121.

The available options in Manage Certificates allow you to create, modify, import, and remove certificates.

4.3.1 Create a certificate

You can create a self-signed certificate or a certificate that is signed by a trusted third party. By default, the HMC comes with a self-signed certificate. Follow these steps to create a certificate that is signed by a certificate authority:

1. Select **HMC Management** → **Manage Certificates** → **Create** → **New Certificate**, as shown in Figure 4-13.

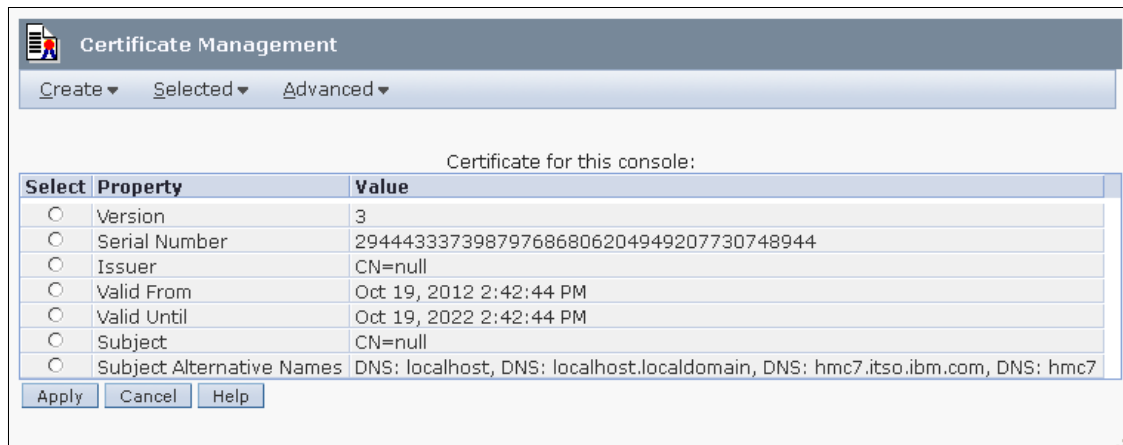


Figure 4-13 Create New Certificate

2. You are given the option of creating a self-signed certificate or a certificate that is signed by a certificate authority.

3. The HMC displays the **New Certificate** window. Complete the New Certificate form and click **OK**.
4. A window displays prompting you for the certificate to be stored. You have the option of storing the certificate on removable media on the console or on the file system on the system running the browser. Make your choice to continue.
5. A message box displays asking for Save verification. Click **OK** to save the Certificate Signing Request as a file. You are then prompted if you want to use a temporary self-signed certificate until your certificate is signed and returned.

Clicking **Yes** creates a self-signed certificate.

You are returned to the **Manage Certificates** window that is shown in Figure 4-13 on page 119. Many of the values will display as Not available.

6. Click **Apply** to apply the new self-signed certificate. *These values are updated after the certificate is applied and the console is restarted.* The next window asks for verification to replace the current certificate.
7. Click **Yes** to proceed. You are then presented with a message box asking if the certificate was replaced successfully or if any errors occurred.
8. Click **OK**. *Clicking OK restarts the console.*
9. After your certificate request is signed and returned, you have to import the certificate and apply by clicking **HMC Management** → **Manage Certificates** → **Advanced** → **Import Certificate**. After the certificate is imported, apply it and restart the console.

4.3.2 Modifying existing certificates

You can modify certain properties of an existing certificate. To modify a certificate, from the Manage Certificates window, select the radio button of the entry that you want to modify, then click **Selected** → **Modify**. Modifiable properties include the following components:

- ▶ Valid Until
- ▶ Subject
- ▶ Subject Alternative Names

For example, to modify the Valid Until property, select the radio button for that property and then select **Selected** → **Modify**.

4.3.3 Advanced options for modifying existing certificates

There are several advanced options available for working with certificates under the Advanced tab. You can do the following actions:

- ▶ Delete and Archive Certificate

You can remove the current certificate. After deleted, the certificate is archived on the HMC.

- ▶ Work with Archived Certificate

You can view and restore archived certificates.

To restore an archived certificate, select **Actions** → **Install**. A window displays asking for verification for restoring the certificate. Click **Yes** to proceed. *This action restarts the console if the installation is successful.*

- ▶ Import certificate

You can import a certificate from media or a remote file system. Select the location of the certificate to import. When the certificate is uploaded, you have to apply and restart the console.

- ▶ View Issuer certificate

Displays available information about the issuer of the certificate.

4.4 Virtualization using HMC

The Hardware Management Console (HMC) is a system that controls Power Systems servers (also called managed systems). It uses a web-browser interface or command-line interface to create and manage logical partitions (LPARs).

LPARs are a virtualized subset of the hardware resources of a physical computer. An IBM Power Systems server and blade can be partitioned into multiple LPARs. Each logical partition has a separate operating system: AIX, Linux, and IBM i.

4.4.1 Processor virtualization

The virtualization of physical processors in IBM Power Systems introduces an abstraction layer that is implemented within the IBM POWER Hypervisor™. The POWER Hypervisor abstracts the physical processors and presents a set of virtual processors to the operating system within the micro-partitions on the system. A micro-partition can have a processor entitlement from a minimum of 0.1 (0.05 with POWER7+ technology-based servers) of a processor up to the

total processor capacity in the system. The granularity of processor entitlement is 0.01 of a processor, allowing entitlement to be precisely determined and configured.

In contrast, dedicated-processor LPARs can be only allocated whole processors, so the maximum number of dedicated-processor LPARs in a system is equal to the number of physical activated processors.

4.4.2 Memory virtualization

IBM Power Systems provide *dedicated* memory allocation, and two features, *Active Memory Sharing* and *Active Memory Expansion* for memory virtualization to increase the flexibility and overall usage of physical memory.

- ▶ **Dedicated**

The physical memory is distributed among the partitions.

- ▶ **Active Memory Sharing**

Active Memory Sharing (AMS) enables the sharing of a pool of physical memory among AIX, IBM i, and Linux partitions on a single IBM POWER6 technology-based server or later, helping to increase memory utilization and drive down system costs. The memory is dynamically allocated among the partitions as needed to optimize the overall physical memory usage in the pool.

- ▶ **Active Memory Expansion**

Active Memory Expansion is the ability to expand the memory that is available to a POWER7 AIX partition beyond the amount of assigned physical memory. Active Memory Expansion compresses memory pages to provide more memory capacity for a partition. The POWER7+ processor chip embeds a new hardware accelerator for AIX memory compression. This tool offloads compression work from processor (as it is on POWER7) cores from doing this task and improves overall performance of the server.

4.4.3 Virtual I/O

Virtual I/O describes the ability to share physical I/O resources between partitions in the form of *virtual adapters* that are in the managed system. Each logical partition typically requires one I/O slot for disk attachment and another I/O slot for network attachment. In the past, these I/O slot requirements would have been physical requirements. To overcome these physical limitations, I/O resources are shared with virtual I/O. With *Virtual Ethernet*, the physical Ethernet adapter is not required to communicate between LPARS. *Virtual Small Computer System Interface (SCSI)* provides the means to share I/O resources for SCSI

storage devices. *Virtual Fibre Channel (NPIV)* provides the means to share a Fibre Channel Adapter for SAN storage and tape devices.

POWER Hypervisor for virtual I/O

The POWER Hypervisor (PHYP) provides the interconnection for the partitions. To use the functionalities of virtual I/O, a partition uses a virtual adapter as shown in Figure 4-14. The PHYP provides the partition with a view of an adapter that has the appearance of an I/O adapter, which might or might not correspond to a physical I/O adapter.

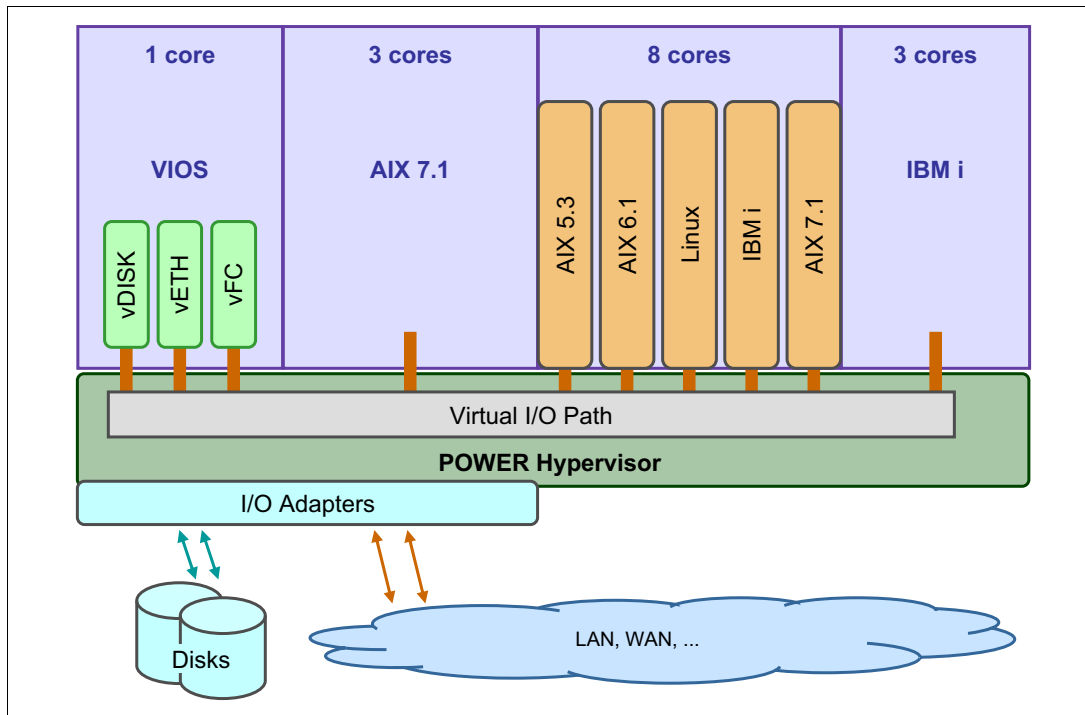


Figure 4-14 Role of PHYP for virtual I/O

Virtual I/O Server

The Virtual I/O Server (VIOS) can link the physical resources to the virtual resources. By this linking, it provides virtual storage and Shared Ethernet Adapter capability to client logical partitions on the system. It allows physical adapters with attached disks on the Virtual I/O Server to be shared by or more client partitions.

Virtual I/O Server mainly provides two functions:

- ▶ Serves virtual SCSI devices to clients, which are described in “Virtual Small Computer System Interface” on page 125.
- ▶ Serves virtual Fibre Channel devices to clients, which are described in “Virtual Fibre Channel” on page 128.
- ▶ Provides a Shared Ethernet Adapter for virtual Ethernet, which is described in “Virtual Ethernet” on page 129.

Virtual I/O Server partitions are not intended to run applications or for general user logins. The Virtual I/O Server is installed in its own partition. The Virtual I/O Server partition is a special type of partition, which is marked as such on the first window of the Create Logical Partitioning Wizard program.

Currently, the Virtual I/O Server is implemented as a customized AIX partition. However, the interface to the system is abstracted by using a secure shell-based command-line interface (CLI). When a partition is created as this type of partition, only the Virtual I/O Server software boot image boots successfully when the partition is activated.

Configure the Virtual I/O Server with enough resources. If a Virtual I/O Server has to host numerous resources to other partitions, ensure that enough processor power and memory are available.

Rule of Thumb: Sizing the Virtual I/O Server: See this IBM developerWorks® website for Nigel Griffiths AIXpert Blog:

https://www.ibm.com/developerworks/mydeveloperworks/blogs/aixpert/entry/rule_of_thumb_sizing_the_virtual_i_o_server78?lang=en

VIOS Performance Advisor: The VIOS Advisor is an application that runs within the client's Virtual I/O Server (VIOS) for a user specified amount of time (hours), which polls and collects key performance metrics before analyzing results. The Advisor also provides a health check report and proposes changes to the environment or areas to investigate further. For more information, see this website:

<https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/Power%20Systems/page/VIOS%20Advisor>

Virtual Small Computer System Interface

Virtual SCSI is based on a client/server relationship. A Virtual I/O Server partition owns the physical resources, and logical client partitions access the virtual SCSI resources that are provided by the Virtual I/O Server partition. The Virtual I/O Server partition has physically attached I/O devices and exports one or more of these devices to other partitions, as shown in Figure 4-15.

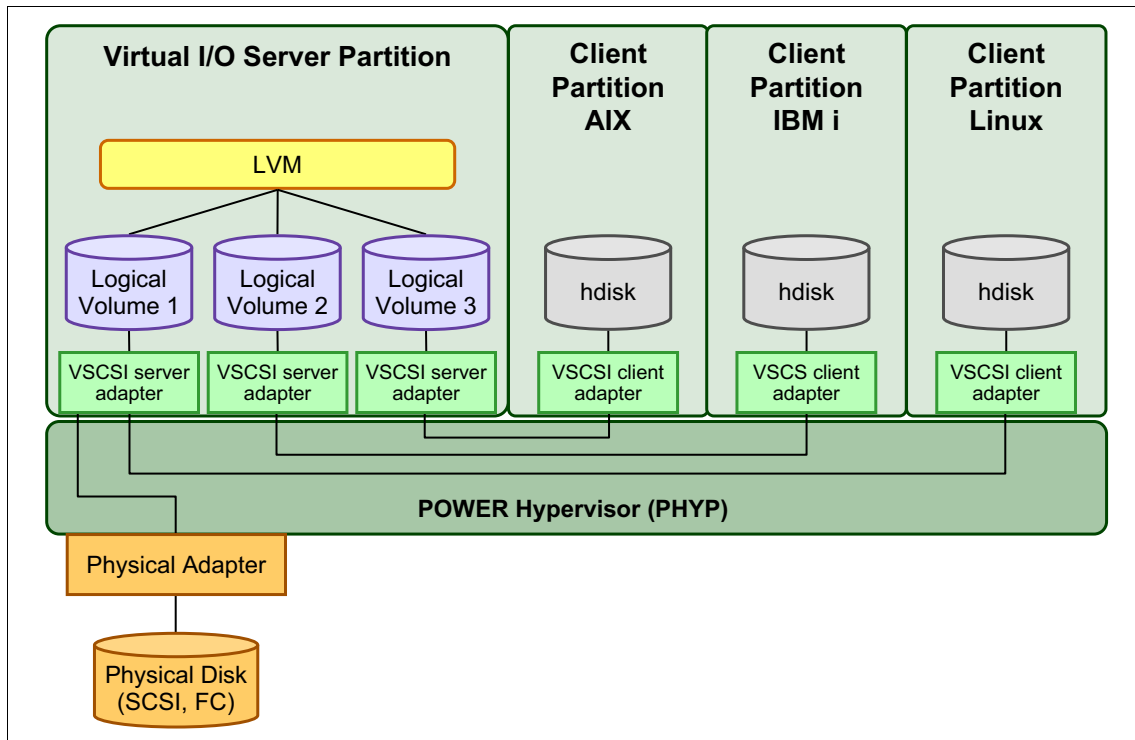


Figure 4-15 Virtual SCSI overview

The client partition is a partition that has a virtual client adapter node that is defined in its device tree and relies on the Virtual I/O Server partition to provide access to one or more block interface devices. Virtual SCSI requires POWER5, POWER6, or POWER7 hardware with the *PowerVM* feature activated.

Client/server communications

In the Figure 4-15, the virtual SCSI adapters on the server and the client are connected through the hypervisor. The virtual SCSI adapter drivers (server and client) communicate control data through the hypervisor.

When data is transferred from the backing storage to the client partition, it is transferred to and from the client's data buffer by the DMA controller on the

physical adapter by using redirected SCSI Remote Direct Memory Access (RDMA) Protocol. This facility enables the Virtual I/O Server to securely target memory pages on the client to support virtual SCSI.

Adding Virtual SCSI adapter

You can create the virtual adapters in two periods. One is to create those adapters during that installation of the Virtual I/O Server. The other is to add those adapters in an existing Virtual I/O Server. In this chapter, we suppose that we already created the Virtual I/O Server.

Before activating a server, you can add the virtual adapter by using the Manage Profiles task. For an activated server, you can do that only through dynamic LPAR operation if you want to use virtual adapters immediately. This procedure requires that the network is configured with connection to the HMC to allow for dynamic LPAR.

Add/Remove commands: With HMC V7R760, during DLPAR addition or removal of virtual I/O adapters to or from a Virtual I/O Server, the HMC now automatically attempts to run the Add/Remove commands (`cfgdev/rmdev`) on the Virtual I/O Server. In earlier releases, you manually had to run these commands on the Virtual I/O Server.

Now, you can add the adapter through DLPAR. To add the adapter:

1. Select the activated Virtual I/O Server partition in HMC. Then, click **Virtual Adapters** in the Dynamic Logical Partitioning section in the Task pane. The Virtual Adapters window opens.
2. Click **Actions** → **Create Virtual Adapter** → **SCSI Adapter**.

- In the next window, you create the new Virtual SCSI adapter, as shown in Figure 4-16. If the clients are not known to the HMC, select **Any client partition can connect**. If you select this option, you have to change the client partition options to the correct name of the client after you create the clients. If you know which client partition is connected, select **Only selected client partition can connect**. Then, choose the client adapter ID number. Click **OK**.

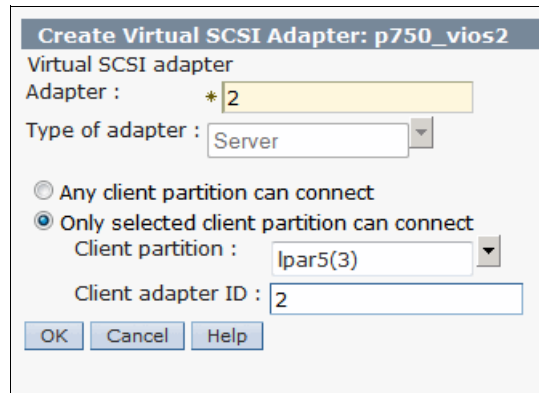


Figure 4-16 Create Virtual SCSI Adapter panel

- Now, you can see the new virtual SCSI adapter in the Virtual Adapters window.
In our example, we set up the adapter to *lpar5* partition with *Adapter ID 2*.
- Run the **cfgdev** command in the Virtual I/O Server (no need to do this command if your HMC is with V7R760 firmware level) to configure the newly created virtual SCSI server device (*vhost#*), and map this device to the SCSI device (disk, logical volume, tape) by using the **mkvdev** Virtual I/O Server command. For more details, see this website:

<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7hcg/mkvdev.htm>

Update the LPAR profile: If you dynamically created a virtual adapter, do not forget to update the LPAR profile with the same values, *Client partition name*, *Adapter ID* and *Client adapter ID*. Otherwise, adapters will disappear after partition off/on.

Virtual Fibre Channel

Just as with the virtual SCSI adapters, virtual Fibre Channel adapter, N_Port ID Virtualization (NPIV), is based on a client/server relationship (see Figure 4-17) and communicates control data through PHYP. Because virtual SCSI is block device (disk, LV, tape) virtualization, NPIV is a Fibre Channel industry standard method for virtualizing a physical Fibre Channel port.

On a Power Systems server, NPIV allows LPARs to have dedicated N_Port ID (WWN), giving the OS a unique identity to the SAN, just as though it had a dedicated physical host bus adapter (HBA). NPIV is supported by AIX, IBM i, and Linux partitions.

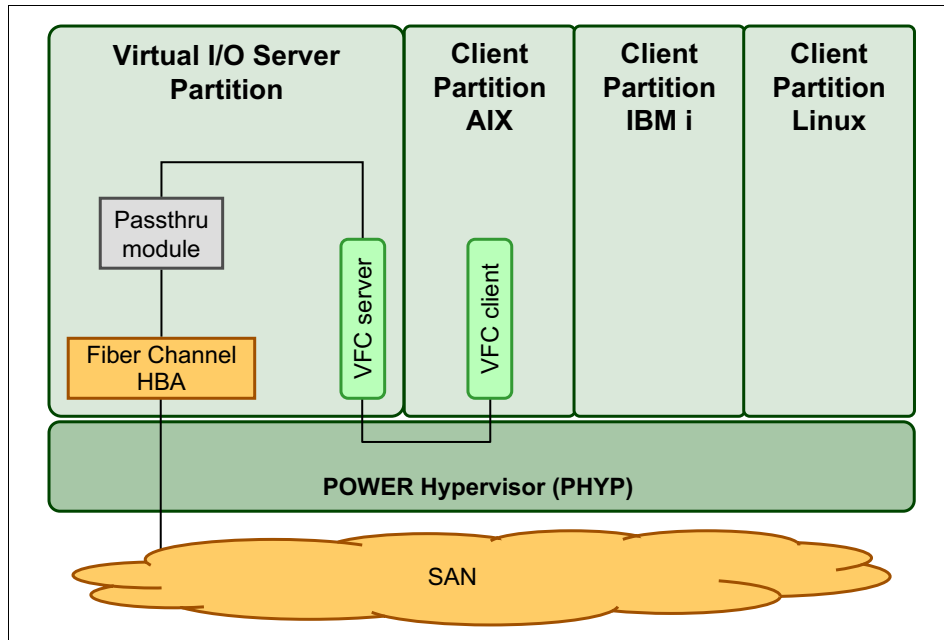


Figure 4-17 NPIV overview

NPIV prerequisites: Prerequisites for NPIV on POWER6 and POWER7 technology-based servers are, **FC5735** (8 Gb PCIe Dual Port Fibre Channel Adapter) or **FC5708** (10 Gb FCoE PCIe Dual Port adapter), AIX 5.3 TL9 and higher, VIOS 2.1 FP 20.1 and higher, and SAN or FCoE switches that support NPIV.

Adding a virtual Fibre Channel adapter

Now, you can add the adapter through DLPAR. To add the adapter:

1. Select the activated Virtual I/O Server partition in HMC. Then, click **Virtual Adapters** in the Dynamic Logical Partitioning section in the Task pane. The Virtual Adapters window opens.
2. Click **Actions** → **Create Virtual Adapter** → **Fibre Channel Adapter**.
3. In the next window, you create the virtual Fibre Channel (vFC) adapter, select **Client Partition** name and the client vFC adapter **ID** number. Ensure that you create a client vFC adapter before you create a vFC adapter on the Virtual I/O Server.

Click **OK**.

4. Now, you can see the new vFC adapter in the Virtual Adapters window
5. Run the **cfgdev** command in the Virtual I/O Server (no need to do this command if your HMC is with V7R760 firmware level) to configure the newly created vFC server device (*vfchost#*), and map this device to a physical HBA by using the **vfcmmap** Virtual I/O Server command. For more information, see this website:

<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7hcg1/vfcmmap.htm>

Virtual Ethernet

Virtual Ethernet enables inter-partition communication without having physical network adapters that are assigned to each partition. It can be used in both shared and dedicated POWER5, POWER6, and POWER7 processor partitions. Virtual Ethernet can be used in these processors if the partition is running AIX 5L V5.3 or higher, IBM i 6.1.1 or higher, or Linux with the 2.6 or higher that supports virtualization. This technology enables IP-based communication between logical partitions on the same system by using a Virtual LAN capable software switch (POWER Hypervisor).

Because of the number of partitions possible on many systems being greater than the number of I/O slots, virtual Ethernet is a convenient and cost saving option to enable partitions within a single system to communicate with one another through a virtual Ethernet LAN. These connections exhibit characteristics that are similar to physical high-bandwidth Ethernet connections and support multiple protocols (IPv4, IPv6, and Internet Control Message Protocol (ICMP)).

Virtual Ethernet does not require the purchase of any additional features or software, such as the PowerVM feature. Virtual Ethernet is different from Shared Ethernet adapter in that, there is no connection to a physical Ethernet adapter

which connects to a physical Ethernet network. To use virtual Ethernet to connect to a physical Ethernet adapter, which connects to a physical Ethernet network, you must implement Shared Ethernet adapter.

Virtual LAN overview

Virtual LAN (VLAN) is a technology that is used for establishing virtual network segments on top of physical switch devices. Multiple VLAN logical devices can be configured on a single system, as shown in Figure 4-18. Each VLAN logical device constitutes an extra Ethernet adapter instance. These logical devices can be used to configure the same types of Ethernet IP interfaces as are used with physical Ethernet adapters.




Figure 4-18 VLAN example: Two VLANs

Virtual Ethernet connection

Virtual Ethernet connections supported in POWER5, POWER6, and POWER7 processor-based systems use VLAN technology to ensure that the partitions can access only data that is directed to them. The POWER Hypervisor provides a virtual Ethernet switch function that is based on the IEEE 802.1Q VLAN standard that enables partition communication within the same server, as shown in Figure 4-19. The connections are based on an implementation internal to the Hypervisor that moves data between partitions.

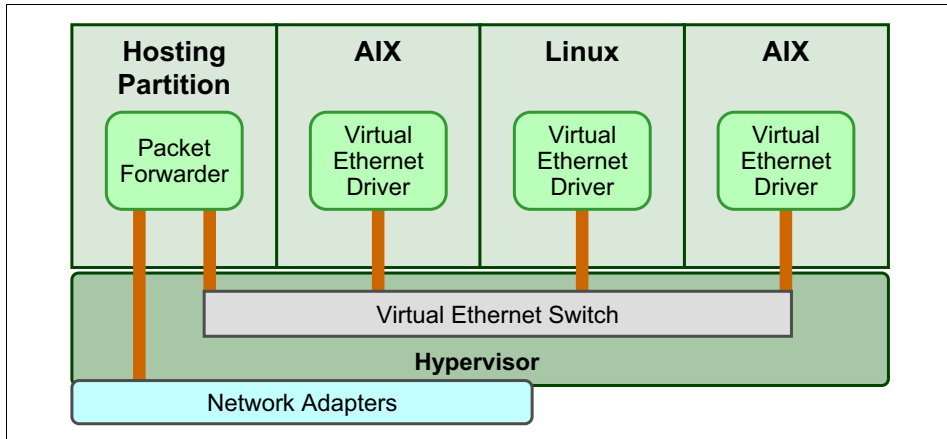


Figure 4-19 Virtual Ethernet connection

Adding virtual Ethernet adapters

You can create virtual Ethernet adapters in the same manner as creating a virtual SCSI adapter, as described in “Adding Virtual SCSI adapter” on page 126.

Shared Ethernet Adapter

Virtual I/O Server partition is not required for implementing a VLAN. Virtual Ethernet adapters can communicate with each other through the POWER Hypervisor without the functionality of the Virtual I/O Server. A Shared Ethernet Adapter bridges external networks to internal VLANs. The Shared Ethernet Adapter hosted in the Virtual I/O Server partition acts as an OSI Layer 2 switch between the internal and external network.

Figure 4-20 shows the Shared Ethernet Adapter that is used as a bridge between the virtual Ethernet and physical Ethernet.

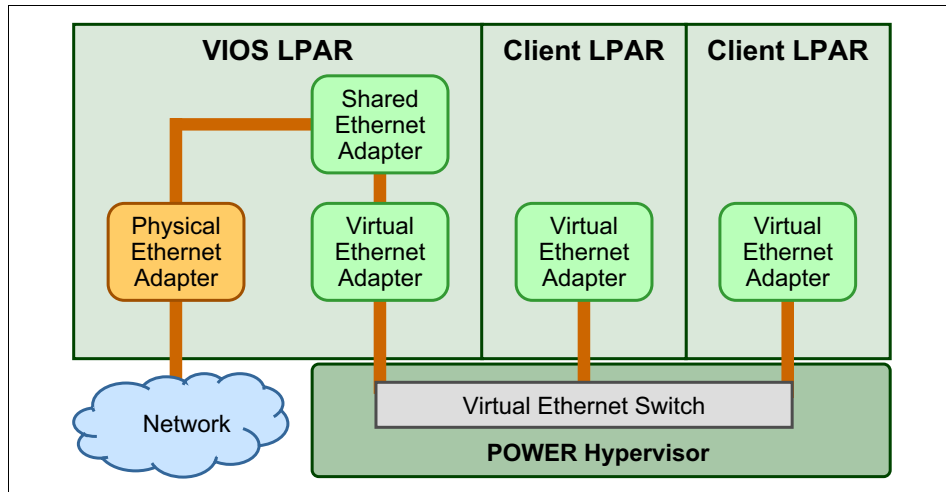


Figure 4-20 Shared Ethernet Adapter configuration

The bridge interconnects the logical and physical LAN segments at the network interface layer level and forwards frames between them. The bridge performs the function of a MAC relay (OSI Layer 2) and is independent of any higher layer protocol. Figure 4-21 is a close-up view of the Virtual I/O Server partition.

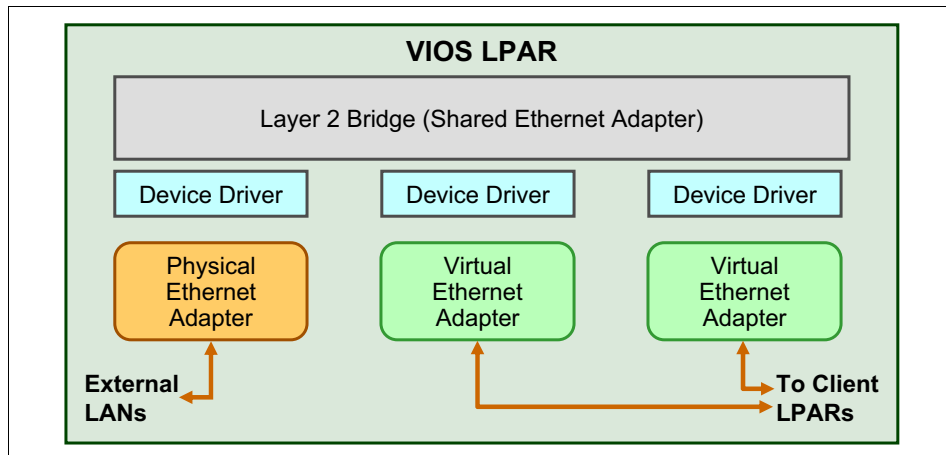


Figure 4-21 Shared Ethernet Adapter OSI layer

The bridge is transparent to the Internet Protocol (IP) layer. For example, when an IP host sends an IP datagram to another host on a network that is connected

by a bridge, it sends the datagram directly to the host. The datagram crosses the bridge without the sending IP host being aware of it.

The Virtual I/O Server partition offers broadcast and multicast support. Address Resolution Protocol (ARP) and Neighbor Discovery Protocol (NDP) also work across the Shared Ethernet Adapter.

The Virtual I/O Server does not reserve bandwidth on the physical adapter for any of the VLAN clients that send data to the external network. Therefore, if one client partition of the Virtual I/O Server sends data, it can take advantage of the full bandwidth of the adapter. This scenario assumes that the other client partitions do not send or receive data over the network adapter at the same time.

4.4.4 Live Partition Mobility

Live Partition Mobility, a component of the PowerVM Enterprise Edition hardware feature, can move AIX, IBM i, and Linux logical partitions from one system to another. The mobility process transfers the system environment that includes the processor state, memory, attached virtual devices, and connected users.

With *active partition mobility*, move AIX, IBM i, and Linux logical partitions that are running, including the operating system and applications, from one system to another. The logical partition and the applications that run on that migrated logical partition do not need to be shut down.

With *inactive partition mobility*, move a powered-off AIX, IBM i, or Linux logical partition from one system to another.

You can use the HMC to move an active or inactive logical partition from one server to another.

Because the HMC always moves the last activated profile, an inactive logical partition that has never been activated cannot be moved. For inactive partition mobility, you can either select the partition state that is defined in the hypervisor, or select the configuration data that is defined in the last activated profile on the source server.

4.4.5 Host Ethernet Adapter

Host Ethernet Adapter (HEA) is a physical Ethernet adapter that is integrated directly into the system bus on a managed system. HEAs offer high throughput, low latency, and virtualization support for Ethernet connections. HEAs are also known as Integrated Virtual Ethernet adapters (IVE adapters). Multiple logical

partitions can connect directly to the HEA and use the HEA resources. This configuration allows these logical partitions to access external networks through the HEA without having to go through an Ethernet bridge (e.g. “Shared Ethernet Adapter” on page 131) on another logical partition (e.g. “Virtual I/O Server” on page 123).

To get to the Host Ethernet Adapter area:

1. In the HMC workplace window, select **Systems Management** → **Servers**, then select the name of the server. Select **Hardware** → **Adapters** → **Host Ethernet** to open the window.
2. Select the Ethernet adapter that you want to configure, and then select **Configure**.
3. In the **HEA Physical Port Configuration** window, you can set the network adapter speed, duplex mode, packet size, and default partition control for the integrated controller. These numbers depend on the type of network switch to which you are connecting the server. When you finish, select **OK** and your selections are saved.

Models with HEA: HEA is available on POWER6 and POWER7 models, 9179-MHB, 9117-MMB, 8233-E8B, 8205-E6B, 8202-E4B, and 8231-E2B. HEA is no longer an option on POWER7, POWER7+ and their follow on systems.

4.4.6 Shared pool usage of dedicated capacity

Beginning with POWER6, HMC V7 R3 allows for the shared use of dedicated processing resources. The option to donate resources is turned off by default when partitions are created as *Dedicated* and can be configured through the partition property window. (To read about how to create a dedicated partition, see “Configuring a dedicated processor partition” on page 146.

To verify that your system can share dedicated capacity:

1. In the HMC workplace window, select **Systems Management** → **Servers**, then select the name of the server. Select **Tasks** → **Properties**, click the **Capabilities** tab, scroll down, and see if the value of capability **Active Partition Processor Sharing Capable** is **True**.
2. To configure the managed server as a processor donor and open the window as shown in Figure 4-22 on page 135, select **Systems Management**, then select the name of the server. Select the name of the partition to view the partition properties.

Go to the Hardware tab to view the settings for processors, memory, and I/O.

3. On the Processors tab, you can select the radio buttons for when you want to allow processor sharing for this particular partition. In this window, *inactive* and *active* refer only to the partition's activation state. Before POWER6, dedicated processors were not shared with other partitions, even if the processors were idle. With POWER6 systems, you can now share idle processing power from dedicated processors.

When the LPAR with dedicated processors is inactive, the processors are always idle. With this version of the HMC and POWER6 systems, you can also share the idle processor cycles to the shared processor pool when the partition is active. This method gives you the performance benefit of configuring dedicated processors to a partition while providing the server utilization benefit of sharing idle resources with other partitions:

- Allow when partition is inactive

When this option is selected, the dedicated resources for this partition are allocated to other active partitions for their shared processor usage.

- Allow when partition is active

When this option is selected, as processors become idle on this partition, the idle processors are allocated to other active partitions for their shared processor usage.

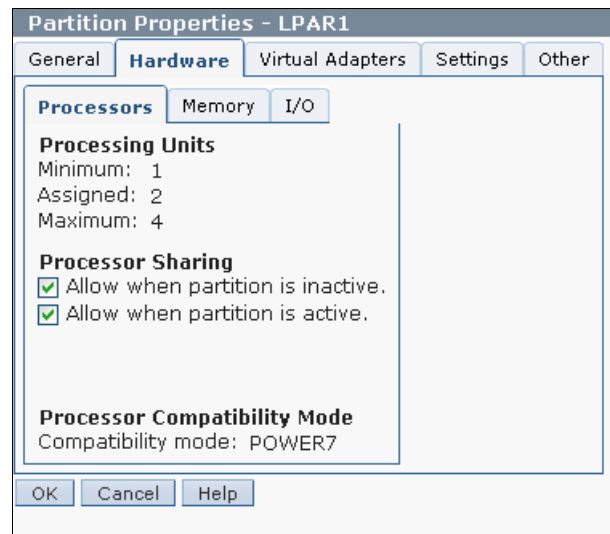


Figure 4-22 Processor sharing

4.4.7 Multiple Shared Processor Pool

Starting with POWER6 processor-based systems, the capability of multiple shared processor pools allows a systems administrator to create Micro-Partitioning with the purpose of controlling the processor capacity that can be consumed from the physical shared-processor pool. For more information about this topic, see *PowerVM Virtualization Introduction and Configuration*, SG24-7940.

4.4.8 Suspend logical partition

With the HMC version 7 release 7.2.0, or later, you can suspend a running AIX or Linux logical partition with its operating system and applications. And with the HMC version 7 release 7.3.0 or later: IBM i logical partition with its operating system and applications.

Only POWER7 processor-based servers support the Suspend/Resume feature. When a logical partition is suspended, the state of the logical partition is saved on persistent storage, and the server resources that were in use by that logical partition are made available for use by other logical partitions. At a later time, the operation of the suspended logical partition and its applications can be resumed. For more information about this topic, see *PowerVM Virtualization Introduction and Configuration*, SG24-7940.

4.4.9 Partition availability priority

Starting with POWER6 technology-based servers, and with the HMC V7, the concept of *partition availability priority* configuration option allows you to set up a hierarchy of partitions to cover for the event of a processor failure and ensures that high priority partitions have a higher guarantee of processor access than other partitions when a processor fails.

To access the window that is shown in Figure 4-23 select **Systems Management** → **Servers** → **Add the Tick to your Servers**, then select the name of the server. Select **Configuration System Plans** → **Partition** → **Availability Priority**.

Partition Availability Priority: 8233-E8B-SN10DD51P

You can change the partition availability priority for the following partitions by first selecting one or more partitions and then choosing an availability priority from the field below the table. Click OK to submit your changes.

Select	Partition Name	Partition Type	Processing units	Processing Mode	Availability priority
<input checked="" type="checkbox"/>	LPAR1	AIX or Linux	2	Dedicated	127
<input type="checkbox"/>	lp1	AIX or Linux	0	Shared	127
<input type="checkbox"/>	lp2	AIX or Linux	0	Shared	127
<input type="checkbox"/>	lp3	AIX or Linux	0	Shared	127
<input type="checkbox"/>	lp4	AIX or Linux	0	Shared	127
<input type="checkbox"/>	lp5	IBM i	0	Dedicated	127
<input type="checkbox"/>	vios1	Virtual I/O Server	0	Shared	191

Availability priority:

- Default (127)
- Minimum (0)
- Low (63)
- High (191)
- Maximum (255)

Figure 4-23 Setting partition availability priority

The managed system uses partition-availability priorities in the case of processor failure. If a processor fails on a logical partition and if there are no unassigned processors available on the managed system, the logical partition can acquire a replacement processor from logical partitions with a lower partition-availability priority. This process allows the logical partition with the higher partition-availability priority to continue running after a processor failure.

When a processor fails on a high-priority logical partition, the managed system follows these steps to acquire a replacement processor for the high-priority logical partition:

1. If there are unassigned processors on the managed system, the managed system replaces the failed processor with an unassigned processor.
2. If there are no unassigned processors on the managed system, the managed system checks the logical partitions with lower partition-availability priorities, starting with the lowest partition-availability priority.

3. If a lower-priority logical partition uses dedicated processors, the managed system shuts down the logical partition and replaces the failed processor with one of the processors from the dedicated-processor partition.
4. If a lower-priority logical partition uses shared processors, and removing a whole processor from the logical partition would not cause the logical partition to go below its minimum value, the managed system removes a whole processor from the shared-processor partition. This process is done by using Dynamic Logical Partitioning and replaces the failed processor with the processor that the managed system removed from the shared-processor partition.
5. If a lower-priority logical partition uses shared processors, but removing a whole processor from the logical partition causes the logical partition to go below its minimum value, the managed system skips that logical partition. The system continues to the logical partition with the next higher partition availability.
6. If the managed system still cannot find a replacement processor, the managed system shuts down as many of the shared-processor partitions as it must to acquire the replacement processor. The managed system shuts down the shared-processor partitions in partition-availability priority order, starting with the lowest partition-availability priority.

A logical partition can take processors only from logical partitions with lower partition-availability priorities. If all of the logical partitions on your managed system have the same partition-availability priority, a logical partition can replace a failed processor only if the managed system has unassigned processors.

By default, the partition availability priority of Virtual I/O Server logical partitions and IBM i logical partitions with virtual SCSI adapters is set to 191. The partition-availability priority of all other logical partitions is set to 127, by default, as shown in Figure 4-23 on page 137.

Priority of partitions: Do not set the priority of Virtual I/O Server logical partitions to be lower than the priority of the logical partitions that use the resources on the Virtual I/O Server logical partition.

Do not set the priority of IBM i logical partitions with virtual SCSI adapters to be lower than the priority of the logical partitions that use the resources on the IBM i logical partition.

4.4.10 Logical partition management

There are various ways to create partitions on POWER5, POWER6, and POWER7 processor-based systems by using the HMC.

System Planning Tool: The System Planning Tool and system plans are closely tied to LPAR management. For more information, see the IBM System Planning Tool website:

<http://www.ibm.com/systems/support/tools/systemplanningtool/>

An LPAR is the division of a computer's processors, memory, and hardware resources into multiple environments so that each environment can be operated independently with its own operating system and applications. The number of logical partitions that can be created depends on the system. Typically, partitions are used for different purposes, such as database operation, client/server operations, web server operations, test environments, and production environments. Each partition can communicate with the other partitions as though each partition were a separate machine.

DLPAR can logically attach and detach a managed system's resources to and from an LPAR's operating system without rebooting. DLPAR requests are built from simple add and remove requests that are directed to logical partitions. The user can run these commands as move requests at the HMC, which manages all DLPAR operations. DLPAR operations are enabled by Power Systems PHYP and AIX, IBM i, and Linux.

You can use the HMC graphical-user interface or the command-line interface to create, modify, delete, move, or suspend the LPARs. Each LPAR has one or more profiles that includes the settings that are used when the LPAR is turned on. Multiple profiles allow for the saving of multiple configurations for a single LPAR. This flexibility gives you the ability to configure an LPAR to handle different workloads and to save that information to make it easily repeatable and scheduled.

4.4.11 Create an AIX or Linux logical partition

Options for creating a Virtual I/O Server: The options and window views for creating a Virtual I/O Server partition are the same as those that we present in this section. Thus, we do not document the steps for the Virtual I/O Server.

To create an AIX partition, follow these steps:

1. Select **System Management** → **Servers**.
2. Select *<Managed Server>*.

3. From **Task**, select **Configuration** → **Create Logical Partition** → **AIX or Linux**, as shown in Figure 4-24.

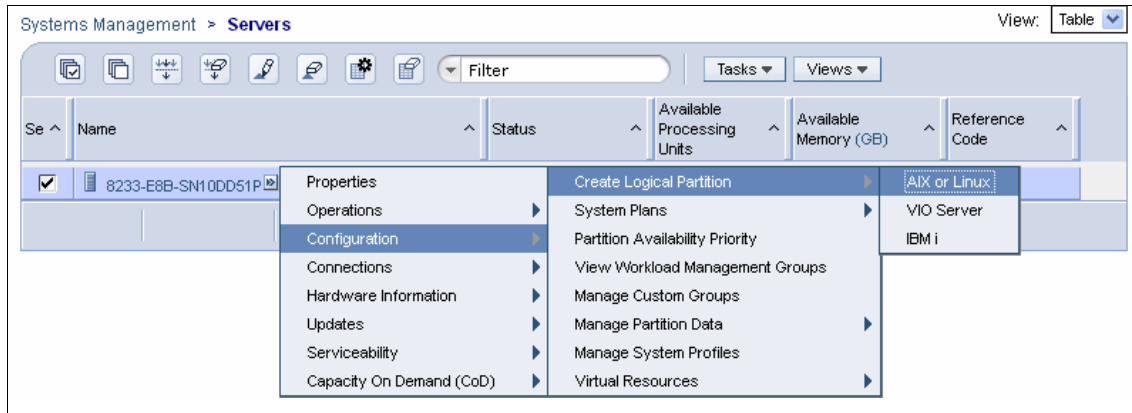


Figure 4-24 Create Virtual I/O Server LPAR

4. You can set **Partition ID** and specify **Partition name**, as shown in Figure 4-25. If you want the partition to be suspended, check the **Allow this partition to be suspended** box. Then, click **Next**.

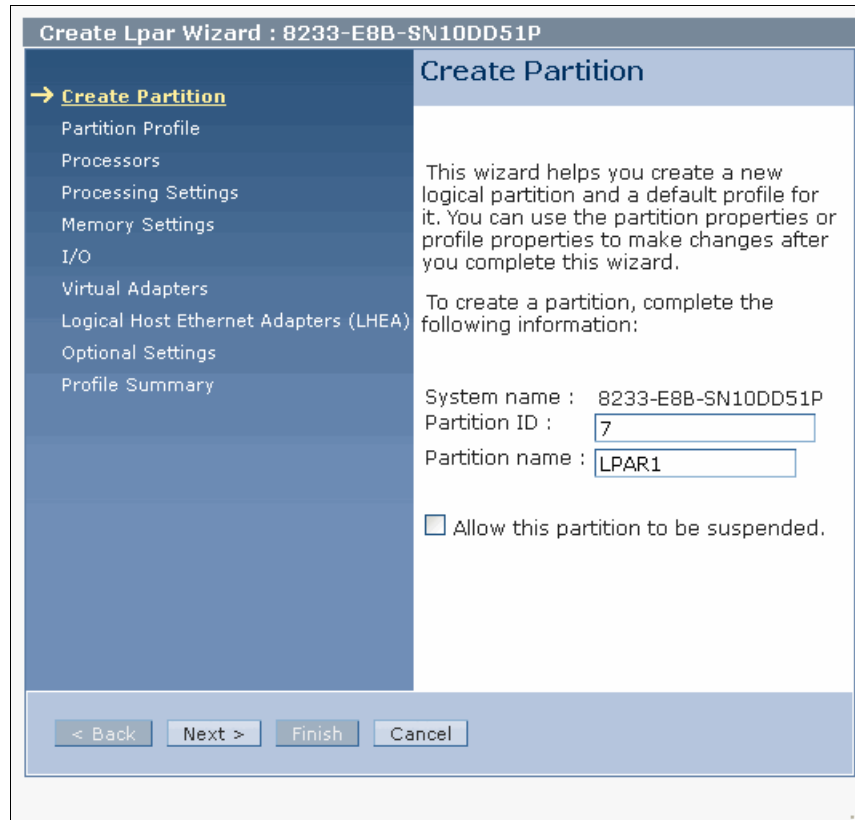


Figure 4-25 Create Partition panel

Virtual I/O clients: VIO clients wait for the Virtual I/O Server to start before continuing to boot. Good practice is to put the VIO clients high in the list (ID > 4).

5. Enter **Profile name** (for example: *default*) for this partition and click **Next**. You can then create a partition with either *shared* or *dedicated* processors on your server.

Configuring a shared processor partition

This section describes how to create a partition with a *shared* processor. If you want to create a partition with a *dedicated* processor, refer to “Configuring a dedicated processor partition” on page 146.

Specify the processing units for the partition and any settings for virtual processors, as shown in Figure 4-26. The sections that immediately follow this figure describe the settings in this figure in detail.

The screenshot shows the 'Create Lpar Wizard' window with the 'Processing Settings' panel selected. The panel contains the following settings:

- Processing Settings:**
 - Total usable processing units: 8.00
 - Minimum processing units: * 0.5
 - Desired processing units: * 1
 - Maximum processing units: * 2
 - Shared processor pool: DefaultPool (0)
- Virtual processors:**
 - Minimum processing units required 0.10 for each virtual processor:
 - Minimum virtual processors: * 1
 - Desired virtual processors: * 2
 - Maximum virtual processors: * 4
 - Uncapped
 - Weight : 128.0

Figure 4-26 Processing Settings panel

Processing Settings area

In the Processing Settings area, you must specify the minimum number of processors that you want the shared processor partition to acquire, the wanted amount, and the maximum upper limit that is allowed for the partition.

The values in each field can range anywhere between 0.1 (0.05 on POWER7+ technology-based servers) and the total number of processors in the managed server and can be any increment in between in tenths of a processor.

Each field defines the following information:

– **Minimum processing units**

The absolute minimum number of processing units that are required from the shared processing pool for this partition to become active. If the number in this field is not available from the shared processing pool, this partition cannot be activated.

This value has a direct bearing on DLPAR, because the minimum processing units value represents the smallest value of processors that the partition can have as the result of a DLPAR deallocation.

– **Desired processing units**

This number must be greater than or equal to the amount set in *Minimum processing units*, and represents a number of processors asked for above the minimum amount. If the minimum is set to 2.3 and the wanted is set to 4.1, then the partition can become active with any number of processors between 4.1 and 2.3, whatever number is greater and available from the shared resource pool.

When a partition is activated, it queries for processing units starting at the wanted value and goes down in 0.1 of a processor until it reaches the minimum value. If the minimum is not met, the partition does not become active.

Wanted processing units govern only the possible number of processing units that a partition can become active with. If the partition is made *uncapped*, then the hypervisor can let the partition exceed its wanted value depending on how great the peak need is and what is available from the shared processing pool.

– **Maximum processing units**

This setting represents the absolute maximum number of processors this partition can own at any specified time, and must be equal to or greater than the *Desired processing units*.

This value has a direct bearing on DLPAR, because the maximum processing units value represents the largest value of processors the partition can have as the result of a DLPAR allocation.

Furthermore, although this value affects DLPAR allocation, it does not affect the processor allocation handled by the hypervisor for idle processor allocation during processing peaks.

Values: Whether your partition is *capped* or *uncapped*, the minimum value for *Maximum processing units* is equal to the value specified for *Desired processing units*.

Uncapped option

The Uncapped option represents whether you want the HMC to consider the partition capped or uncapped. Whether a partition is *capped* or *uncapped*, when it is activated it takes on a processor value equal to a number somewhere between the minimum and wanted processing units, depending on what is available from the shared resource pool. However, if a partition is capped, it can gain processing power only through a DLPAR allocation and otherwise stays at the value given to it at time of activation.

Capped partition: When your partition is *capped* and it does not use the entitled processor capacity, the free processor cycles can be used by *uncapped* partitions if they are in the same processor pool.

If the partition is *uncapped*, it can exceed the value that is set in the *Desired number of processing units* area (while it is running, these units are referred to as the LPAR entitlement) from the shared processor pool that it needs. This value is not seen from the HMC view of the partition, but you can check the value of processors that are owned by the partition from the operating system level with the appropriate commands.

The *weight* field defaults to 128 and can range from 0 to 255. Setting this number below 128 decreases a partition's priority for processor allocation, and increasing it above 128, up to 256, increases a partition's priority for processor allocation.

If all partitions are set to 128 (or another equivalent number), then all partitions have equal access to the shared processor pool. If a partition's *uncapped weight* is set to 0, then that partition is considered *capped*, and it never owns a number of processors greater than that specified in *Desired processing units*.

Virtual processors area

The values that are set in the Virtual processors area of this window govern how many processors to present to the operating system of the partition. You must show a minimum of one virtual processor per actual processor, and you can have as many as 10 (20 in POWER7+ technology-based servers) virtual processors per physical processing unit.

As a general recommendation, a partition requires at least as many virtual processors as you have actual processors, and configure a partition with

no more than twice the number of virtual processors as you have actual processors.

Each field defines the following information:

– **Minimum virtual processors**

Your partition must have at least one virtual processor for every part of a physical processor that is assigned to the partition. For example, if you assigned 2.5 processing units to the partition, the minimum number of virtual processors is three, and the maximum is 25 (50 in a POWER7+ technology-based server).

Furthermore, this value represents the lowest number of virtual processors that can be owned by this partition as the result of a DLPAR operation.

– **Desired virtual processors**

The desired virtual processors value must be greater than or equal to the value set in *Minimum virtual processors*. And as a general guideline, about twice the amount that is set in *Desired processing units*. Performance with virtual processing can vary depending on the application, and you might need to experiment with the wanted virtual processors value before you find the perfect value for this field and your implementation.

Desired virtual processors value: The desired virtual processors value, along with the resources available in the shared resource pool, is the only value that can set an effective limit on the amount of resources that can be used by an uncapped partition.

– **Maximum virtual processors**

You can have only 10 (20 in POWER7+ technology-based server) virtual processors per processing unit. Therefore, you cannot assign a value greater than 10 (20 in POWER7+ technology-based server) times the *Maximum processing units* value as set in “Processing Settings area” on page 142. It is recommended, though not required, to set this number to twice the value that is entered in *Maximum processing units*.

Upper limit of virtual processors: With POWER7 and AIX7, the upper limit of virtual processors is 256. With AIX6, the upper limit is 64 processors.

Finally, this value represents the maximum number of virtual processors that this partition can have as the result of a DLPAR operation.

Configuring a dedicated processor partition

This section describes how to create a partition with a *dedicated* processor. If you want to create a partition with a *shared* processor, refer to “Configuring a shared processor partition” on page 142.

To configure a dedicated processor partition:

1. Select **Dedicated** and then select **Next**, as shown in Figure 4-27.



Figure 4-27 Create dedicated processor partition

- Specify the number of minimum, desired, and maximum processors for the partition, as shown in Figure 4-28, and click **Next**.

Figure 4-28 shows the 'Processing Settings' step in the 'Create Lpar Wizard' for system 8233-E8B-SN10DD51P. The wizard is currently on the 'Processing Settings' step, which is highlighted in the left sidebar. The main area displays the following settings:

- Total number of processors : 8.00
- Minimum processors: * 1
- Desired processors: * 2
- Maximum processors: * 4

The 'Next >' button is visible at the bottom of the dialog.

Figure 4-28 Processor setting with dedicated processors

- Setting Partition Memory

Set the partition memory, as shown in Figure 4-29 on page 148.

The *minimum*, *desired*, and *maximum* settings are similar to their processor counterparts:

- **Minimum memory**

Represents the absolute memory that is required to make the partition active. If the amount of memory that is specified under *minimum* is not available on the managed server, then the partition cannot become active.

- **Desired memory**

Specifies the amount of memory beyond the minimum that can be allocated to the partition. If the minimum is set at 256 MB and the desired is set at 4 GB, then the partition in question can become active with anywhere between 256 MB and 4 GB.

– **Maximum memory**

Represents the absolute maximum amount of memory for this partition. This value can be a value greater than or equal to the number specified in *Desired memory*. If this value is set at the same amount as desired, then the partition is considered *capped*. If this number is equal to the total amount of memory in the server, this partition is considered *uncapped*.

After you made your memory selections, select **Next**.

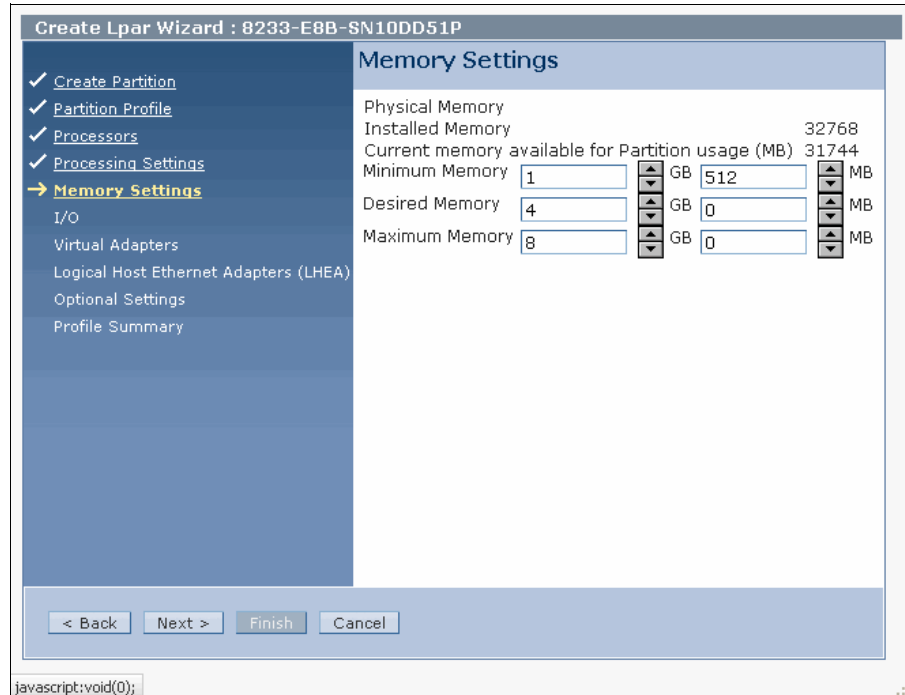


Figure 4-29 Partition memory setting

4. Configure Physical I/O

On the I/O window, as shown in Figure 4-30 on page 149, you select physical I/O resources for the partition to own. If you want the partition to own virtual resources, refer to item 5., “Configuring virtual resources” on page 149.

You can define the I/O resources as *Required* or *Desired*.

– **Required**

Represents the I/O resource that is required to make the partition active. Required I/O resource cannot be dynamically (DLPAR) removed from the partition.

– **Desired**

If during the partition startup the desired I/O resource is not assigned to any other running partitions, it is assigned to that partition. The desired I/O resources can be dynamically (DLPAR) removed from the partition.

Figure 4-30 shows the physical I/O setting of the partition.

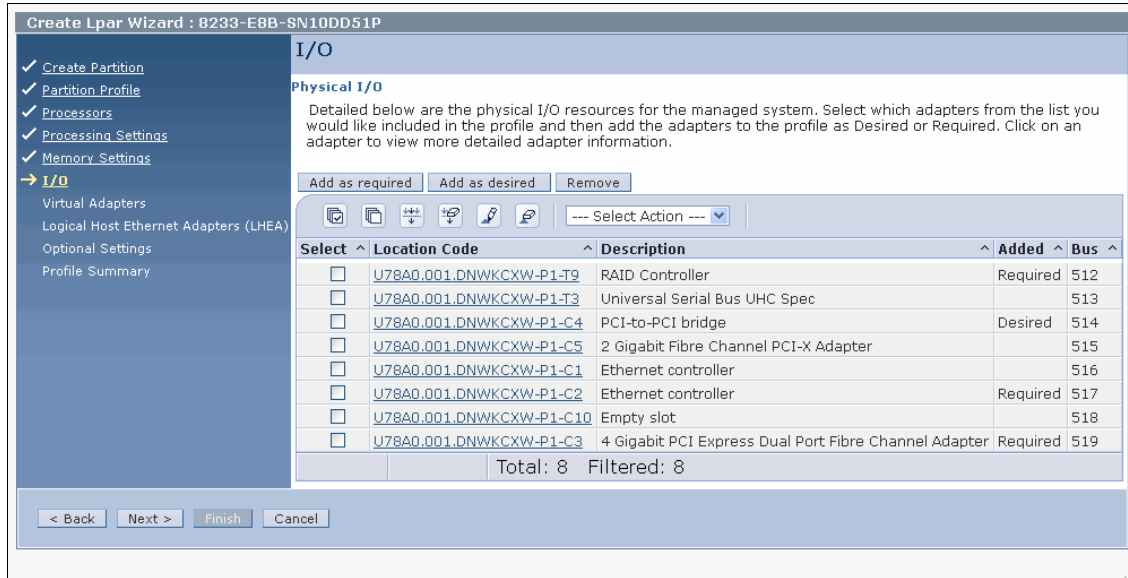


Figure 4-30 Partition physical I/O setting

5. Configuring virtual resources

If you have adapters that are assigned to the Virtual I/O Server (as explained in “Virtual Small Computer System Interface” on page 125), you can create a virtual adapter share for your partition. Take the following steps:

- a. Select **Actions** → **Create** → **Ethernet Adapter** to create a virtual Ethernet share. See Figure 4-31.

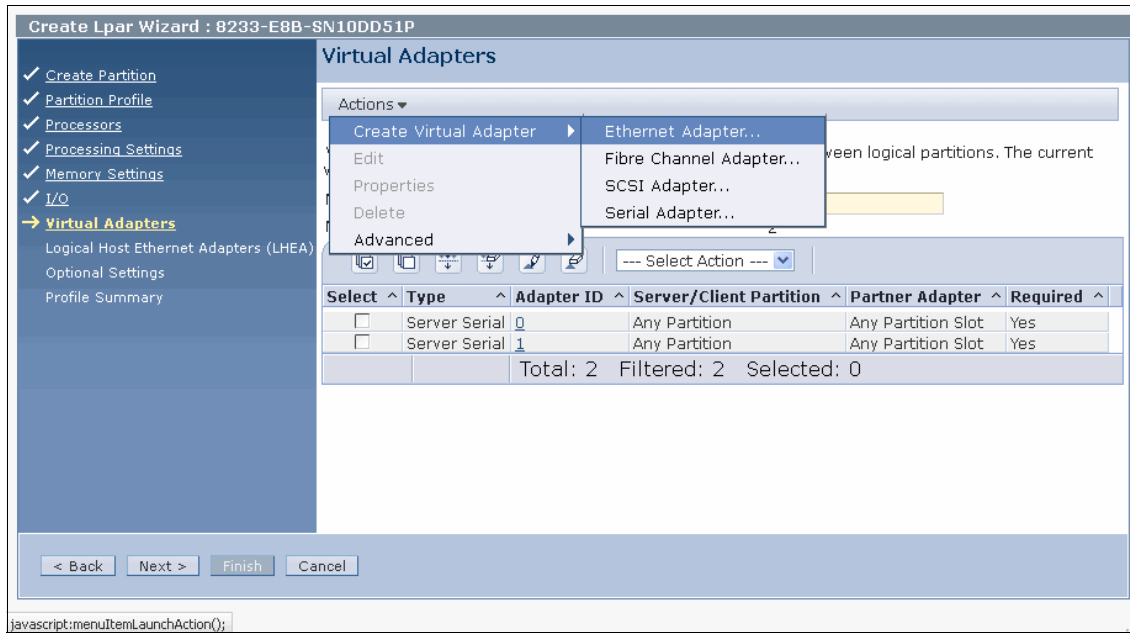


Figure 4-31 Configuring virtual resource

- If more than one VSwitch is defined in Hypervisor, select the appropriate Vswitch for this adapter.

- If required, add more VLAN (IEEE 802.1q compatible adapter) tags, as shown in Figure 4-32.




Figure 4-32 Create virtual Ethernet adapter

- b. Select **Actions** → **Create** → **SCSI** to create a virtual SCSI share. Alternatively, select **Actions** → **Create** → **Fibre Channel Adapter** to create a virtual Fibre Channel (NPIV) share.

You can specify your server partition, get system Virtual I/O Server information, and specify a tag for adapter identification, as shown in Figure 4-33. After you enter all of the data, select **OK**.

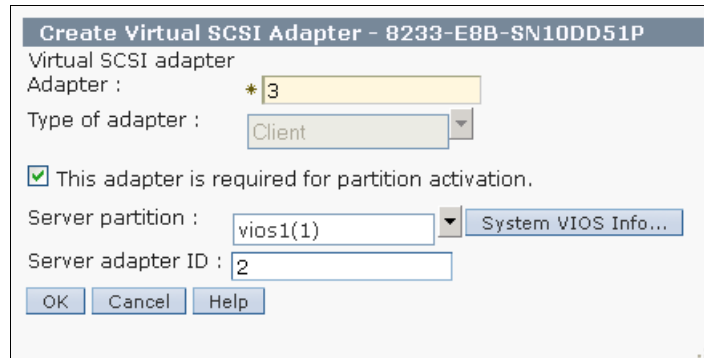


Figure 4-33 Create virtual SCSI adapter

You are returned to the virtual adapters window, as shown in Figure 4-33. When you are done creating all the virtual resources, select **Next**.

6. Configure **Host Ethernet Adapter**

If your server is POWER6 or POWER7 model 8231-E2B, 8202-E4B, 8205-E6B, 8233-E8B, 9117-MMB, or 9179-MHB, the server has a Host Ethernet Adapter, as shown in Figure 4-34. See 4.4.5, “Host Ethernet Adapter” on page 133 for configuration information.

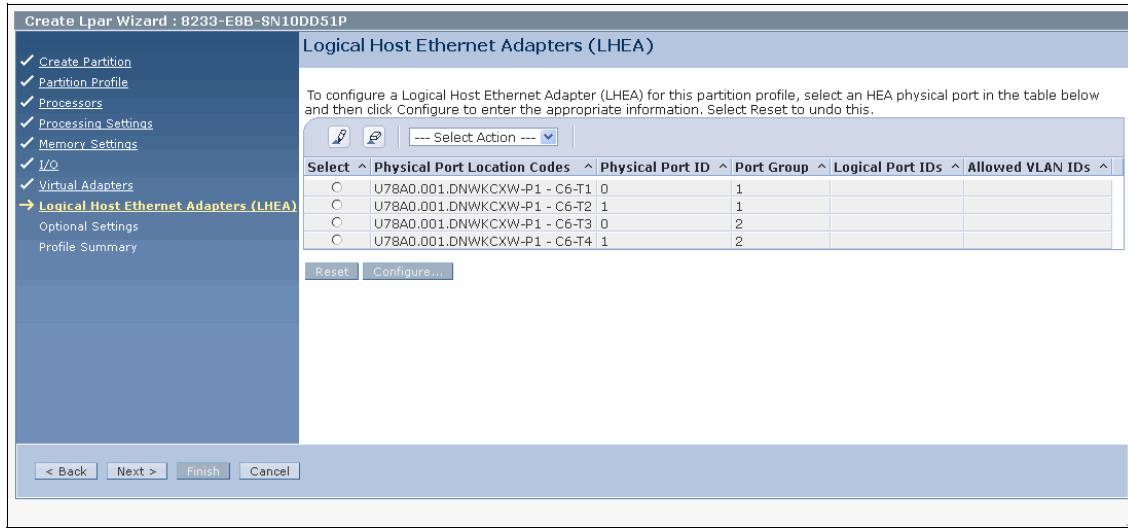


Figure 4-34 Logical Host Ethernet Adapters (LHEAs)

7. Optional Setting window

On the Optional Settings window that is shown in Figure 4-35, you can do the following functions:

- Enable connection monitoring.
- Start the partition with the managed system automatically.
- Enable redundant error path reporting.

You can also specify one of the various boot modes that are available.

After you make your selections in this window, click **Next** to continue.

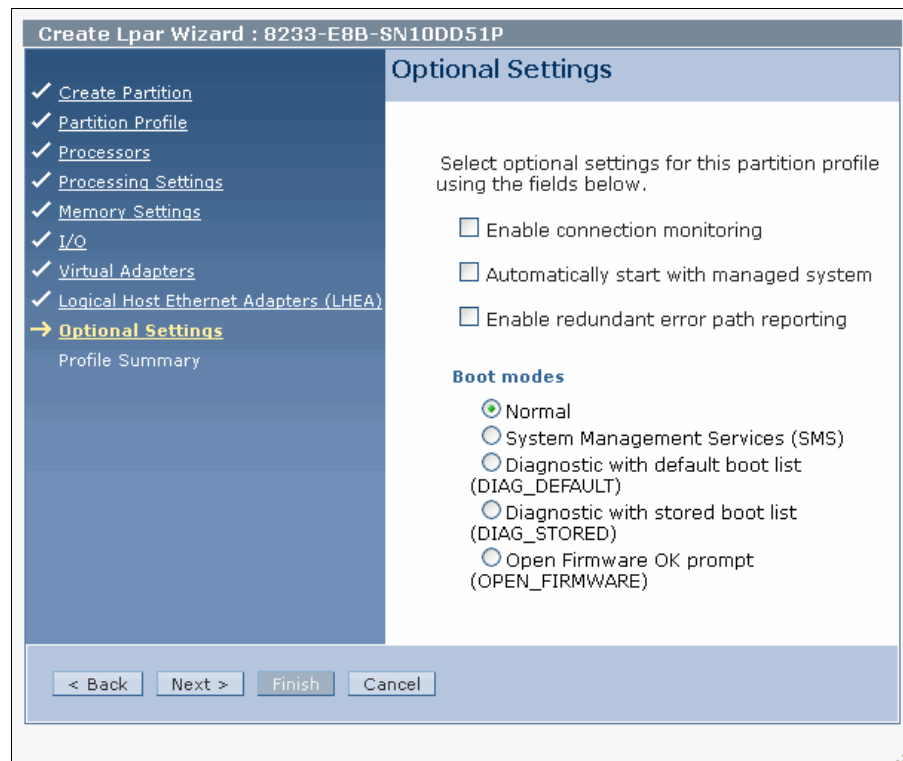


Figure 4-35 Optional settings

□ Enabling connection monitoring

Select this option to enable connection monitoring between the HMC and the logical partition that is associated with this partition profile. When connection monitoring is enabled, the Service Focal Point (SFP) application periodically tests the communications channel between this logical partition and the HMC. If the channel does not work, the SFP application generates a serviceable event in the SFP log. This step

ensures that the communications channel can carry service requests from the logical partition to the HMC when needed.

If this option is not selected, the SFP application still collects service request information when there are issues on the managed system. This option controls only whether the SFP application automatically tests the connection and generates a serviceable event if the channel does not work.

Clear this option if you do not want the SFP application to monitor the communications channel between the HMC and the logical partition that is associated with this partition profile.

□ Starting with managed system automatically

This option shows whether this partition profile sets the managed system to activate the logical partition that is associated with this partition profile automatically when you power on the managed system.

When you power on a managed system, the managed system is set to activate certain logical partitions automatically. After these logical partitions are activated, you must activate any remaining logical partitions manually. When you activate this partition profile, the partition profile overwrites the current setting for this logical partition with this setting.

If this option is selected, the partition profile sets the managed system to activate this logical partition automatically the next time the managed system is powered on.

If this option is not selected, the partition profile sets the managed system so that you must activate this logical partition manually the next time the managed system is powered on.

□ Enabling Redundant error path reporting

Select this option to enable the reporting of server common hardware errors from this logical partition to the HMC. The service processor is the primary path for reporting server common hardware errors to the HMC. Selecting this option allows you to set up redundant error reporting paths in addition to the error reporting path provided by the service processor.

Server common hardware errors include errors in processors, memory, power subsystems, the service processor, the system unit vital product data (VPD), nonvolatile random access memory (NVRAM), I/O unit bus transport (RIO and PCI), clustering hardware, and switch hardware. Server common hardware errors do not include errors in I/O processors (IOPs), I/O adapters (IOAs), or I/O device hardware.

If this option is selected, this logical partition reports server common hardware errors and partition hardware errors to the HMC. If this option is

not selected, this logical partition reports only partition hardware errors to the HMC.

This option is available only if the server firmware allows for the enabling of redundant error path reporting (the Redundant Error Path Reporting Capable option on the Capabilities tab in Managed System Properties is True).

– **Boot modes**

Select the default boot mode that is associated with this partition profile. When you activate this partition profile, the system uses this boot mode to start the operating system on the logical partition unless you specify otherwise when activating the partition profile. (The boot mode applies only to AIX, Linux, and Virtual I/O Server logical partitions. This area is unavailable for IBM i logical partitions.) Valid boot modes are as follows:

- **Normal**

The logical partition starts as normal. (This is the mode that you use to do most everyday tasks.)

- **System Management Services (SMS)**

The logical partition boots to the System Management Services (SMS) menu.

- **Diagnostic with default boot list (DIAG_DEFAULT)**

The logical partition boots using the default boot list that is stored in the system firmware. This mode is normally used to boot client diagnostics from the CD-ROM drive. Use this boot mode to run stand-alone diagnostics.

- **Diagnostic with stored boot list (DIAG_STORED)**

The logical partition performs a service mode boot using the service mode boot list that is saved in NVRAM. Use this boot mode to run online diagnostics.

- **Open Firmware OK prompt (OPEN_FIRMWARE)**

The logical partition boots to the open firmware prompt. This option is used by service personnel to obtain more debug information.

8. Profile Summary

When you arrive at the profile summary as shown in Figure 4-36, you can review your partition profile selections. If you see anything that you want to change, select **Back** to get to the appropriate window and to make changes. If you are satisfied with the data that is represented in the Profile Summary, select **Finish** to create your partition.

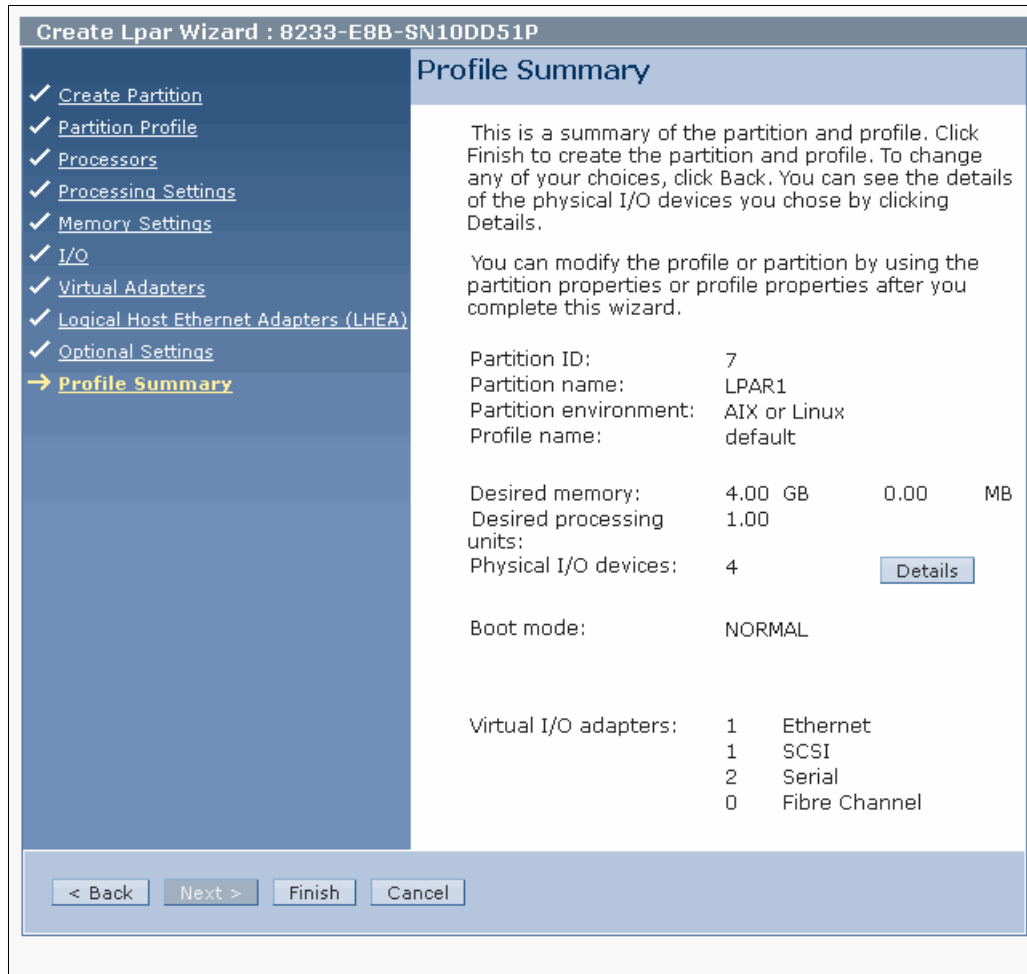


Figure 4-36 Profile Summary window

4.4.12 Create an IBM i logical partition

To create an IBM i partition, follow these steps:

1. Select **System Management** → **Servers**.
2. Select <Managed Server>.
3. From **Task**, select **Configuration** → **Create Logical Partition** → **IBM i**.
4. You can set **Partition ID** and specify **Partition name**, as shown in Figure 4-37. If you want the partition to be suspended, check the *Allow this partition to be suspended* box.

The *Restricted IO partition* check box in the IBM i LPAR creation is to determine if the IBM i LPAR is used as a live partition mobility. If the check box is selected, you can move the IBM i mobile partition. If the check box is cleared, you cannot move the IBM i mobile partition.

Then, click **Next**.

Create Lpar Wizard : 8233-E8B-SN10DD51P

→ Create Partition

Partition Profile
Processors
Processing Settings
Memory Settings
I/O
Virtual Adapters
Logical Host Ethernet Adapters (LHEA)
OptiConnect Settings
Tagged I/O
Optional Settings
Profile Summary

Create Partition

This wizard helps you create a new logical partition and a default profile for it. You can use the partition properties or profile properties to make changes after you complete this wizard.

To create a partition, complete the following information:

System name : 8233-E8B-SN10DD51P
Partition ID :
Partition name :

Allow this partition to be suspended.
 Restricted IO partition

< Back Next > Finish Cancel

Figure 4-37 Restricted IO Partition check box Option

5. Follow steps from “Configuring a shared processor partition” on page 142 through 6., “Configure Host Ethernet Adapter” on page 153.
6. Configure **OptiConnect Settings**.

The virtual OptiConnect feature provides high-speed interpartition communication within a managed system. The virtual OptiConnect feature emulates external OptiConnect hardware by providing a virtual bus between logical partitions. The virtual OptiConnect works only between IBM i LPAR, and you must install OptiConnect IBM i (a priced optional feature) on each IBM i logical partition that uses virtual OptiConnect. Ensure that the **Use virtual OptiConnect** option is selected on the **OptiConnect** tab.

Click **Next**.

7. Tagged I/O

A Tagged resource is an Input/Output Adapter (IOA) that is selected because it controls a device that performs a specific function for a logical partition. The HMC and IBM i operating system use this tagging to locate and use the correct I/O device for each I/O function.

Tagged I/O devices for this partition profile are defined as follows in Table 4-4.

Table 4-4 Devices associates with tagged IOAs

No	Device	Description	Required
1	Load source	Each IBM i logical partition must have one disk unit that is controlled by an IOA designated as the load source. The system uses the load source to start the logical partition.	Yes
2	Alternate restart device	This device can be a tape drive or an optical device. The media in the alternate restart device is what the system uses to start from when you perform a D-mode initial program load (IPL). The alternate restart device loads the License Internal Code that is contained on the source disk unit.	Yes
3	Console	The first workstation that the system activates in the logical partition and the only device it activates on a manual IPL. The logical partition is always available for use.	Yes (HMC is the default)
4	Alternate console	The alternate console of the first workstation that is stated in <i>Console</i> option.	No
5	Operation console	The option if the system uses the operation console to be the IBM i LPAR console. Leave it <i>None</i> if its console is controlled by the HMC.	No

Click **Next** and click **Finish** to complete creating an IBM i LPAR.

4.5 Capacity on Demand

This chapter describes the various types of Capacity on Demand (CoD), how to acquire enablement and activation codes, and finally how to enter these enablement and activation codes on your HMC to gain the benefits of the various CoD types.

4.5.1 Advantages of Capacity on Demand

CoD provides several advantages to IBM clients:

- ▶ Clients can plan for later expansion.

An IBM client can order a 64 way p770 system now with 16-way active and then can scale their system performance granularity up to a 64-way without human intervention or more hardware installation.

Similarly, clients can order a p780 or p795 with 4 TB of system memory and 1 TB active, and then can increase their memory capacity without extra hardware installation.

- ▶ Clients can work around budget constraints by taking advantage of CoD.

A client might want a 32-way p770 now but can afford only an 8-way within the current budget. In this scenario, the client can buy the 32-way with 8-way active. Then, when the budget allows, the client can activate the additional processors without having to order more hardware, schedule a CE, and so forth.

- ▶ Clients can plan for scaled usage or billing of POWER6 and POWER7 technology-based servers.

Clients can use utility Capacity on Demand or On/Off Capacity on Demand to have resources in reserve and can save money on servers by paying just for what they use.

- ▶ Clients can take advantage of increased reliability, availability, and serviceability (RAS).

Processor sparing allows for inactive processors to be activated immediately in the event of a processor failure. Processor sparing incurs no activation charge to the client.

Based on your current workload and foreseeable growth, the following chart (Figure 4-38) can help you to decide what CoD offering best fits your needs.

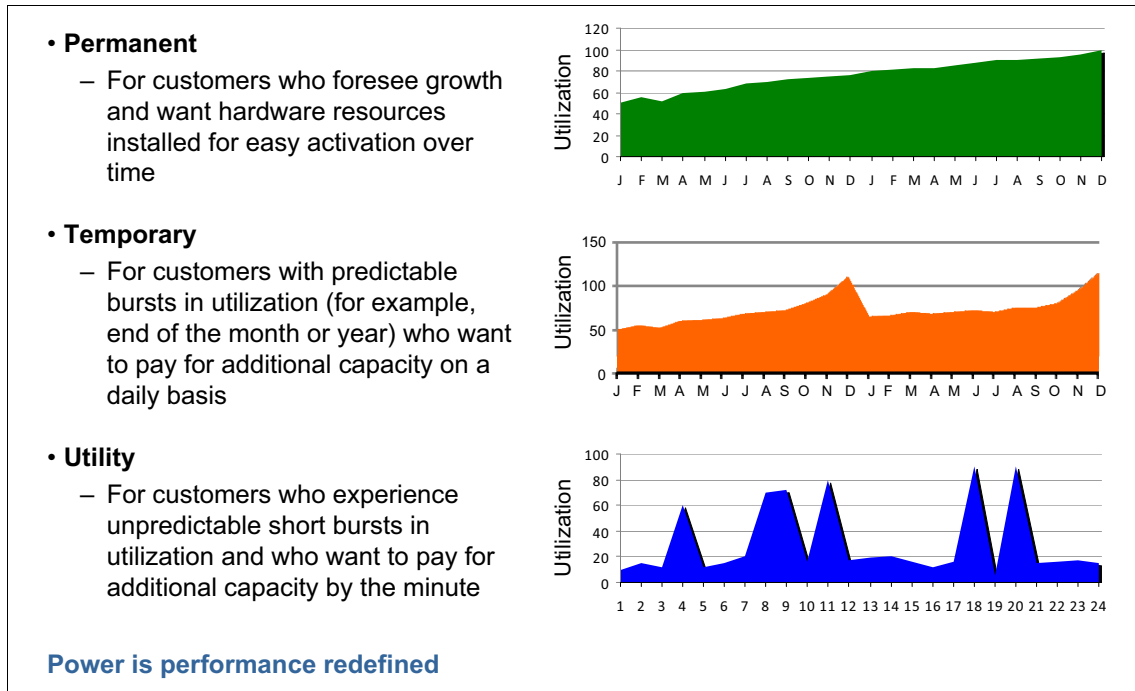


Figure 4-38 IBM Capacity on Demand offering

4.5.2 Permanent types of Capacity on Demand

Permanent types of CoD are permanent activations of inactive resources. On large scale Power Systems, resources (processors and sometimes memory) can be delivered inactive. For these inactive resources, permanent activations can be purchased either on initial order or through an upgrade. There are two types of permanent activation for these resources:

- ▶ Capacity Upgrade on Demand (CUoD)
- ▶ Mobile CoD

Capacity Upgrade on Demand

With Capacity Upgrade on Demand, bring new capacity online quickly and easily. Processors and memory can be activated dynamically without interrupting system or partition operations. Processors can be activated in increments of one processor. Memory can be activated in increments of 1 GB. As your workload demands require more processing power, you can activate inactive processors or

memory by placing an order for an activation feature. Over the Internet, you can retrieve an electronically encrypted activation code that unlocks the wanted amount of capacity. There is no hardware to ship and install, and no additional contract is required.

Mobile Capacity on Demand

Mobile CoD is the ability to move, at no charge, resource (processor *and* memory) activations between systems, and it is handled by the CoD Project Office.

Guidelines for Mobile Capacity on Demand

1. Process the Record Purpose Only (RPO) orders for the donor and target servers and send the completed RPO orders to the Capacity on Demand Project Office at pcod@us.ibm.com.

Be aware that there might be software changes and hardware service records also requiring updates for processor movement.

2. In reply to RPO orders, you must receive a deactivation code for the donor server from the Capacity on Demand Project Office. The deactivation code lowers the number of activated resources to align with the request for price quotation (RPQ) request.
3. Enter the deactivation code on the donor server. After you enter this code, send a listing of the updated VPD to the Capacity on Demand Project Office at pcod@us.ibm.com.

Collect the VPD by using the HMC command-line instruction:

```
Processor move: lscod -m <managed system> -t code -r proc -m  
mobile
```

```
Memory move: lscod -m <managed system> -t code -r mem-m mobile
```

4. With the receipt of the lscod profile, the Capacity on Demand Project Office provides an activation code for the target server.

Limitations:

- ▶ Movement must be between same machine type and models of servers (for example 770 to 770, but not 770 to 780).
- ▶ Movement must result in valid configuration for both the donor and target servers (the same hardware feature code that is removed must be added).
- ▶ Movement must be with the same country and the same enterprise.

4.5.3 Temporary types of Capacity on Demand

On some system types, some resources (processors and memory) can be temporarily activated for brief or extended periods of time. This method allows for greater resource allocation flexibility, and provide more choices to clients in how to pay for resource usage by allowing them to pay for just what they use.

There are several different types of temporary CoD:

- ▶ Trial CoD
- ▶ Utility CoD
- ▶ Capacity BackUp (CBU)
- ▶ On/Off CoD

Trial Capacity on Demand

Trial Capacity on Demand provides the flexibility to evaluate how more resources affect system workloads. A standard request is easily made for a set number of processor core activations and a set number of memory activations. The standard requests can be made after system installation and again after each purchase of permanent processor activation. POWER5 and POWER6 technology-based servers, except the POWER6 595, can activate up to two processor cores and up to 4 GB of memory. POWER7 technology-based servers and the POWER6 595 can activate up to eight processor cores and up to 64 GB of memory.

An exception request can be made one time over the life of the server and enables all available processor cores or memory.

Both standard and exception requests are available at no additional charge.

Utility Capacity on Demand

Utility CoD provides automated use of on-demand processors from the shared processor pool for short-term workloads on IBM POWER6 and POWER7 processor-based systems. Utility CoD is for clients with unpredictable, short workload spikes who need an automated and affordable way to help assure that adequate server performance is available as needed.

- ▶ Usage is measured in processor minute increments.
- ▶ Capacity can be paid for either before *or* after usage.
- ▶ Resource usage reporting is *required*.
- ▶ Capacity can be turned on and off by the client.
- ▶ Requires AIX 5.3 and higher, and PowerVM.

When you add Utility CoD processor cores, they are automatically placed in the default shared processor pool. These processor cores are available to any uncapped partition in any shared processor pool.

The processor cores become available to the resource manager of the pool. When the system recognizes that the combined processor utilization within the shared pool exceeds 100% of the level of base (purchased or active) processor cores that are assigned across uncapped partitions, a Utility CoD Processor Minute is charged. This level of performance is available for the next minute of use. If more workload requires a higher level of performance, the system automatically allows the additional Utility CoD processor cores to be used. The system automatically and continuously monitors and charges for the performance that is needed above the base (permanent) level.

When action on your part is required after you implement this CoD offering, the HMC displays messages on the HMC desktop.

For more information about Utility CoD, see this website:

<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7ha2/utilitycapacityondemandkick.htm>

Capacity BackUp

Capacity BackUp (CBU) uses On/Off CoD capabilities to provide an off-site, disaster-recovery server.

The CBU offering has a minimum set of active processor cores that can be used for any workload and many inactive processor cores that can be activated by using On/Off CoD in the event of a disaster. A specified number of no-charge On/Off CoD processor days is provided with CBU.

For more information about On/Off CoD, see “On/Off Capacity on Demand” on page 164.

On/Off Capacity on Demand

On/Off CoD allows you to temporarily activate and deactivate processor cores and memory units to help meet the demands of business peaks. After you request that a number of processor cores or memory units are to be made temporarily available for a specified number of days, those processor cores and memory units are available immediately. You can start and stop requests for On/Off CoD, and you are billed for usage at the end of each quarter.

You can change the number of resources and number of days in a running On/Off CoD request. Instead of having to stop the current request and start a new request, or wait until the current request expires, you can change the number of resources and number of days in the current request. For more

information about how billing works when changing a current request, see *Billing when changing a running On/Off Capacity on Demand request* at this website:

<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7ha2/onoffcodbillchange.htm>

For more information about how to change a current request, see *Changing a running On/Off Capacity on Demand request* at this website:

<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7ha2/onoffcodchangerequest.htm>

4.5.4 Capacity on Demand website navigation

The Power Systems Capacity on Demand website is at this link:

<http://www.ibm.com/systems/power/hardware/cod/index.html>

IBM Systems > Power Systems > Hardware >

Power Systems Capacity on Demand

Capacity on Demand solutions for Power Systems, System i and System p

Overview Offerings Activations Resources

Power Systems Capacity on Demand
Imagine you launch a dynamite new Web application for a peak season, and it's getting more traffic than you expected. Or say your business takes off and you need growth quickly but can't afford to take the system down for an upgrade.

With Capacity on Demand from IBM Power Systems, it's easy to activate dormant processor and memory resources within your system, without taking your system or application down. Whether your need is temporary or permanent, the solution is fast, it's easy and it's available today.

Financing your Capacity on Demand solutions

IBM Global Financing can help match your payments with your usage, with competitive financing for fixed and variable costs related to IBM Capacity on Demand offerings. By financing your Capacity on Demand costs and associated charges together with your base lease, spikes in demand need not become spikes in your budget.

Figure 4-39 Power Systems Capacity on Demand website

The Power Systems CoD website offers the following views:

- ▶ **Overview:** Capacity on Demand overview.
- ▶ **Offerings:** Description of CUoD, On/Off CoD, Utility CoD, Trial CoD, and Trial Active Memory Expansion.

- ▶ **Activations** - How to *request, enable, activate,* and *report* Permanent CoD, On/Off CoD, Utility CoD, Trial CoD, and Trial Active Memory Expansion. For example, see Chapter 4.5.5, “Entering CoD codes” on page 166.
- ▶ **Resources** - Additional CoD information for Power Systems and earlier System i and System p models.

4.5.5 Entering CoD codes

The following steps describe how to activate CUoD (permanent):

1. Order one or more processor cores and memory activation features.
2. After the order is fulfilled, get the code from this website:

<http://www-912.ibm.com/pod/pod/>

Enter your system type and serial number, as shown in Figure 4-40 and Figure 4-41 on page 167.

Capacity on Demand
Activation code

To search for an activation code for a specific system, enter the information below and click **Submit**.

Search for an activation code

System Type: 9117

Serial Number: 10 - B8FC4

Submit

Figure 4-40 Capacity on Demand activation code panel

Figure 4-41 shows the Capacity on Demand activation code output.

Capacity on Demand

Activation code

To search for an activation code for a specific system, enter the information below and click **Submit**.

Search results

System Type: 9117 Serial Number: 10-B8FC4

Type	Activation Code	Posted Date (MM/DD/YYYY)
MOD	5DC8D880EC3420DF568000000032004372	10/26/2012
MOD	1238260767D0052E56800000003200428E	07/25/2012
POD	DBD84196E631ECFA56720000008004189	06/28/2009
MOD	59D02E3583B887915680000006400417F	06/28/2009

Activation type definitions

POD: CUoD Processor Activation Code

MOD: CUoD Memory Activation Code

TCOD: On/Off CoD Enablement Code

TMOD: On/Off CoD Memory Enablement Code

VET: Virtualization Technology Code
(PowerVM, Enterprise Enablement, WWPN, Active Memory Expansion)

STDP: Standard Trial CoD Processor Activation Code

STDM: Standard Trial CoD Memory Activation Code

STME: Standard Trial Active Memory Expansion Code

EXCP: Exception Trial CoD Processor Activation Code

EXCM: Exception Trial CoD Memory Activation Code

USTA: Utility CoD Enablement Code

USTO: Utility CoD Termination Code

URPT: Utility CoD Reporting Code

PAID: Utility CoD PrePaid Code

Figure 4-41 Capacity on Demand activation code output

3. Enter the code through the HMC or Advanced System Management Interface (ASMI).
 - a. Through the HMC:
 - i. Select **System Management** → **Servers** → **<required managed_server>**.
 - ii. In the Tasks pane, select **Capacity on Demand (CoD)** → **Enter CoD Code**.

- iii. Enter the code for each resource, processor core, or memory, as shown in Figure 4-42.

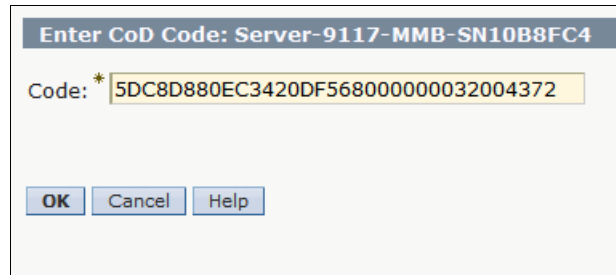


Figure 4-42 HMC Capacity on Demand enter code panel

- b. Through the ASMI:

Important: Shut down LPARs and the server to enter CoD codes through ASMI.

- i. From the HMC, select **System Management** → **Servers** → **<required managed_server>**.
- ii. In the Tasks pane, select **Operation** → **Launch Advanced System Management (ASM)**.
- iii. In the **Launch ASM Interface** panel, select **Service Processor IP Address** (primary, or secondary if redundant HMC exists).
- iv. Log in to ASMI and select **On Demand Utilities** → **CoD Activation**.

- v. Enter the code for each resource, processor core, or memory, as shown in Figure 4-43.

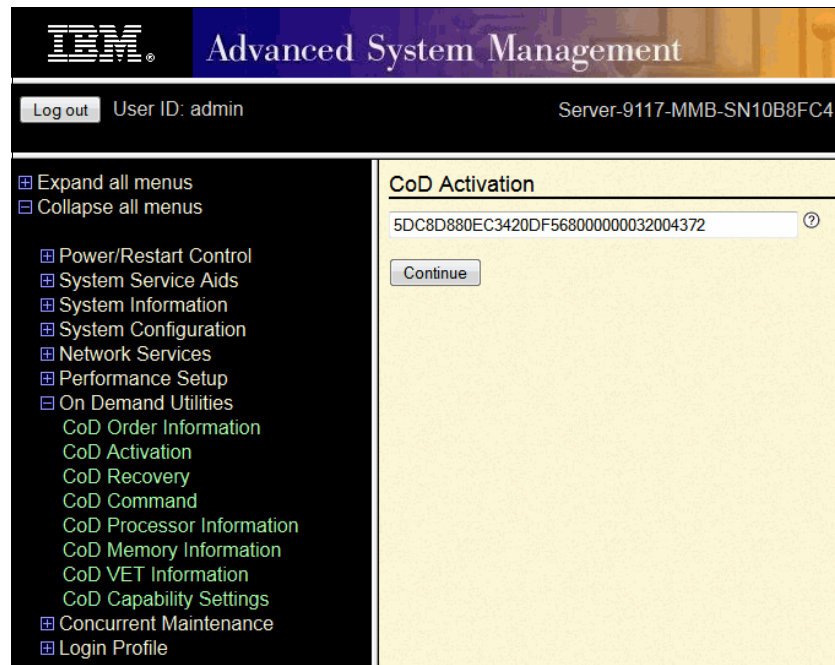


Figure 4-43 ASMI Capacity on Demand enter code panel

Accessing the ASMI without an HMC: To learn how to enter codes for a server that is not managed by an HMC, see *Accessing the ASMI without an HMC* at this website:

http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7hby/connect_asmi.htm

The *Accessing the ASMI without an HMC* website shows how to activate Capacity Upgrade on Demand (permanent) resources.

The Power Systems Capacity on Demand website provides instruction on how to do the following functions:

- ▶ Enable, use, report, and pay for *On/Off Capacity on Demand*
- ▶ Enable, use, report, and pay for *Utility Capacity on Demand*
- ▶ Request *Trial Capacity on Demand*
- ▶ Request *Trial Active Memory Expansion*

See the Power Systems Capacity on Demand website:

<http://www-03.ibm.com/systems/power/hardware/cod/activations.html>

4.5.6 Stopping Trial CoD

Trial Capacity on Demand ends when the trial period is over and the resources are reclaimed by the server. Return the resources before the trial period ends.

You can use the HMC to stop (see “Steps to stop CoD:” on page 170 that follow) a current Capacity on Demand trial for processors or memory units before the trial automatically expires. If you choose to stop the trial before it expires, you cannot restart it and you forfeit any remaining days.

If your server is powered off or loses power before the resources are removed from the logical partitions, you might need to do recovery actions to successfully power on your server. For more information, see *Recovery actions*:

<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/ipha2/tcodend.htm>

Steps to stop CoD:

1. Through the HMC, select **System Management** → **Servers** → **<required managed_server>**.
2. In the **Tasks Pane**, select **Capacity on Demand (CoD)** → **Processor** → **Trial CoD**.

3. Click **Stop** (as shown in Figure 4-44) to return resources to the pool.

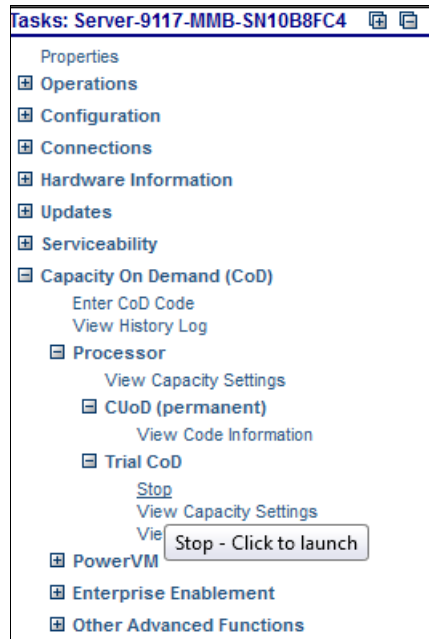


Figure 4-44 Stopping Trial CoD panel



Operating

This chapter describes a general overview of the Hardware Management Console (HMC) operation, including resource (server, partition, frame) management, HMC management, remote access logical partitions from HMC, and partition data management.

5.1 Basic operation

This section describes how to operate resources such as servers, partitions, and frames, which are managed by the HMC.

5.1.1 Using the web-based user interface

This section describes how to use the web-based user interface to do tasks on your managed resources.

HMC Version 7 is migrated to a new framework, the IBM System z® HMC framework. All existing management functions and commands remain unchanged. However, there are new user interface improvements and changes because of the new framework.

A major change is the web-based user interface. With this interface, you do not have to install an application to access the HMC remotely, and you can connect to the HMC by using your browser. Firefox and Internet Explorer are supported.

Starting the HMC

First, start the HMC by setting both the display and system units to the on position. You then see the initialization window that includes the IBM logo and copyright information.

After you finish the initialization step, the Welcome window displays, as shown in Figure 5-1. This page includes the link to log on to view the online help, and the summarized HMC status information.

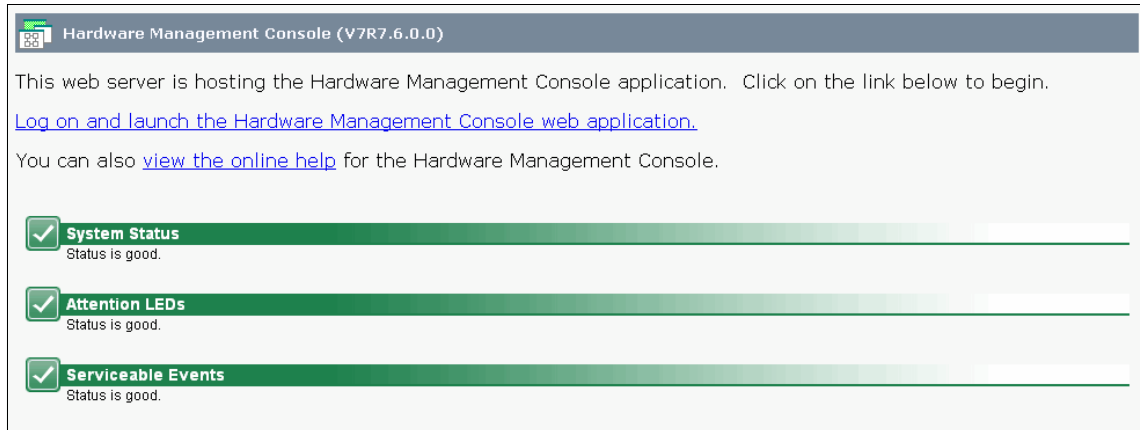


Figure 5-1 HMC Welcome window

To log on to the HMC, click **Log on and launch the Hardware Management Console web application** from the Welcome window. The Logon window opens, as shown in Figure 5-2.

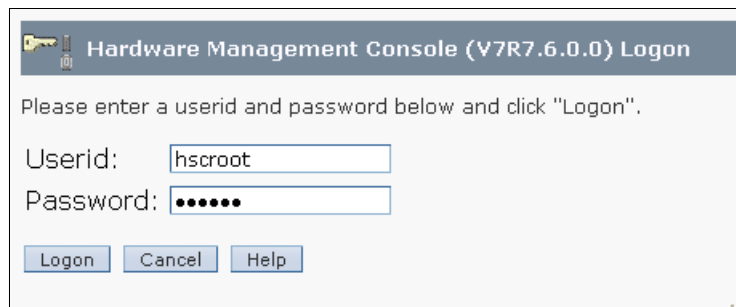


Figure 5-2 HMC Log on window

The HMC is supplied with a predefined user ID, *hscroot*, and the default password *abc123*. When you update your password, you can no longer keep it six characters. The minimum length for a password is now seven characters.

User ID and password are case-sensitive: Both the user ID and password are case-sensitive and must be entered exactly.

Session preservation

With HMC Version 7, you can remain in the graphical user interface (GUI) session across logins, as shown in Figure 5-3. If you want to preserve your session, choose **Disconnect** and click **OK**.

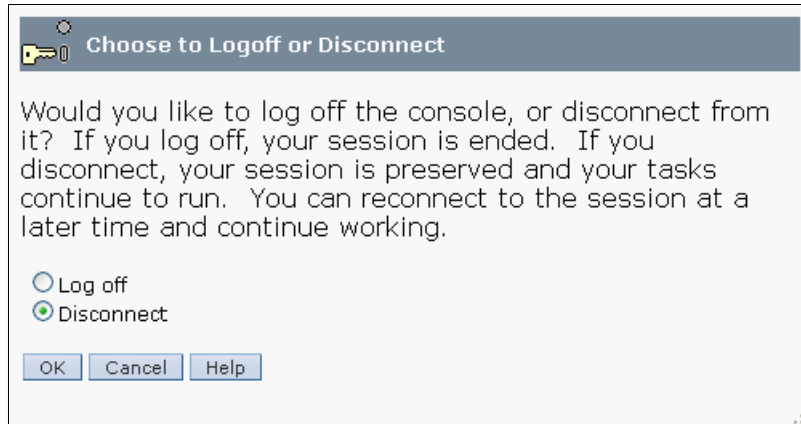


Figure 5-3 HMC logoff or disconnect window

After disconnecting from the session, you can reconnect to the session by selecting the session that you want to connect. As shown in Figure 5-4, session ID 28 has two running jobs. When you reconnect that session, the jobs that you were doing previously are displayed. You also see that there are three disconnected sessions for the user ID *hscroot*. This is a typical situation when all users log in with the same user ID (for example, *hsroot*). The disconnect feature provides another reason to use separate user IDs for each user.

Choose a Disconnected Session

The following disconnected sessions are available to user "hscroot". You can choose to either reconnect to one of these sessions, or start a new session. To reconnect, select the session to which you wish to reconnect, then click **Reconnect**. To create a new session, click **New Session**.

You can also delete a disconnected session by selecting the session you wish to delete, and then clicking **Delete**.

If you'd rather cancel connecting, click **Cancel**.

Select	Session Id	Disconnect Time	Creation Time	Running Tasks
<input checked="" type="radio"/>	28	Oct 25, 2012 8:34:50 AM	Oct 25, 2012 8:15:11 AM	2
<input type="radio"/>	24	Oct 24, 2012 5:40:19 PM	Oct 24, 2012 3:55:59 PM	1
<input type="radio"/>	14	Oct 24, 2012 5:36:36 PM	Oct 24, 2012 8:46:27 AM	0

Figure 5-4 Reconnecting the previous session

Components of the web-based user interface

The HMC workplace window consists of several major components, as shown in Figure 5-5.

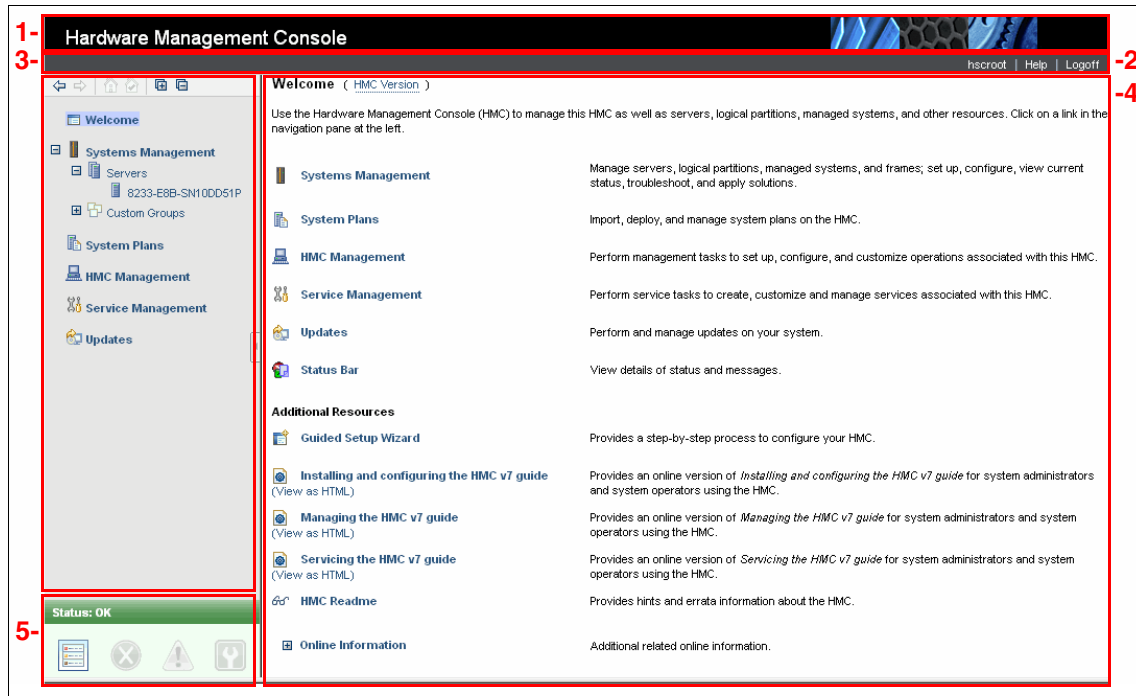


Figure 5-5 HMC workplace window

Table 5-1 also shows HMC workplace window components.

Table 5-1 Components of the HMC workplace window

No.	Component
1	Banner
2	Task bar
3	Navigation pane
4	Work pane
5	Status bar

Banner

The *banner*, across the top of the workplace window, identifies the product and logo. It is optionally displayed and is set by using the **Change User Interface Setting** task.

Task bar

The *task bar* is located below the banner. It displays the names of any tasks that are running, the user ID you are logged in as, online help information, and the ability to log off or disconnect from the console. The task bar provides the capability of an active task switcher. You can move between tasks that were started and are not yet closed. However, the task switcher does not pause or resume existing tasks. For example, when you run three tasks on the HMC, you can see tasks name in the Task bar and click to switch them, as shown in Figure 5-6.



Figure 5-6 Active tasks in the Task bar

Navigation pane

The *navigation pane*, in the left portion of the window, contains the primary navigation links for managing your system resources and the HMC. The following links can be found on the navigation pane:

- ▶ Welcome
- ▶ Systems Management
- ▶ System Plans
- ▶ HMC Management
- ▶ Service Management
- ▶ Updates

Work pane

The *work pane*, in the right portion of the window, displays information that is based on the current selection from the navigation pane. For example, when you select **Welcome** in the navigation pane, the welcome window content displays in the work pane, as shown in Figure 5-5 on page 178.

Status bar

The *status bar*, in the lower left portion of the window, provides visual indicators of current overall system status. It also includes a status overview icon that can be selected to display more detailed status information in the work pane.

5.1.2 Systems Management: Servers

Systems Management displays tasks to manage servers, logical partitions, and frames. Use these tasks to set up, configure, view status, troubleshoot, and apply solutions for servers.

This section describes the tasks to manage a server.

Servers

The *servers* node represents the servers that are managed by this HMC. To add servers, take the following steps:

Before you begin: You must assign an IP address or host name to the service processor on the managed system. We explain the task of connecting managed systems in 3.1.2, “HMC Guided Setup wizard” on page 64.

1. Select **Systems Management** → **Servers** in the navigation pane.
2. Click **Connections** → **Add Managed Systems** in the work pane.
3. Select **Add a managed system** and enter an IP address or host name, then click **OK**, as shown in Figure 5-7.

Add Managed Systems

Use this panel to add systems in the network to the systems managed by this HMC.

If you know the name or IP address of the system you want to add, enter its specific name or IP address and click Ok.

If you want to find the IP addresses of systems in the network, you can specify a range of IP addresses and click Ok to view the list of IP addresses with their system names that were discovered in the network. You can then select one or more systems from the list to add to the managed systems of this HMC. The discovery process will take a long time.

Add a managed system
IP Address/Host name: *
Password:

Find managed systems
Enter a range of IP addresses to search for managed systems.
Beginning IP Address: *
Ending IP Address: *

Figure 5-7 Add Managed Systems window

4. Click **Servers** to see a listing of individually defined servers in table form in the work pane, as shown in Figure 5-8.
5. Select the server that you want to add to the HMC.

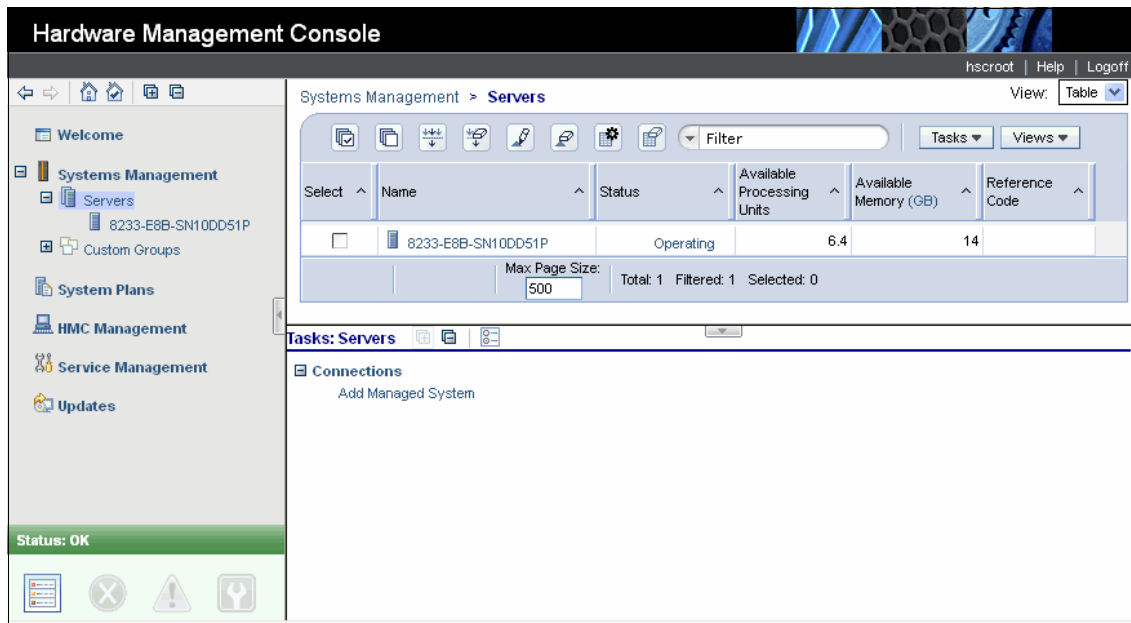


Figure 5-8 System Management servers window

By default, the contents of servers displays the following attributes:

- ▶ **Name**
Specifies the user-defined name of the managed system.
- ▶ **Status**
Displays the status of the managed system (for example, *operating*, *power off*, *initializing*). In addition, displays icons that represent an unacceptable state and active attention LED.
- ▶ **Available Processing Units**
Displays the number of processing units that are available for assignment to logical partitions on the managed system. This number is the total number of processing units that are activated on the managed system minus the number of processing units that are assigned to the logical partitions. This number includes the logical partitions that are shut down on the managed system. This number does not include any processing units that are not yet activated with CoD.

▶ Available Memory

Displays the amount of memory that is available for assignment to logical partitions on the managed system. This amount is the total amount of memory that is activated on the managed system minus the amount of memory that is needed by managed system firmware minus the amount of memory that is assigned to the logical partitions. This amount includes the logical partitions that are shut down on the managed system. This number does not include any memory that is not yet activated with CoD. The available memory amount can be shown in MB or GB. Click **MB** or **GB** in the Available Memory column title.

▶ Reference Code

Displays the progress system reference code (SRC). By clicking the displayed SRC, you can receive more information.

The table can also display the following optional attributes:

- ▶ Name
- ▶ Status
- ▶ Available Processing Units
- ▶ Available Memory
- ▶ Reference Code
- ▶ Configurable Processing Units
- ▶ Configurable Memory
- ▶ Serial Number
- ▶ Type-Model
- ▶ CoD Processor Capable
- ▶ CoD Memory Capable
- ▶ Permanent Processors
- ▶ On/Off CoD Processor State
- ▶ Trial CoD Processor State
- ▶ Reserve CoD Processor State
- ▶ Utility CoD Processor State
- ▶ Permanent Memory
- ▶ On/Off CoD Memory State
- ▶ Trial CoD Memory State
- ▶ Current Power Server Mode
- ▶ Desired Power Server Mode

These attributes display when you select the *Column configuration* icon on the table toolbar, as shown in Figure 5-9. This function allows you to select extra attributes that you want displayed as columns in the table.

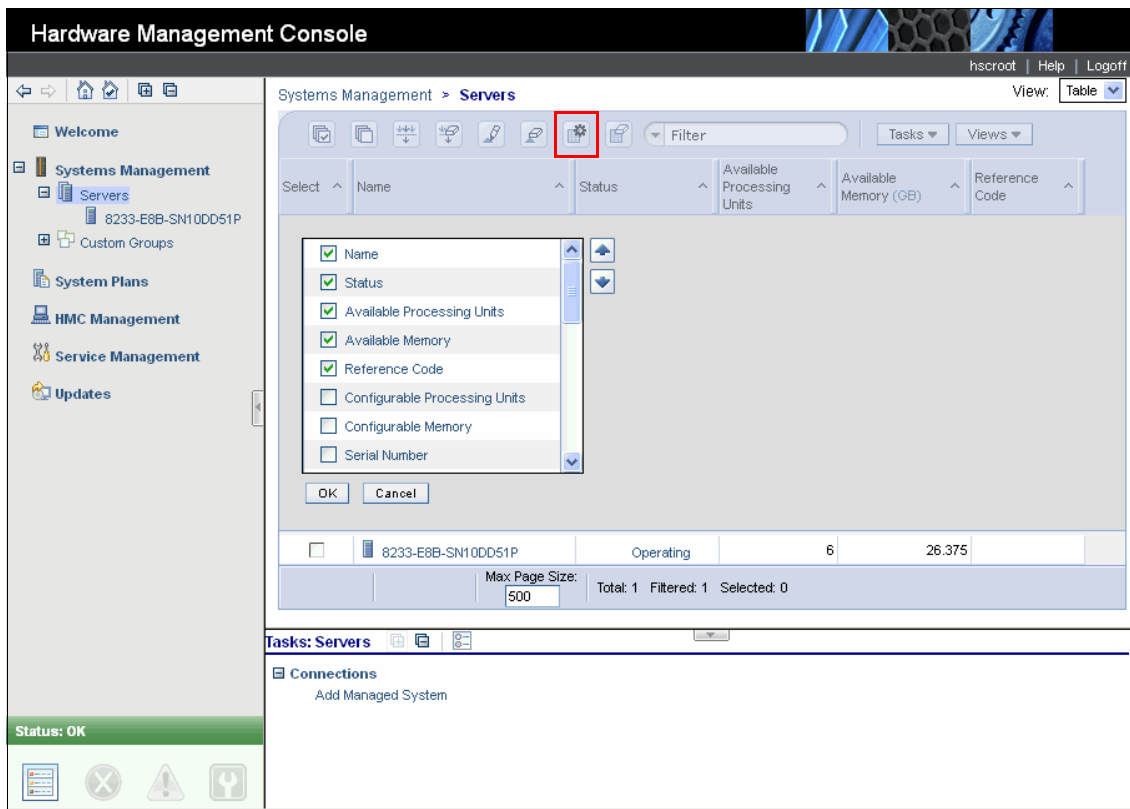


Figure 5-9 Column configuration

You can also use the **Views** menu from the table toolbar to change the view between default server attributes and the CoD server attributes, as shown in Figure 5-10.




Figure 5-10 Views option

If you select the **Manage Views** option from the **Views** menu, this option also allows you to create new custom view, as shown in Figure 5-11.

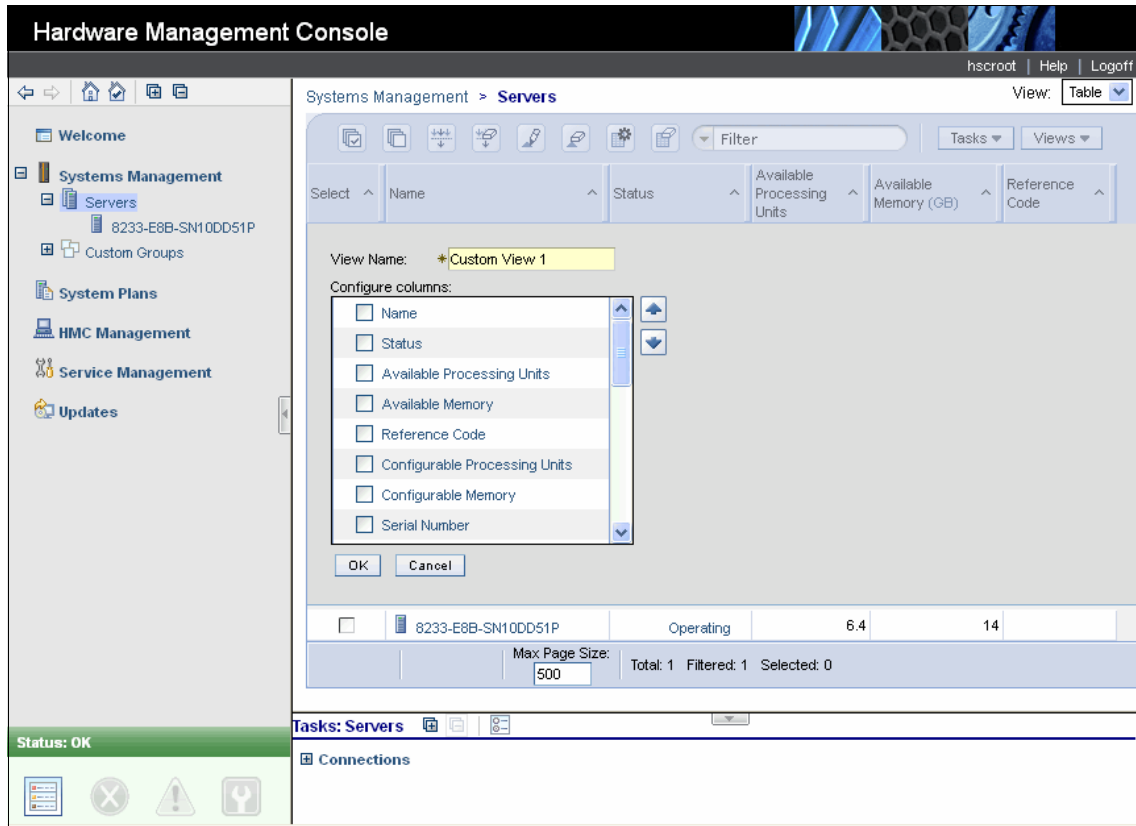


Figure 5-11 Create new custom view

Properties

Properties include the tasks to display the selected properties of the managed system, as shown in Figure 5-12.

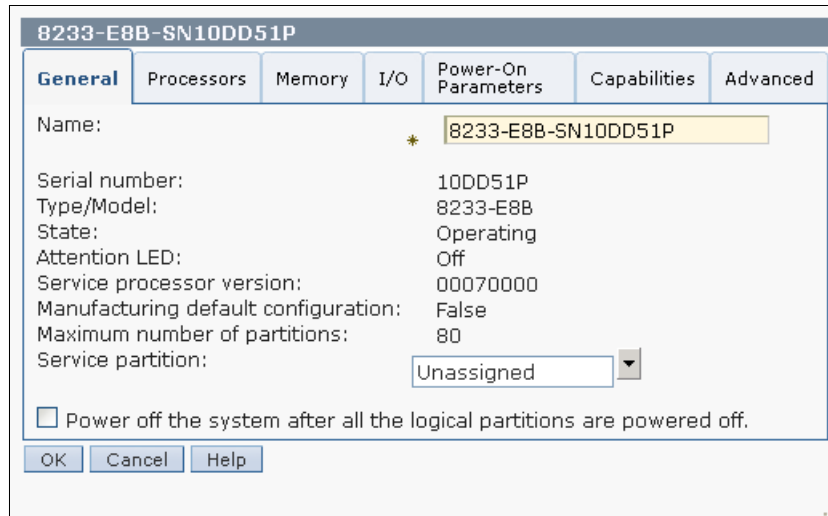


Figure 5-12 Properties tasks of the server

The following list defines the properties:

- ▶ **General**

The *General* tab displays the system's name, serial number, model and type, state, attention LED state, service processor version, maximum number of partitions, assigned service partition (if designated), and power off policy information.

- ▶ **Processors**

The *Processor* tab displays information about the managed system's processors including installed processing units, unconfigured processing units, available processing units, configurable processing units, minimum number of processing units per virtual processor and maximum number of shared processor pools.

- ▶ **Memory**

The *Memory* tab displays information about the managed system's memory including installed memory, unconfigured memory, available memory, configurable memory, memory region size, current memory available for partition usage, and system firmware current memory. A tab describes the maximum number of memory pools.

- ▶ I/O

The *I/O* tab displays the physical I/O resources for the managed system. The assignment of I/O slots and partition and adaptor-type information are displayed, grouped by units. Select the link in the Slot column to display the physical I/O properties of each resource. Select I/O Pools to display all of the I/O pools found in the system and the partitions that are participating in the pools.

- ▶ Migration

The *Migration* tab displays partition migration information.

Migration tab: If your managed system is not partition-migration capable, the **Migration** tab is not shown.

- ▶ Power-On Parameters

The *Power-On Parameters* tab displays the initial program load (IPL) source mode for restarting, and allows you to change the power-on parameters for the next restart by changing the values in the Next fields. These changes are only valid for the next managed system restart.

- ▶ Capabilities

The *Capabilities* tab displays the runtime capabilities of this server. Select Help for more information about the capabilities listed.

- ▶ Advanced

The *Advanced* tab displays huge page memory capabilities on the managed system, including available huge page memory, configurable huge page memory, current page size, and current maximum huge page memory. To change memory allocation on systems with huge page table support, set the Requested huge page memory (in pages) field to the wanted memory. To change the requested value for huge page memory, the system must be powered off. The Barrier Synchronization Register (BSR) displays array information.

Operations

Operations includes the tasks for server operations. These following list provides operations tasks:

- ▶ Power on
- ▶ Power off
- ▶ Power management
- ▶ LED status
- ▶ Schedule operations
- ▶ Launch Advanced System Management (ASM)

- ▶ Utilization data
- ▶ Rebuild
- ▶ Change password

These tasks are now described.

Power On task

Use this task to power on the managed system. You can choose from the following options (see Figure 5-13):

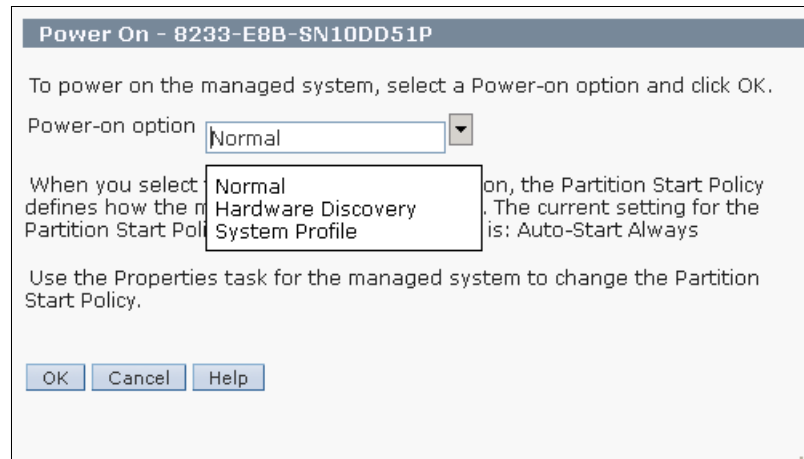


Figure 5-13 Power On task

▶ **Normal**

Turns on the managed system with the current setting for the partition start policy to determine how to power on the managed system. You can change the partition start policy from the **Power On Parameters** tab of the **Properties** task for the managed system. The current setting can be one of the following values:

- Auto-Start Always
- Stop at Partition Standby
- Auto-Start for Auto-Recovery
- User-Initiated

▶ **Hardware Discovery**

Turns on the managed system into a special mode, which performs the hardware discovery. After the Hardware Discovery process is complete, the system is in Operating state with any partitions in the power-off state. The Hardware Discovery process records the hardware inventory in a cache on the managed system. The collected information is then available for use when

displaying data for I/O devices or when creating a system plan based on the managed system. This option is available only if the system can use the hardware discovery process to capture I/O hardware inventory for the managed system.

► **System Profile**

Turns on the managed system and its logical partitions based on a predefined system profile. When you select this option, you must select the partition profile that you want the HMC to use to activate logical partitions on the managed system.

System profiles: If the HMC does not have any system profiles, the **System Profile** option is not shown. System profiles are explained in “Manage System Profile task” on page 197.

Power Off task

Use this task to shut down the managed system. Turning off the managed system makes all partitions unavailable until the system is turned on again.

Before you turn off the managed system, ensure that all logical partitions are shut down and that their states have changed from *Running* to *Not Activated*.

If you do not shut down all logical partitions on the managed system before you turn off the managed system, the managed system shuts down each logical partition before the managed system itself turns off. This process can cause a substantial delay in turning off the managed system, particularly if the logical partitions are not responsive. Further, the logical partitions might shut down abnormally, which can result in data loss and further delays when you activate the logical partitions again.

You can choose from the following options (see Figure 5-14):

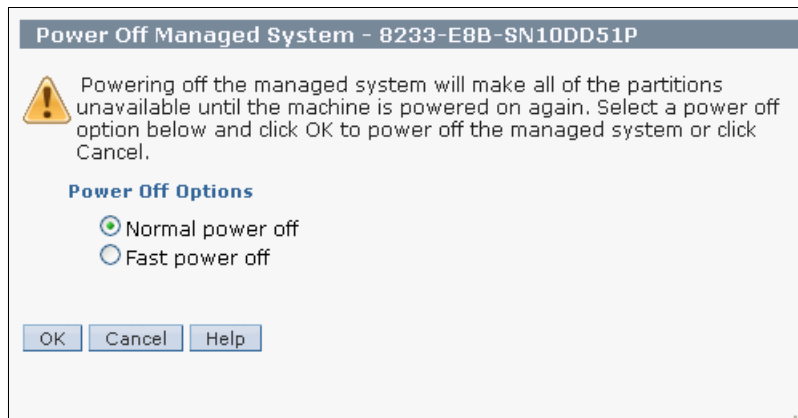


Figure 5-14 Power Off task

► **Normal power off**

The system ends all active jobs in a controlled manner. During that time, programs running in those jobs are allowed to do cleanup (end-of-job processing).

► **Fast power off**

The system ends all active jobs immediately. The programs running in those jobs are not allowed to do any cleanup. Some applications, such as web servers that are providing information, might not have a problem with fast power off. Other applications, such as databases with cached information, might lose data if the application cannot do cleanup before the application ends.

Power Management task

Use this task to change the power saver mode of the managed system. You can reduce power consumption by enabling the power saver mode on. You can check the current power saver mode and change it, as shown in Figure 5-15.

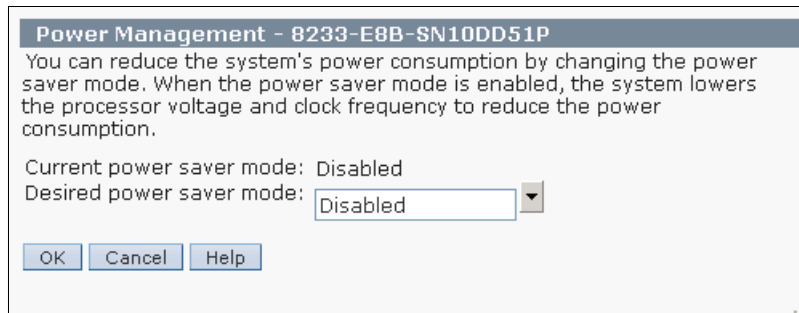


Figure 5-15 Power Management task

Power saver mode: Enabling the power saver mode on a managed system might affect the accuracy of any performance monitoring tools that are running on the managed system.

LED Status task

Use this task to view system attention LED information, light specific LEDs to identify a system component, and test all LEDs on a managed system.

You can choose from the following options:

► **Deactivated Attention LED**

Deactivates all system attention LEDs and logical partition LEDs.

► Identify LED

Displays the current Identify LED states for all the location codes that are contained in the selected enclosure, as shown in Figure 5-16. From this task, you can select a single location code or multiple location codes to operate against and activate or deactivate the LEDs by selecting the corresponding box.

Identify LED, Select Location - 8233-E8B-SN10DD51P

The current Identify LED states for all the location codes contained in the selected enclosure are displayed below. Select a single location code or multiple location codes to operate against and activate or deactivate the LED(s) by selecting the corresponding button.

Selected System: 8233-E8B*10DD51P
Selected Enclosure: System Unit, Model E8B, 78A0-001/DNWKCXW

Select ^	Location ^	Description ^	Identify LED State ^
<input type="checkbox"/>	U78A0.001.DNWKCXW-A1	Air Moving Device	Off
<input type="checkbox"/>	U78A0.001.DNWKCXW-A2	Air Moving Device	Off
<input type="checkbox"/>	U78A0.001.DNWKCXW-A3	Air Moving Device	Off
<input type="checkbox"/>	U78A0.001.DNWKCXW-A4	Air Moving Device	Off
<input type="checkbox"/>	U78A0.001.DNWKCXW-D1	Operator Panel	no LED present
<input type="checkbox"/>	U78A0.001.DNWKCXW-D1-T1	USB Port	no LED present
<input type="checkbox"/>	U78A0.001.DNWKCXW-E1	Power Supply	Off
<input type="checkbox"/>	U78A0.001.DNWKCXW-E1-T1	unknown	no LED present
<input type="checkbox"/>	U78A0.001.DNWKCXW-E2	Power Supply	Off
<input type="checkbox"/>	U78A0.001.DNWKCXW-E2-T1	unknown	no LED present
<input type="checkbox"/>	U78A0.001.DNWKCXW-P1	System Backplane	Off
<input type="checkbox"/>	U78A0.001.DNWKCXW-P1-C1	PCI Adapter Card	Off
<input type="checkbox"/>	U78A0.001.DNWKCXW-P1-C10	RAID Enablement Card	no LED present
<input type="checkbox"/>	U78A0.001.DNWKCXW-P1-C11	RAID Enablement Card	no LED present
<input type="checkbox"/>	U78A0.001.DNWKCXW-P1-C12	TPMD Card	Off
<input type="checkbox"/>	U78A0.001.DNWKCXW-P1-C13	System Processor Assembly	Off
<input type="checkbox"/>	U78A0.001.DNWKCXW-P1-C13-C1	Voltage Regulator	Off
<input type="checkbox"/>	U78A0.001.DNWKCXW-P1-C13-C10	Voltage Regulator	Off
<input type="checkbox"/>	U78A0.001.DNWKCXW-P1-C13-C2	Memory DIMM	Off
<input type="checkbox"/>	U78A0.001.DNWKCXW-P1-C13-C3	Memory DIMM	Off

Activate LED Deactivate LED Refresh Cancel Help

Figure 5-16 Identify LED task

► Test LED

Initiates an LED Lamp Test against the selected system. All LEDs activate for several minutes.

Schedule Operations task

Use this task to create a schedule for certain operations to be performed on the managed system without operator assistance. 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. For example, you can schedule power on or off operations for a managed system.

You can choose following tasks:

- ▶ Activate on a System Profile
- ▶ Backup Profile Data
- ▶ Power Off Managed System
- ▶ Power On Managed System
- ▶ Manage Utility CoD processors
- ▶ Manage Utility CoD processor minute usage limit
- ▶ Modify a Shared Processor Pool
- ▶ Move a partition to a different pool
- ▶ Change the power saver mode on a managed system

From the **Set up a Scheduled Operation** window, you can set up the task, as shown in Figure 5-17.

The screenshot shows a dialog box titled "Set up a Scheduled Operation - 8233-E8B-SN10DD51P". It has two tabs: "Date and Time" (selected) and "Repeat". The main text reads: "The scheduled operations will be created for the selected objects listed below: 8233-E8B-SN10DD51P. Select the date and time, and select a time window. The scheduled operation will start at the specified date and time unless an existing condition prevents its execution, a resource constraint, for example. In this case, an attempt will be made to start the scheduled operation within the time window starting at the specified date and time." Below this, there are two input fields: "Date *" with the value "11/3/12" and "Time *" with the value "10:00 AM". To the right, under the heading "Time Window", there are six radio button options: "10 minutes" (selected), "20 minutes", "30 minutes", "40 minutes", "50 minutes", and "60 minutes". At the bottom, there are three buttons: "Save", "Cancel", and "Help".

Figure 5-17 Set up a Scheduled Operation task: Date and Time

Figure 5-18 shows how to set up a scheduled operation to repeat.

The screenshot shows the same dialog box as Figure 5-17, but with the "Repeat" tab selected. The main text is the same. Below it, under the heading "Single or Repeated", there are two radio button options: "Set up a single scheduled operation" (selected) and "Set up a repeated scheduled operation". Below this, there are two sections: "Days of the Week" and "Options". "Days of the Week" has checkboxes for Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, and Sunday. "Options" has an "Interval" spinner set to "1" with a range of "1 to 26 weeks", a "Repetitions" spinner set to "1" with a range of "1 to 100", and a checkbox for "Repeat indefinitely" which is currently unchecked. At the bottom, there are three buttons: "Save", "Cancel", and "Help".

Figure 5-18 Set up a Scheduled Operation task: Repeat

You can also display defined tasks in the **Customize Scheduled Operations** window, as shown in Figure 5-19.

Select	Target	Date	Time	Operation	Remaining Repetitions
<input type="checkbox"/>	8233-E8B-SN10DD51P	5/2/13	10:00 AM	Backup Profile Data	1

Figure 5-19 Customize Scheduled Operations task

Launch Advanced System Management

If configured to do so, the HMC connects directly to the Advanced System Management (ASM) interface for a selected system from this task. We explain this task in 6.3, “Advanced System Management Interface” on page 341.

Utilization data

Utilization data is records that include information about the memory and processor usage on a managed system at a particular time. You can set the HMC to collect usage data in this task.

You can choose following tasks:

- ▶ Change Sampling Rate
- ▶ View

Rebuild task

Use this task to extract the configuration information from the managed system and to rebuild the information about the HMC. Rebuilding the managed system means that you update, or refresh, the information about the HMC about the managed system. Rebuilding the managed system can be helpful when the state of the managed system is *incomplete*. The incomplete state means that the HMC lost communication with the managed server and no longer has complete information. Rebuilding the managed system is different from simply refreshing the HMC window. When the managed system is rebuilt, the HMC extracts the information from the managed system.

Rebuild can take several minutes: You cannot start other tasks while the HMC rebuilds the managed system. This process can take several minutes.

Change Password task

Use this task to change the HMC access password on the selected managed system. After the password is changed, you must update the HMC access password for all other HMCs from which you want to access this managed system.

Configuration task

Configuration includes the tasks for server configuration. The following list provides configuration tasks:

- ▶ Create logical partition
- ▶ System plans
- ▶ Partition availability priority
- ▶ View workload management groups
- ▶ Manage custom groups
- ▶ Manage partition data
- ▶ Manage system profiles
- ▶ Virtual resources

This section describes these tasks.

Create Logical Partition task

Use this task to create an AIX, Linux, Virtual I/O Server (VIOS), or IBM i logical partition on a managed system. The *create LPAR* wizard helps you to create a new logical partition and a default profile for the partition. We explain this task in 4.4.10, “Logical partition management” on page 138.

System Plans task

System plans records or deploys specifications for logical partitions, partition profiles, or hardware specifications on a chosen system. We explain this task in 2.1, “System planning tools” on page 14.

Partition Availability Priority task

Use this task to specify the partition availability priority of each logical partition on a managed system. The managed system uses partition availability priorities in the case of processor failure.

If a processor fails on a logical partition, and there are no unassigned processors available on the managed system, the logical partition can acquire a replacement processor from logical partitions with a lower partition-availability priority. This allows the logical partition with the higher partition-availability priority to continue running after a processor failure. We explain this task in 4.4.9, “Partition availability priority” on page 136.

View Workload Management Groups task

Use this task to display a detailed view of the workload management groups that you specified for a managed system, as shown in Figure 5-20. Each group displays the total number of processors, processing units for partitions using shared mode processing, and the total amount of memory that is allocated to the partitions in the group.

	Processors	Memory(MB)	State
▼ 8233-E8B-SN10DD51P			
▼ (None)	1.60	17,408	
lp1(3)	0.20	2,048	Running
lp2(4)	0.20	2,048	Not Activated
lp3(5)	0.20	2,048	Running
lp4(6)	0.50	10,240	Not Activated
vios1(1)	0.50	1,024	Not Activated
Total: 7			

Figure 5-20 Partition Workload Groups task

Manage Custom Groups task

The custom groups provide a mechanism to group system resources together in a single view. You can also nest groups (a group that is contained within a group) to provide hierarchical or topology views.

You can create others, delete the ones that were created, add to created groups, create groups by using the pattern match method, or delete from created groups by using this task.

Do the following steps to create a group:

1. Select one or more resources (for example: servers, partitions, or frames) that you want to include in the group.
2. Select **Configuration** → **Manage Custom Groups**.

3. Select **Create a new group**, specify a group name and description, and then click **OK**, as shown in Figure 5-21.
4. The new user-defined group is displayed in the navigation pane under Custom Groups and includes a selected resource.

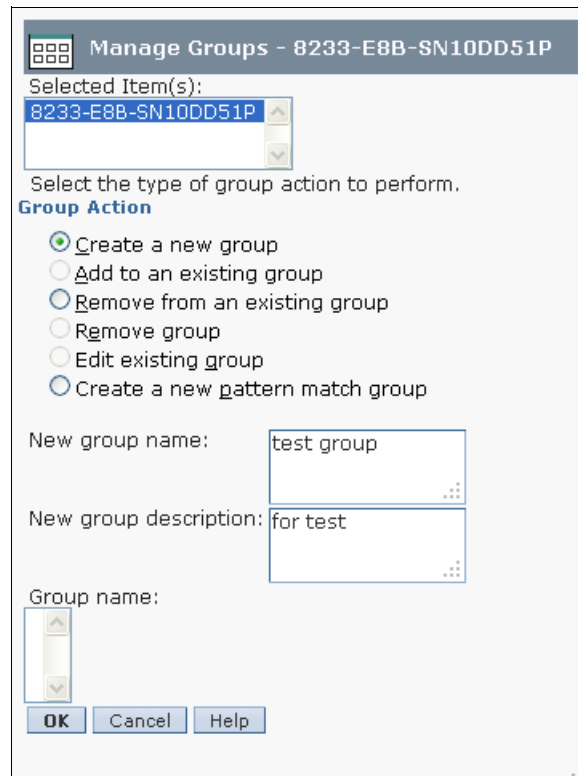


Figure 5-21 Manage Custom Groups task

Manage Partition Data task

Use this task to provide four operations, restore, initialize, backup, and delete to manage profile data. We explain this task in 5.4, “Managing partition data” on page 262.

Manage System Profile task

A system profile is an ordered list of partition profiles that is used by the HMC to start the logical partitions on a managed system in a specific configuration.

When you activate the system profile, the managed system attempts to activate each partition profile in the system profile in the order specified. A system profile helps you activate or change the managed system from one complete set of logical partition configurations to another. They can also be used to validate the

resource configuration of multiple partitions to ensure that resource conflicts do not exist between partitions. You can create system profile, as shown in Figure 5-22.

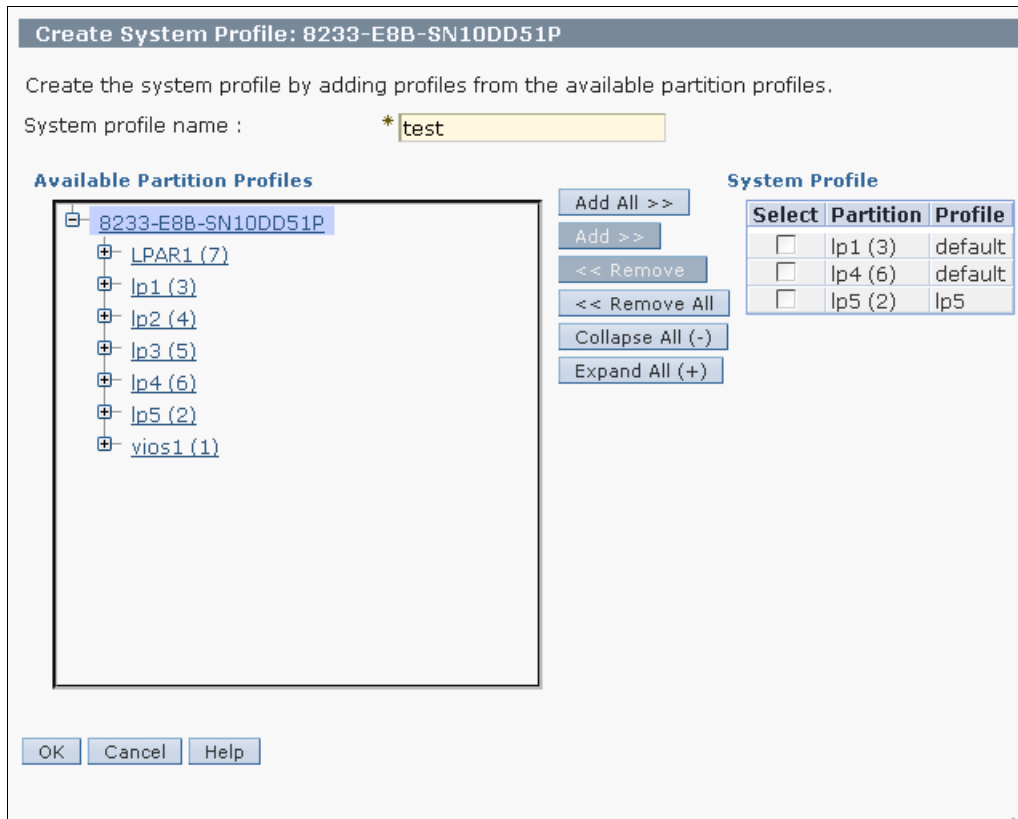


Figure 5-22 Manage System Profile task

Virtual Resources task

Use this task to manage the following virtual resources:

- ▶ Shared processor pool
- ▶ Shared memory pool
- ▶ Virtual storage
- ▶ Virtual network
- ▶ Reserved storage device pool

We explain this task in 4.4, “Virtualization using HMC” on page 121.

Connection task

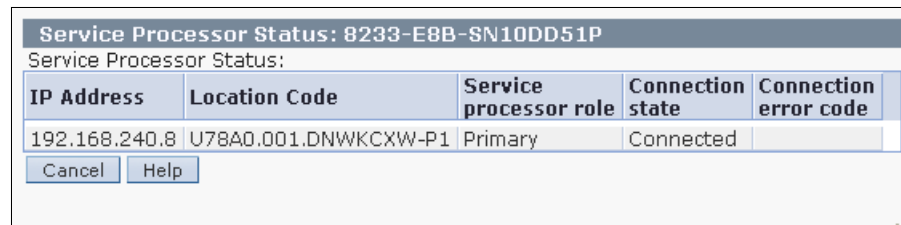
Connection tasks allow you to view the HMC connection status to service processors or frames, reset those connections, connect another HMC to the selected managed system, or connect another managed system to the HMC. This list provides connection tasks:

- ▶ Service processor status
- ▶ Reset or remove connections
- ▶ Disconnect another HMC
- ▶ Add managed system

This section describes these tasks.

Service Processor Status task

Use this task to display the HMC connection status to the service processor of a selected managed system, as shown in Figure 5-23. If you selected a frame, *service processor status* displays the state of the connection from the HMC to side A and side B of the bulk power assembly.



IP Address	Location Code	Service processor role	Connection state	Connection error code
192.168.240.8	U78A0.001.DNWKCXW-P1	Primary	Connected	

Figure 5-23 *Service processor status task*

Reset or Remove Connections task

Use this task to remove or reset a managed system from the contents area of the HMC, as shown in Figure 5-24.

When you remove the connection with a managed system, the connection is broken between the HMC and the managed system. Remove the connection with the managed system if you no longer want to manage the managed system by using this HMC. Use this window to remove the connection before you physically disconnect the HMC from the managed system (or from the network).

When you reset the connection with a managed system, the connection is broken and then reconnected. Reset the connection with the managed system if the managed system is in a *no connection* state and you verified that the network settings are correct on both the HMC and the managed system.

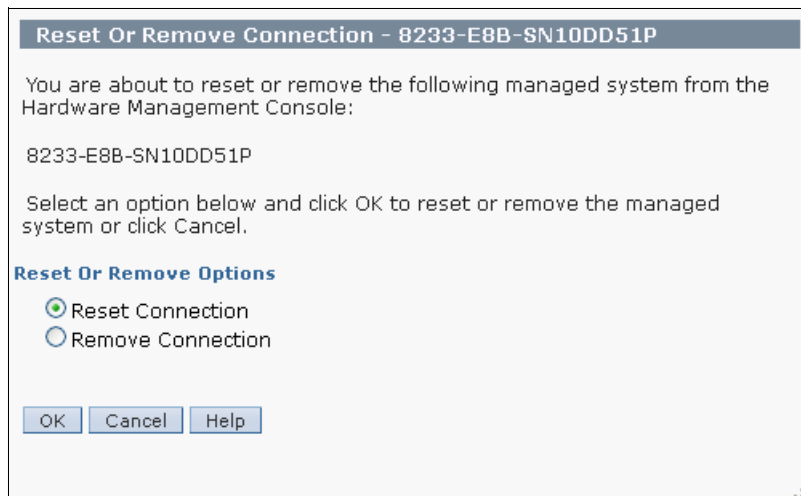


Figure 5-24 *Reset or remove connections task*

Disconnect Another HMC task

Use this task to disconnect another HMC from the selected managed system. Also, you can find which HMC locked the selected managed system. This task releases any lock that the other HMC might have on the selected managed system. After the disconnection is complete, the other HMC automatically attempts to reconnect to the managed system.

When you use an HMC to change a managed system, the HMC locks the managed system so that no other HMC can make conflicting changes at the same time. Normally, the HMC unlocks the managed system after the change is complete. If there is an error, and the managed system remains locked, you must

disconnect the HMC from the managed system to reset the lock before other HMCs can change the managed system.

Add Managed System task

Use this task to guide you through adding systems in the network to systems managed by this HMC. We explain this task in “Servers” on page 180.

Hardware Information task

Hardware information includes tasks to display information about the hardware that is attached to a selected managed system. Hardware information includes the following tasks:

- ▶ Adapter
- ▶ View hardware topology

This section describes these tasks.

Adapter task

Use this task to view information about the Host Ethernet Adapters (HEA, also referred to as Integrated Virtual Ethernet adapters) or host channel adapters (HCA) for a selected managed system.

- ▶ Host Ethernet

Use this task to display the port configuration and status of the physical Host Ethernet on the managed system, as shown in Figure 5-25.

Host Ethernet Adapters : 8233-E8B-SN10DD51P												
Select a physical port in the table below to display the port's current partition usage.												
Current Status												
Select	Physical Port	Location Codes	Port ID	Port Type	Port Group ID	Port Group MCS Value	Connection State	Speed	Duplex	Transmit Flow Control	Receive Flow Control	Maximum
<input type="radio"/>	U78A0.001.DNWKCXW-P1 - C6-T1	0	1 G	1	4		up	1 Gbps	full	disabled	disabled	1500
<input type="radio"/>	U78A0.001.DNWKCXW-P1 - C6-T2	1	1 G	1	4		down	Auto	full	disabled	disabled	1500
<input type="radio"/>	U78A0.001.DNWKCXW-P1 - C6-T3	0	1 G	2	4		down	Auto	full	disabled	disabled	1500
<input type="radio"/>	U78A0.001.DNWKCXW-P1 - C6-T4	1	1 G	2	4		down	Auto	full	disabled	disabled	1500
Configure...												
Logical Partition Usage												
Logical Partition	Logical Port ID	Logical Port DRC Name	Logical Port burned-in MAC / user-defined MAC	Capability	Allowed MAC Addresses	Allowed VLAN IDs	Promiscuous	LPAR				
OK Cancel Help												

Figure 5-25 Host Ethernet list

You can change the configuration of any of the ports on a Host Ethernet by selecting the Host Ethernet, selecting the port under Current Status, and clicking **Configure**. Then, the port configuration displays as shown in Figure 5-26. We explain these tasks in 4.4.5, “Host Ethernet Adapter” on page 133.

HEA Physical Port Configuration : 8233-E8B-SN10DD51P

Use the fields below to specify the configuration for the selected physical port.

Speed: Duplex:

Maximum receiving packet size: Pending Port Group Multi-Core Scaling value:

Flow control enabled Promiscuous LPAR:

Figure 5-26 Host Ethernet Physical Port Configuration

► Host Channel

Host channel adapters (HCAs) provide a managed system with port connections to other devices. That port can be connected to another HCA, a target device, or a switch that redirects the data coming in on one of its ports out to a device attached to another of its ports. You can show a list of the HCAs for the managed system. You can select an HCA from the list to display the current partition usage for the HCA. From this task you can display the following components:

- The physical location of each HCA on the managed system.
- The number of Globally Unique Identifiers (GUIDs) that are in use on each HCA.
- The number of GUIDs on each HCA that are available to be assigned to logical partitions.
- HMC management status. HCAs that are unable to be managed by an HMC are in an error state.
- The logical partition usage for a selected HCA.

View Hardware Topology task

Use this task to display the current RIO topology of the selected managed system. Current Topology displays the current topology. Any discrepancies

between the current topology and the last valid topology are identified as errors. The following information is shown:

- ▶ The *starting location* of the physical RIO cable and the RIO connection (cable to port).
- ▶ The *ending location* of the physical RIO cable and the RIO connection (cable to port).
- ▶ *Starting node type* displays the values of the node. Possible values are Local Bridge, Local NIC, Remote Bridge, and Remote NIC.
- ▶ *Link status* displays the leading port status.
- ▶ *Cable length* displays the length of the RIO cables. Errors occur when the actual cable lengths are different from the expected cable lengths.
- ▶ The serial number of the power-controlling managed system.
- ▶ The serial number of the function-controlling managed system.

Updates task

Use this task to do a guided update of managed system, power, or I/O Licensed Internal Code. We explain these tasks in 6.2, “Software maintenance” on page 309.

Serviceability task

Use this task to view specific events for selected systems. Problem Analysis on the HMC automatically detects error conditions and reports to you any problem that requires service to repair it. These problems are reported to you as serviceable events. We explain this task in 6.1, “Service Management” on page 268.

Capacity on Demand

Capacity on demand (CoD) allows you to nondisruptively activate (no boot required) processors and memory. CoD also gives you the option to temporarily activate capacity to meet intermittent performance needs to activate extra capacity on a trial basis, and to access capacity to support operations in times of need. We explain this task in 4.5, “Capacity on Demand” on page 160.

5.1.3 Systems Management: Partitions task

Systems Management displays tasks to manage servers, logical partitions, and frames. Use these tasks to set up, configure, view status, troubleshoot, and apply solutions for partitions.

This section describes the tasks to manage a partition.

Properties task

Properties include the tasks to display the properties of the selected partition, as shown in Figure 5-27.

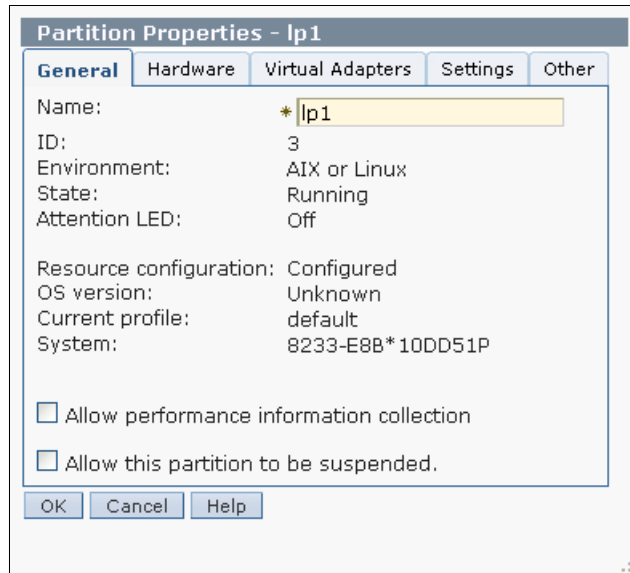


Figure 5-27 Partition Properties task

The partition properties task includes the following tabs:

- ▶ **General**

The *General* tab displays the name of the partition, ID, environment, state, resource configuration, operating system, the current profile that is used when starting the partition, and the system on which the partition is located.
- ▶ **Hardware**

The *Hardware* tab displays the current usage of processors, memory, and I/O on the partition.
- ▶ **Virtual Adapters**

The *Virtual Adapters* tab displays the current configuration of virtual adapters. Virtual adapters allow for the sharing of resources between partitions. From this tab, you can view, create, and edit virtual adapters on the partition.
- ▶ **Settings**

The *Settings* tab displays the boot mode and keylock position of the partition. Also displayed are the current service and support settings for the partition.
- ▶ **Other**

The *Other* tab displays the partition's Workload Management Group (if applicable), and the partition's controlling partitions.

Change Default Profile task

Change Default Profile includes the task to change the default profile for the partition. Select a profile from the drop-down list to be the new default profile, as shown in Figure 5-28.

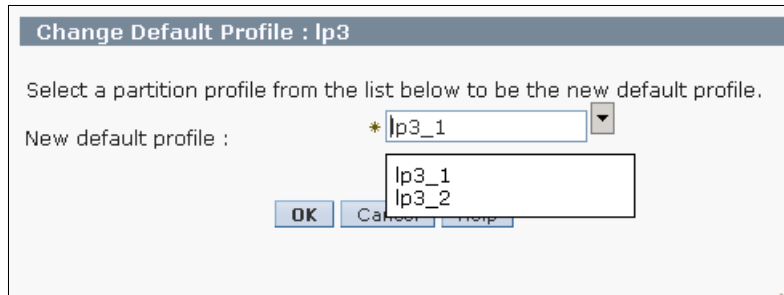


Figure 5-28 *Change Default Profile* task

Operations task

Operations include the tasks for partition operations. The following list provides the operations tasks:

- ▶ Activate
- ▶ Restart
- ▶ Shut Down
- ▶ Deactivate Attention LED
- ▶ Schedule Operations
- ▶ Delete
- ▶ Mobility
- ▶ Suspend Operations

This section describes these tasks.

Activate task

Use the Activate task to activate a partition on the managed system in the *Not Activated* state. Choose from the following options:

► Profile

This option is to select a profile to use when activating the partition from the list of profiles displayed. Select from the list of profiles and click **OK** to activate the partition, as shown in Figure 5-29.

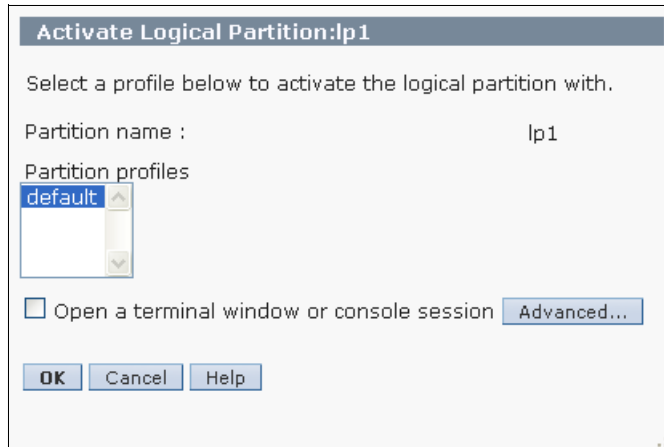


Figure 5-29 Activate Logical Partition task

► Current configuration

You can use this option to quickly start multiple logical partitions by using the current configuration information for the partitions that is available on the hypervisor.

Current configuration: The HMC cannot activate a partition by using its current configuration if the partition does not have an active profile that is associated with it, such as a newly created partition.

Restart task

Use this task to restart the selected logical partition or partitions, as shown in Figure 5-30.

For IBM i logical partitions, use this window only if you cannot restart the IBM i logical partition from the command line of the operating system. Using this window to restart an IBM i logical partition results in an abnormal IPL.

If you choose to restart Virtual I/O Server partitions that are acting as the Paging Service Partition (PSP) for a number of client partitions, a warning displays. This warning indicates to should shut down the client partitions before shutting down the Virtual I/O Server partition.

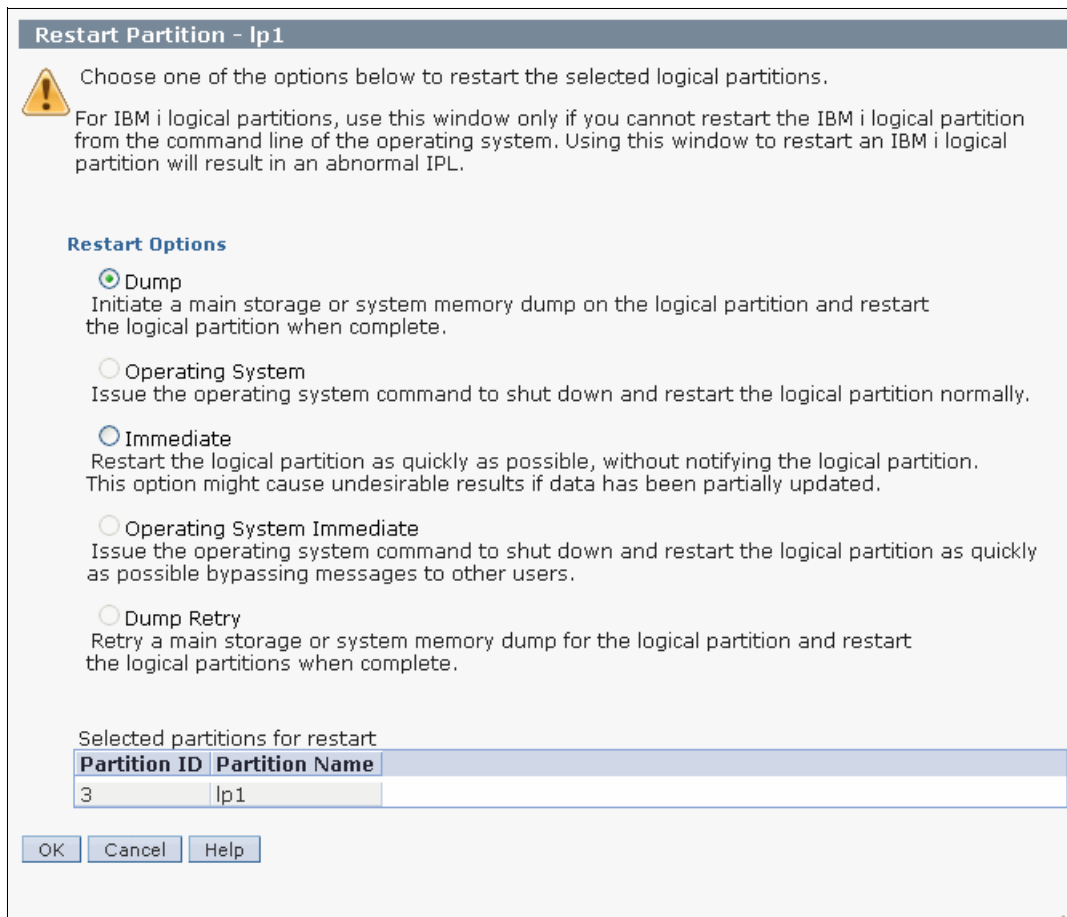


Figure 5-30 Restart Partition task

Choose from the following options:

Operating System options: The *Operating System* option and the *Operating System Immediate* option are only enabled if Resource Monitoring and Control (RMC) is up and configured.

▶ **Dump**

The HMC shuts down the logical partition and initiates a main storage or system memory dump. For AIX and Linux logical partitions, the HMC also notifies the logical partition that it is shut down. For IBM i logical partitions, the processors are stopped immediately. After the shutdown is complete, the logical partition is immediately restarted. (IBM i logical partitions are restarted multiple times so that the logical partition can store the memory dump information.) Use this option if a portion of the operation system appears hung and you want a memory dump of the logical partition for analysis.

▶ **Operating System**

The HMC shuts down the logical partition normally by issuing a **shutdown -r** command to the logical partition. During this operation, the logical partition does any necessary shutdown activities. After the shutdown is complete, the logical partition is immediately restarted. This option is only available for AIX logical partitions.

▶ **Immediate**

The HMC shuts down the logical partition immediately. The HMC ends all active jobs immediately. The programs running in those jobs are not allowed to do any job cleanup. This option might cause undesirable results if data is partially updated. Use this option only after a controlled end was unsuccessfully attempted.

▶ **Operating System Immediate**

The HMC shuts down the logical partition immediately by issuing a **shutdown -Fr** command to the logical partition. During this operation, the logical partition bypasses messages to other users and other shutdown activities. After the shutdown is complete, the logical partition is immediately restarted. This option is only available for AIX logical partitions.

▶ **Dump Retry**

The HMC tries a main storage or system memory dump again on the logical partition. After this step is complete, the logical partition is shut down and restarted. Use this option only if you previously tried the memory dump option without success. This option is only available for IBM i logical partitions.

Shut Down Partitions task

Use this task to shut down the selected logical partition or partitions, as shown in Figure 5-31.

For IBM i logical partitions, use this window only if you cannot shut down the IBM i logical partition from the command line of the operating system. Using this window to shut down an IBM i logical partition results in an abnormal IPL.

If you choose to shut down Virtual I/O Server partitions that are acting as the Paging Service Partition (PSP) for a number of client partitions, a warning displays. This warning indicates to shut down the client partitions before shutting down the Virtual I/O Server partition.

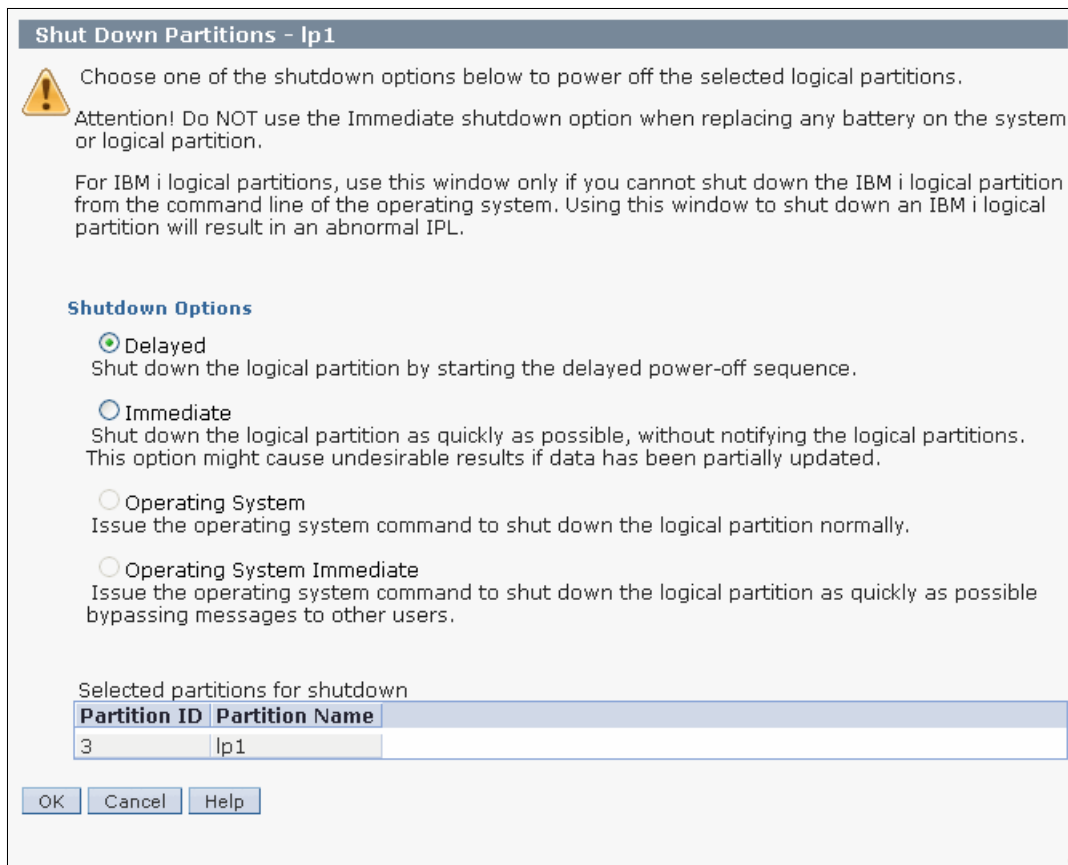


Figure 5-31 Shut Down Partitions task

Choose from the following options:

▶ **Delayed**

The HMC shuts down the logical partition by using the delayed power-off sequence. This process allows the logical partition time to end jobs and write data to disks. If the logical partition is unable to shut down within the predetermined amount of time, it ends abnormally and the next restart might be longer than normal.

▶ **Immediate**

The HMC shuts down the logical partition immediately. The HMC ends all active jobs immediately. The programs running in those jobs are not allowed to do any job cleanup. This option might cause undesirable results if data is partially updated. Use this option only after a controlled shutdown was unsuccessfully attempted.

▶ **Operating System**

The HMC shuts down the logical partition normally by issuing a **shutdown** command to the logical partition. During this operation, the logical partition does any necessary shutdown activities. This option is only available for AIX logical partitions.

▶ **Operating System Immediate**

The HMC shuts down the logical partition immediately by issuing a **shutdown -F** command to the logical partition. During this operation, the logical partition bypasses messages to other users and other shutdown activities. This option is only available for AIX logical partitions.

Deactivate Attention LED task

Use this task to activate or deactivate an attention LED on your partition. All attention LEDs for the partitions on the managed system are listed. Select an LED and choose to activate or deactivate.

Schedule Operations task

Use this task to create a schedule for certain operations to be done on the logical partition without operator assistance. 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. For example, you can schedule an operation to remove resources from a logical partition or move resources from one logical partition to another.

You can choose the following tasks:

- ▶ Activate an LPAR
- ▶ Dynamic Reconfiguration
- ▶ Operating System Shutdown (on a partition)

How to create and check tasks is explained in “Schedule Operations task” on page 192.

Delete task

Use this task to delete the selected partition, as shown in Figure 5-32.

The Delete task deletes the selected partition and all of the partition profiles associated with the partition from the managed system. When you delete a partition, all hardware resources currently assigned to that partition become available to other partitions.

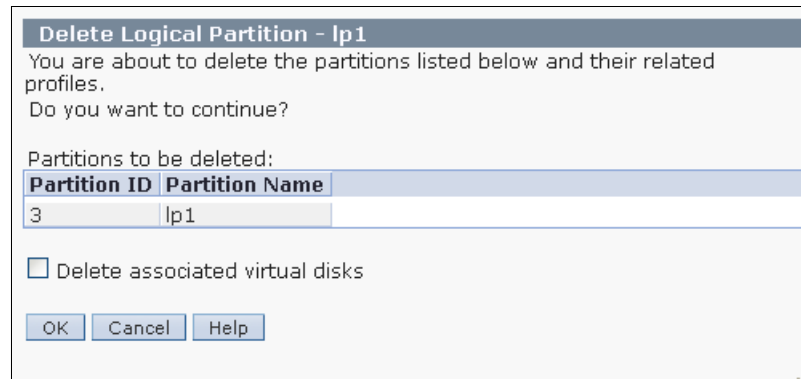


Figure 5-32 Delete Logical Partition task

Mobility task

Use this task to migrate your partition to another server, ensure that the requirements for the migration are met, and recover if the partition is in an invalid state.

Partition mobility is the ability to migrate a logical partition, including its operating system and applications, from one system to another system. Partition migration can be active or inactive. Active partition migration lets you move a running logical partition from one system to another system without shutting down the partition or applications running on the partition. Inactive partition migration lets you move a powered-off logical partition from one system to another system.

In this task, you can choose from the following options:

- ▶ Migrate
- ▶ Validate
- ▶ Recover

For more information about this task, refer to *IBM PowerVM Live Partition Mobility, SG24-7460-01*.

Suspend Operations task

Use this task to suspend a running logical partition or to resume a suspended logical partition. You can also validate that the logical partition meets all requirements for being suspended or for being resumed.

You can also recover from a failed suspend and resume actions. Recovery cleans up partitions that are in a not valid state as a result of an incomplete suspend or resume.

In this task, you can choose from the following options:

- ▶ Suspend
- ▶ Resume
- ▶ Recover

For more information about this task, refer to *IBM PowerVM Virtualization Managing and Monitoring, SG24-7590-03*.

Configuration task

Configuration includes the tasks for partition configuration. These tasks include:

- ▶ Manage profiles
- ▶ Manage custom groups
- ▶ Save current configuration

This section describes these tasks.

Manage Profiles task

Use this task to create, edit, copy, delete, or activate a profile for the selected partition. A partition profile contains the resource configuration for the partition. You can modify the processor, memory, and adapter assignments for a profile by editing the profile.

The default partition profile for a logical partition is the partition profile that is used to activate the logical partition if no other partition profile is selected. You cannot delete the default partition profile unless you first designate another partition profile as the default partition profile. The default profile is defined in the status column, as shown in Figure 5-33.

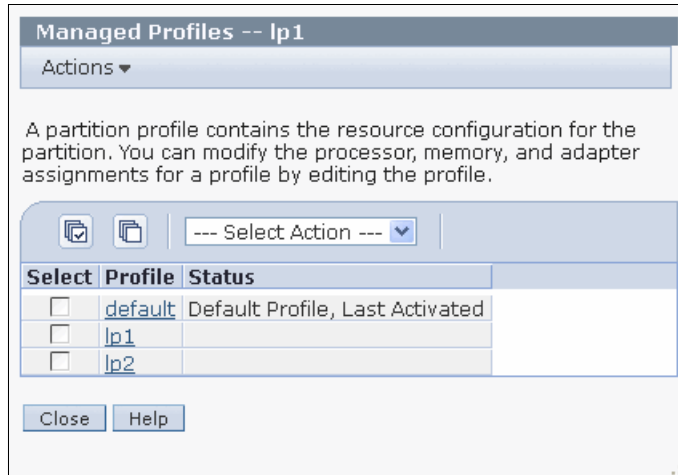


Figure 5-33 Managed Profiles task

If you want to know how to create a logical partition and a partition profile, we explain in 4.4.10, “Logical partition management” on page 138.

Manage Custom Groups task

Custom groups are composed of logical collections of objects. You can report status on a group basis, which allows you to monitor your system in a way that you prefer. You can also nest groups (a group that is contained within a group) to provide hierarchical or topology views. Use this task to create a custom group with selected partition or add the selected partition to an existing custom group.

We explain this task in “Manage Custom Groups task” on page 196.

Save Current Configuration task

Use this task to save the current configuration of a logical partition, as shown in Figure 5-34. You can choose to create a partition profile or to overwrite an existing partition profile. This procedure is useful if you change the configuration of a logical partition by using dynamic logical partitioning and you do not want to lose the changes when you restart the logical partition. You can do this procedure at any time after you initially activate a logical partition.

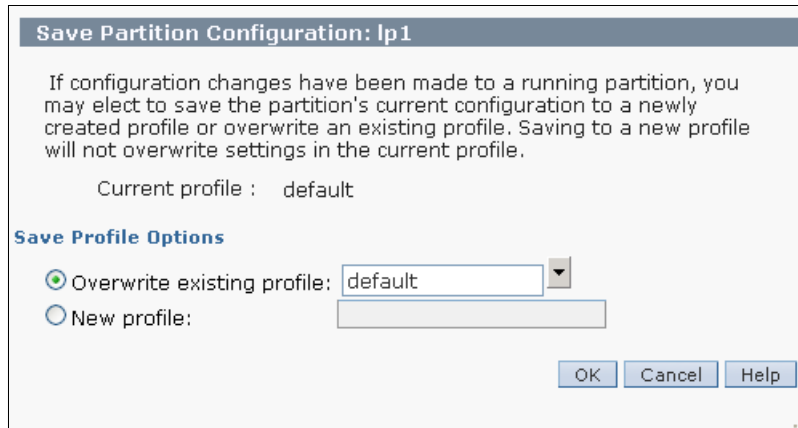


Figure 5-34 Save Partition Configuration task

Hardware Information task

Hardware information includes tasks to display information about the hardware that is attached to a selected managed system. Hardware information includes the following tasks:

- ▶ Adapter
- ▶ Virtual I/O adapter

This section describes these tasks.

Adapter task

Use this task to view information about the Host Ethernet Adapters (HEAs, also referred to as Integrated Virtual Ethernet adapters) or host channel adapters (HCAs) for a selected managed system.

- ▶ Host Channel

Host Channel Adapters (HCAs) provide a managed system with port connections to other devices. That port can be connected to another HCA, a target device, or a switch that redirects the data coming in on one of its ports out to a device attached to another of its ports. You can show a list of the HCAs for the managed system. You can select an HCA from the list to display

the current partition usage for the HCA. From this task you can display the following components:

- The physical location of each HCA on the managed system.
- The number of GUIDs that are in use on each HCA.
- The number of GUIDs on each HCA that are available to be assigned to logical partitions.
- HMC management status. HCAs that are unable to be managed by an HMC are in an error state.
- The logical partition usage for a selected HCA.

► Host Ethernet

Use this task to display the port configuration and status of the physical Host Ethernet on the managed system.

We explain these tasks in 4.4.5, “Host Ethernet Adapter” on page 133.

► Switch Network Interface

Use this task to display a list of the Switch Network Interface (SNI) adapters for the selected managed system. The following is displayed, the SNI adapter handle, the name of the partition to which the adapter is assigned, the physical location of the adapter, and the host name or IP address of the adapter.

Virtual I/O Adapter task

Use this task to view the topology of currently configured virtual Small Computer System Interface (SCSI) and virtual Ethernet adapters on a selected partition.

► SCSI

Use the *SCSI* task to view the topology of virtual SCSI adapters on a partition. The following information is displayed:

- Adapter name
- Backing device
- Remote partition
- Remote adapter
- Remote backing device

► Ethernet

Use the *Ethernet* task to view the current virtual Ethernet configuration for the partition. The following information is displayed:

- Adapter name
- Virtual LANs
- I/O Server
- Server virtual adapter
- Shared adapter

Dynamic Logical Partitioning task

Dynamic logical partitioning includes tasks that allow you to dynamically add or remove processors, memory, and adapters to and from logical partitions.

This section describes these tasks.

Processor task

Use this task to add or remove processor resources from a logical partition or to move processor resources from one logical partition to another. These tasks include:

- ▶ Add or Remove

Use this task to add processor resources to or remove processor resources from the selected logical partition without restarting the logical partition.

- ▶ Move

Use this task to move processor resources from the selected logical partition to another logical partition without restarting either logical partition.

Memory task

Use this task to add or remove memory resources from a logical partition or to move memory resources from one logical partition to another. These tasks include:

- ▶ Add or Remove

Use this task to add memory to or remove memory from the selected logical partition without restarting the logical partition.

- ▶ Move

Use this task to move memory from the selected logical partition to another logical partition without restarting either logical partition.

Physical Adapters task

Use this task to add I/O slots to a logical partition without restarting the partition or to move or remove I/O slots from a logical partition without restarting the partition. These tasks include:

- ▶ Add

Use this task to add I/O slots to a logical partition without restarting the partition. When you add an I/O slot to a logical partition, the I/O adapter in that I/O slot and the devices that are controlled by the I/O adapter can be used by the logical partition. This function is typically used to share infrequently used devices among logical partitions by moving these devices from one logical partition to another.

► Move or Remove

Use this task to remove I/O slots from a logical partition or move I/O slots between logical partitions without restarting the logical partitions. When you remove an I/O slot from a logical partition, the I/O adapter in that I/O slot and the devices that are controlled by the I/O adapter are also removed from the logical partition. If you choose to move the I/O slot to another logical partition, the I/O adapter and the devices that are controlled by the I/O adapter are also moved to the other logical partition. This function is typically used to share infrequently used devices among logical partitions by moving these devices from one logical partition to another.

Recommendation: It is recommended that you vary off the I/O slot and all I/O adapters and devices that are connected to the I/O slot before you remove the I/O slot from the logical partition.

Virtual Adapters task

Use this task to display a list of all of the virtual adapters that currently exist for this logical partition or partition profile. You can also create, change, or remove a virtual adapter on a logical partition or in a partition profile.

Host Ethernet task

Use this task to add Logical Host Ethernet Adapter (LHEA) logical ports dynamically to a running logical partition.

► Add

Use this task to add LHEA logical ports dynamically to a running logical partition. These logical ports allow the logical partition to access and use the physical port resources on a physical HEA.

Some operating system or system software versions do not allow for the adding of logical ports dynamically. For more information, consult the documentation for the operating system or system software.

► Move or Remove

Use this task to move LHEA logical ports dynamically from the selected logical partition. These logical ports allow the logical partition to access and use the physical port resources on a physical HEA. You can move the logical ports dynamically to another running logical partition, or you can leave the logical ports as unassigned. Some operating system or system software versions do not allow you to move or remove logical ports dynamically. For more information, consult the documentation for the operating system or system software.

Console Window task

Console window includes the task to open a terminal window to the operating system that is running on the selected partition.

Choose from the following options:

- ▶ Open terminal window
- ▶ Close terminal window
- ▶ Open shared 5250 console
- ▶ Open dedicated 5250 console

Serviceability task

Use this task to view specific events for selected systems. Problem analysis on the HMC automatically detects error conditions and reports any problem to you that requires service to repair it. These problems are reported to you as serviceable events. We explain this task in 6.1, “Service Management” on page 268.

5.1.4 Systems Management: Frames task

Systems Management displays tasks to manage servers, logical partitions, and frames. Use these tasks to set up, configure, view status, troubleshoot, and apply solutions for frames.

This section describes the tasks to manage a frame.

Frames option: The *Frames* option is shown in the navigation pane only when the HMC manages the server that consists of one or more frames such as the IBM Power 795, as shown in Figure 5-35.

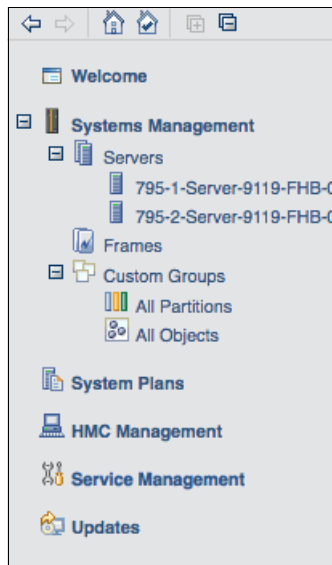


Figure 5-35 Frames option in the navigation pane

Properties task

Properties include the tasks to display the selected properties of a managed frame.

These properties include:

- General

The *General* tab displays the frame name and number, state, type, model, and serial number.

- ▶ **Managed Systems**

The *Managed Systems* tab displays all of the managed systems included in the frame and their cage numbers. A cage is a division of the enclosure that holds the managed systems, the I/O units, and the bulk power assemblies (BPAs).

- ▶ **I/O Units**

The *I/O Units* tab displays all of the I/O units that are contained in the frame, their cage numbers, and their assigned managed systems. A cage is a division of the enclosure that holds the managed systems, the I/O units, and the BPAs. *Not owned* in the System column indicates that the corresponding I/O unit is not assigned to a managed system.

Update Password task

Update password includes the task to update HMC access and Advanced System Management Interface (ASMI) passwords on the managed system. The first time that you access a managed system by using an HMC, the system prompts you to enter passwords for each of the following components:

- ▶ Hardware Management Console: HMC access
- ▶ Advanced System Management Interface: General
- ▶ Advanced System Management Interface: Admin

If you are using an HMC to access the managed system before all required passwords are set, enter the appropriate password for each password that is presented in this task.

Operations task

Operations includes the tasks for frame operations. The following list provides the operations tasks:

- ▶ Initialize frames
- ▶ Initialize all frames
- ▶ Rebuild
- ▶ Change password
- ▶ Power on or power off I/O unit

This section describes these tasks.

Initialize Frames task

Use this task to initialize a frame. When you initialize a managed frame, all of the frames that are managed by the HMC are powered on. As each individual frame is powered on, the I/O units that are contained within the frame are powered on as well. When all the I/O units for the frame are powered on, the managed systems that are contained within the frame are powered on. The complete initialization process can take several minutes to complete.

Power: Managed systems that are already powered on are not affected. They are not powered off and back on again.

Initialize All Frames task

Use this task to initialize all your frames. The unowned I/O units are first powered on within each managed frame, then the managed systems are powered on within each managed frame.

Frames:

- ▶ This operation task is available when no managed frame is selected and the frames tab on the navigation area is highlighted.
- ▶ Frames are already powered on when they are connected to HMC. Initializing frames does not power on the frames.

Rebuild task

Use this task to update information about the frame on the HMC interface. Updating, or rebuilding, the frame acts much like a refresh of the frame information. Rebuilding the frame is useful when the system's state indicator in the work pane of the HMC is shown as *Incomplete*. The Incomplete indicator signifies that the HMC cannot gather complete resource information from the managed system within the frame. No other tasks can be performed on the HMC during this process, which can take several minutes.

Change Password task

Use this task to change the HMC access password on the selected managed frame. After the password is changed, you must update the HMC access password for all other HMCs from which you want to access this managed frame.

Power On or Power Off I/O Unit task

Use this task to power off an I/O unit. Only units or slots that are in a power domain can be turned off. The corresponding power-on or off buttons are disabled for location codes that are not controllable by the HMC.

Configuration task

Configuration includes the task for server configuration. Manage custom groups is part of the Configuration task.

This task is now described.

Manage Custom Groups task

Custom groups are composed of logical collections of objects. You can report status on a group basis, allowing you to monitor your system in a way that you prefer. You can also nest groups (a group that is contained within a group) to provide hierarchical or topology views. Use this task to create a custom group with selected frame or add the selected frame to an existing custom group.

We explain this task in “Manage Custom Groups task” on page 196.

Connections task

Connections tasks allow you to view the HMC connection status to frames or reset those connections. The following components are part of the connections task:

- ▶ Bulk Power Assembly (BPA) status
- ▶ Reset

These tasks are now described.

Bulk Power Assembly Status task

Use this task to view the state of the connection from the HMC to side A and side B of the Bulk Power Assembly (BPA). The HMC operates normally with a connection to either side A or side B. However, for code update operations and some concurrent maintenance operations, the HMC needs connections to both sides.

Reset task

Use this task to reset the connection between the HMC and the selected managed frame. When you reset the connection with a managed frame, the connection is broken and then reconnected. Reset the connection with the managed frame if the managed frame is in a *No Connection* state. Also ensure that you verified that the network settings are correct on both the HMC and the managed frame.

Hardware Information task

Hardware information includes the tasks to display information about the hardware that is attached to a selected managed system. View remote input/output (RIO) topology is part of the Hardware Information task. This task is now described.

View RIO Topology task

Use this task to display the current RIO topology of the selected managed system. Current Topology displays the current topology. Any discrepancies

between the current topology and the last valid topology are identified as errors. The following information is shown:

- ▶ The *starting location* of the physical RIO cable and the RIO connection (cable to port).
- ▶ The *ending location* of the physical RIO cable and the RIO connection (cable to port).
- ▶ *Starting node type* displays the values of the node. Possible values are Local Bridge, Local NIC, Remote Bridge, and Remote NIC.
- ▶ *Link status* displays the leading port status.
- ▶ *Cable length* displays the length of the RIO cables. Errors occur when the actual cable lengths are different from the expected cable lengths.
- ▶ The serial number of the power-controlling managed system.
- ▶ The serial number of the function-controlling managed system.

Serviceability task

Use this task to view specific events for selected systems. Problem analysis on the HMC automatically detects error conditions and reports any problem to you that requires service to repair it. These problems are reported to you as serviceable events. This task is explained in 6.1, “Service Management” on page 268.

5.1.5 Using the command-line interface

The command-line interface (CLI) is an alternative to performing tasks on the HMC graphical user interface. This section provides information about the most common command-line options and usage. You can use the command-line interface in the following situations:

- ▶ Consistent results are required
If you have to administer several managed systems, you can achieve consistent results by using the command-line interface. The command sequence can be stored in scripts and run remotely.
- ▶ Automated operations are required
After you develop a consistent way to manage the managed systems, you can automate the operations by starting the scripts from batch-processing applications, such as the cron daemon, from other systems.

Before using remote command-line interface, you must enable the access to the HMC by using the SSH facility. We explain this task in 5.2.28, “Remote Command Execution task” on page 245.

You can generally choose from two options to use CLI on the HMC: Restricted Shell Terminal task on the local HMC or Secure Shell (SSH) client.

We explain the *Open Restricted Shell Terminal* task in 5.2.4, “Open Restricted Shell Terminal task” on page 233. If you use SSH client to connect the HMC, you must ensure that your script executions between SSH clients and the HMC are secure. For details about each SSH client setting, check documentation about them. For general information to set up the SSH client, see the Systems Hardware information center at this website:

<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7ha1/settingupsecurescriptexecution.htm>

Most common command-line options and usage

Some of the most commonly used command-line options are now described.

HMC command information, including examples, is available from the HMC command line by using the `man` command. You can find a man page for every command. To view the command information, type `man` and then the command name.

You can also access a PDF that includes all the man pages at the HMC website:

<http://www14.software.ibm.com/webapp/set2/sas/f/hmc1/resources.html>

Commands to manage HMC itself

► **lshmc**

Definition: List HMC configuration information

Example:

- List the BIOS level of the HMC
`lshmc -b`
- List the network settings for the HMC
`lshmc -n`
- List the VPD information for the HMC
`lshmc -v`
- List the version information for the HMC
`lshmc -V`

► **chhmc**

Definition: Change HMC configuration information

Example:

- Change the HMC host name
`chhmc -c network -s modify -h newhostname`
- Set the IP address and network mask for network interface eth0
`chhmc -c network -s modify -i eth0 -a 10.10.10.1 -nm 255.255.255.0`

► hmcshutdown**Definition:** Shut down the HMC**Example:**

- Reboot the HMC after 3 minutes
`hmcshutdown -t 3 -r`
- Halt the HMC immediately
`hmcshutdown -t now`

Commands to manage system configuration**► lssyscfg****Definition:** List system resources**Examples:**

- List all systems that are managed by this HMC
`lssyscfg -r sys`
- List all partitions in the managed system, and display only attribute values for each partition, following a header of attribute names
`lssyscfg -r lpar -m system1 -F --header`

► mksyscfg**Definition:** Create system resources**Example:**

- Create an AIX or Linux partition
`mksyscfg -r lpar -m system1 -i "name=aix_lpar2,profile_name=prof1,lpar_env=aixlinux,min_mem=256,desired_mem=1024,max_mem=1024,proc_mode=ded,min_procs=1,desired_procs=1,max_procs=2,sharing_mode=share_idle_procs,auto_start=1,boot_mode=norm,lpar_io_pool_ids=3,"io_slots=21010003/3/1,21030003//0"`

- Create partition profiles by using the configuration data in the file /tmp/profcfg
`mksyscfg -r prof -m system1 -f /tmp/profcfg`
- Create a partition profile by saving the current configuration of a partition
`mksyscfg -r prof -m system1 -o save -p p1 -n newProfile`

► **chsyscfg**

Definition: Change system resources

Examples:

- Change a partition profile's memory amounts (reduce the profile's current memory amounts each by 256 MB), and number of wanted processors
`chsyscfg -r prof -m system1 -i
 "name=profile1,lpar_name=part1,min_mem=256,desired_mem=256,max_mem=256,desired_procs=2"`
- Change a system profile (add two new partition profiles)
`chsyscfg -r sysprof -m system1 -i
 "name=sysprof1,"lpar_names+=part3,part4","profile_names+=3_prof1,
 4_defaultProf"`

► **rmsyscfg**

Definition: Remove a system resource

Example:

- Remove the partition lpar1
`rmsyscfg -r lpar -m system1 -n lpar1`
- Remove the partition profile test_profile for partition lpar1
`rmsyscfg -r prof -m system1 -n test_profile -p lpar1`

Commands to manage hardware resources

► **lshwres**

Definition: List hardware resources

Examples:

- List all system level memory information
`lshwres -r mem -m system1 --level sys`
- List all virtual slots for partition lpar1
`lshwres -r virtualio --subtype slot -m system1 --level slot
 --filter "lpar_names=lpar1"`

► **chhwres**

Definition: Change hardware resources

Examples:

- Move 0.5 processing units from the partition with ID 1 to the partition with ID 2 (both partitions are using shared processors)

```
chhwres -r proc -m system1 -o m --id 1 --tid 2 --procunits .5
```

- Add 128 MB of memory to the partition with ID 1, and time out after 10 minutes

```
chhwres -r mem -m system1 -o a --id 1 -q 128 -w 10
```

- Move the partition sharedlpar1 to shared processor pool pool1

```
chhwres -r procpool -m system1 -o s -p sharedlpar1 -a  
"shared_proc_pool_name=pool1"
```

- Add a virtual Ethernet adapter to the partition with ID 3

```
chhwres -r virtualio -m system1 -o a --id 3 --rsubtype eth -a  
"ieee_virtual_eth=1,port_vlan_id=4,"addl_vlan_ids=5,6",is_trunk=1  
,trunk_priority=1"
```

Commands to manage system connection

► **lssysconn**

Definition: List system connections

Example:

- Lists connection information for all systems and frames that are managed by this HMC

```
lssysconn -r all
```

► **mksysconn**

Definition: Create system connection

Example:

- Connect to and add the system with the IP address 9.3.152.145 (the HMC Access password for the system must be entered when prompted)

```
mksysconn --ip 9.3.152.145
```

- To enable all systems and frames to be automatically discovered by the HMC when using DHCP

```
mksysconn -o auto
```

► **rmsysconn**

Definition: Remove system connection

Example:

- Disconnect from the managed system system1 and remove it from the HMC

```
rmsysconn -o remove -m system1
```

Commands to manage users on the HMC**▶ lshmcusr**

Definition: List HMC user information

Example:

- List all HMC users

```
lshmcusr
```

- List only the user names and managed resource roles for all HMC users, and separate the output values with a colon

```
lshmcusr -F name:resourcerole
```

- List the HMC users hscroot and user1

```
lshmcusr --filter '"names=hscroot,user1"'
```

▶ mkhmcusr

Definition: Create HMC user information

Example:

- Create the user myhmcuser (the user's password must be entered when prompted)

```
mkhmcusr -u myhmcuser -a hmcviewer
```

or

```
mkhmcusr -i "name=myhmcuser,taskrole=hmcviewer"
```

▶ chhmcusr

Definition: Change HMC user information

Example:

- Change the password for the user tester (the new password must be entered when prompted)

```
chhmcusr -u tester -t passwd
```

► **rmhmcusr**

Definition: Remove HMC user information

Example:

- Remove the user tester
`rmhmcusr -u tester`

Commands for Capacity on Demand

► **lscod**

Definition: List CoD information

Example:

- Display CUoD processor capacity information
`lscod -m system1 -t cap -r proc -c cuod`
- Display CUoD processor activation code generation information
`lscod -m system1 -t code -r proc -c cuod`
- Display the CoD history log
`lscod -m system1 -t hist`

► **chcod**

Definition: Change CoD

Example:

- Enter a CoD code
`chcod -m system1 -o e -k code`
- Activate 2 GB of On/Off CoD memory for 10 days
`chcod -m system1 -o a -c onoff -r mem -q 2048 -d 10`
- Deactivate all On/Off CoD processors
`chcod -m system1 -o d -c onoff -r proc`

Commands for virtual terminals

► **mkvterm**

Definition: Open a virtual terminal session

Example:

- Open a virtual terminal session for partition lpar1
`mkvterm -m system1 -p lpar1`

► **rmvterm**

Definition: Close a virtual terminal session

Example:

- Close a virtual terminal session for partition lpar1

```
rmvterm -m system1 -p lpar1
```

5.2 HMC Management task

This task includes a categorized or alphabetical view of HMC management tasks and their descriptions. These tasks are used for setting up the HMC, maintaining its internal code, and securing the HMC.

To display the tasks in the work pane:

1. Select the **HMC Management** node in the navigation pane.
2. From the work pane, select the task that you want to perform.
3. By default, a categorized listing of the tasks displays. The following functions are displayed:
 - Operations
 - Administration

If you want to see that level of the HMC with which you are currently working, point your mouse over **HMC Version** at the top of the work pane, as shown in Figure 5-36.

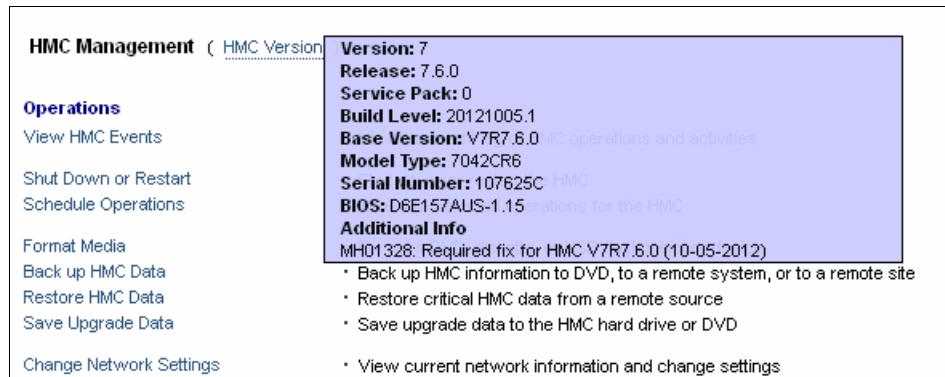


Figure 5-36 Confirm the HMC version

5.2.1 Lock HMC Screen task

Use this task to lock the HMC window. When you choose this task, the HMC window is locked immediately.

To unlock the window and return to the HMC workplace, press Enter and specify the password for the user ID for which you are logged in.

Remote HMC access: If you access the HMC remotely, this task is not shown in the work pane.

5.2.2 View HMC Events task

Use this task to show the console event logs on the HMC. The HMC keeps a log of significant operations and activities automatically, referred to as *console events* that occur while the application is running. Thus, system events are individual activities that indicate when processes occur, begin and end, and succeed or fail.

When an event occurs, the date and time it occurs, and a brief description of the event are recorded in the event log, as shown in Figure 5-37.

Date	Time	Console Event
10/29/2012	14:02:16.030	User hscroot has logged off from session id 9 for the reason: The user logged off.
		[EVENT 1831]: A TIMER TASK HAS BEEN SCHEDULED ON A NAMED TIMER Named Timer Thread: PM-TaskPortal Timer Timer Task ID: 380638896 Date: Delay (s): 0 Period (s):
10/29/2012	14:01:40.180	
10/29/2012	13:59:16.910	User hscroot of session 9 is using user interface "Tree Style".
10/29/2012	13:59:16.870	User hscroot has logged on from the console to session id 9. The user's maximum role is "hmcsuperadmin".
		[EVENT 1831]: A TIMER TASK HAS BEEN SCHEDULED ON A NAMED TIMER Named Timer Thread: PM-TaskPortal Timer Timer Task ID: 971520488 Date: Delay (s): 0 Period (s):
10/29/2012	13:21:33.270	
10/29/2012	12:34:44.200	User hscroot of session 8 is using user interface "Tree Style".
10/29/2012	12:34:44.150	User hscroot has logged on from location 172.16.254.6 to session id 8. The user's maximum role is "hmcsuperadmin".
10/29/2012	11:36:17.050	User hscroot has logged off from session id 2 for the reason: The user logged off.
		[EVENT 1831]: A TIMER TASK HAS BEEN SCHEDULED ON A NAMED TIMER Named Timer Thread: PM-TaskPortal Timer Timer Task ID: 1914466844
		Total: 549 Filtered: 549

Figure 5-37 View HMC events task

5.2.3 Open 5250 Console task

Use this task to open a 5250 emulator session so you can communicate with an IBM i logical partition.

To open a 5250 console, open the Open 5250 Console task from the HMC Management work pane. The 5250 Setup window displays. From the 5250 Setup window, you can configure and start your 5250 emulator.

Remote HMC access: If you access the HMC remotely, this task is not shown in the work pane.

5.2.4 Open Restricted Shell Terminal task

Use this task to acquire a command-line session.

To open a restricted shell terminal, open the *Open Restricted Shell Terminal* task from the HMC Management work pane. The Restricted Shell window displays. From the Restricted Shell window, you can issue commands remotely through Secure Shell access to the managed system. This process provides consistent results and automates administration of managed systems.

Remote HMC access: If you access the HMC remotely, this task is not shown in the work pane.

5.2.5 Shut Down or Restart task

Use this task to shut down (turn off the console) or to restart the console. To shut down the console, ensure that Shutdown HMC is selected, then click **OK** to proceed with the shutdown. To restart the console, ensure that Restart HMC is selected, then click **OK** to proceed with the shutdown. You can choose options, as shown in Figure 5-38.

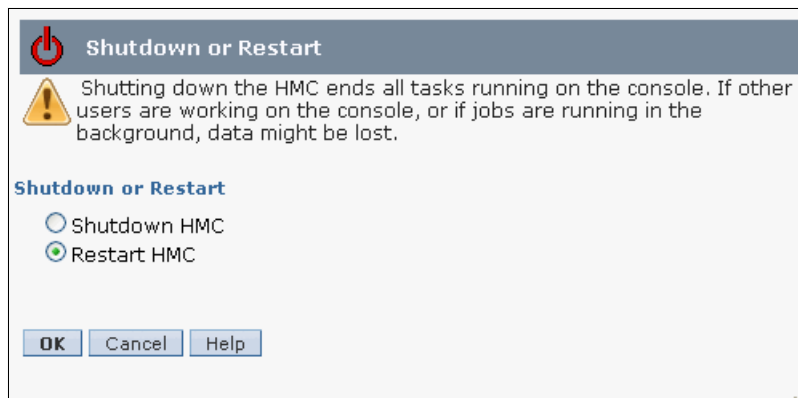


Figure 5-38 Shut down or restart task

5.2.6 Schedule Operations task

Use this task to create a schedule for certain operations to be performed on the logical partition without operator assistance. 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. For example, you can schedule a backup of important HMC information to DVD to occur once, or set up a repeating schedule.

You can choose the Backup Critical Console Data task.

We explain how to create and check tasks in “Schedule Operations task” on page 192.

5.2.7 Format Media task

Use this task to format a DVD-RAM, diskette, or USB 2.0 Flash Drive Memory Key for data backup. We explain this task in 6.1.1, “Management tasks” on page 270.

Remote formatting of media: You cannot do this task remotely unless you already placed the media in the system.

5.2.8 Back up HMC Data task

Use this task to back up (or archive) the data that is stored on your HMC hard disk that is critical to support HMC operations. Back up the HMC data after changes are made to the HMC or information that is associated with logical partitions. We explain this task in 6.2.1, “HMC Data backup” on page 309.

Backs up only the critical data: This task backs up only the critical data that is associated with HMC.

5.2.9 Restore HMC Data task

Use this task to restore critical backup data for this HMC. You can choose to restore data from a Network File System (NFS) server, a File Transfer Protocol (FTP) server, an SSH File Transfer Protocol (SFTP) server or removable media. We explain this task in 6.2.2, “Restoring HMC Data” on page 314.

5.2.10 Save Upgrade Data task

Use this task to use a wizard to save all of the customizable data for the HMC to the hard disk drive or to a DVD before performing an HMC software upgrade. We explain this task in 6.2.3, “HMC software maintenance” on page 315.

5.2.11 Change Network Settings task

Use this task to view the current network information for the HMC and to change network settings. We explain this task in 4.1.2, “Configuring the HMC network setting” on page 98.

5.2.12 Test Network Connectivity task

Use this task to view network diagnostic information for the console’s TCP/IP connection. This task also allows you to send a ping to a remote host. We explain this task in 4.1.3, “Testing network connectivity” on page 108.

5.2.13 View Network Topology task

Use this task to see a tree view of the network nodes that are known to this HMC. Examples of such nodes are managed systems, logical partitions, storage, and other HMCs. When you run this task, you can see the progress window. We explain this task in 4.1.4, “Viewing Network Topology” on page 110.

5.2.14 Tip of the Day task

Use this task to view information about using the HMC. When you enable this feature, a different fact or tip is displayed each time you log in, as shown in Figure 5-39 on page 236.

The Tip of the Day window opens if the **Show tips each time you log on** option is selected. If you prefer not to have this window display each time you log on, you can clear this option, and then click **Close**.

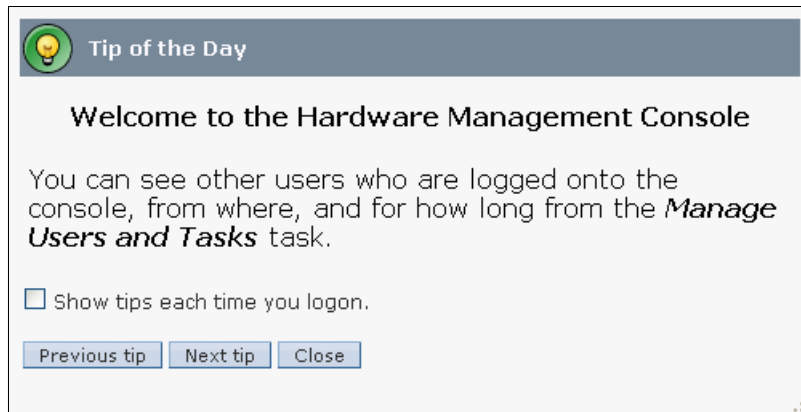


Figure 5-39 Sample of Tip of the Day window

5.2.15 View License task

Use this task to view the Licensed Internal Code (LIC) that you agreed to for this HMC. You can see the window has the link of *Third Party License Agreement* and *Additional License Agreement*, as shown in Figure 5-40.



Figure 5-40 View license task

5.2.16 Change Default User Interface Settings task

Use this task to customize settings that control how the HMC interface is displayed. This task customizes the default applying “all HMC user”. You can display or hide certain user interface components and icons, display or hide specific navigation nodes, and determine whether to save user interface settings changes, as shown in Figure 5-41. You can also restore to factory default settings by clicking **Factory Defaults**.

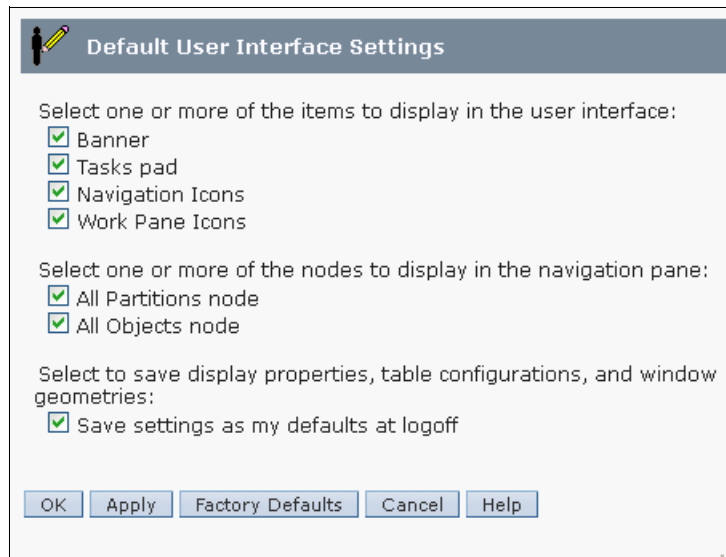
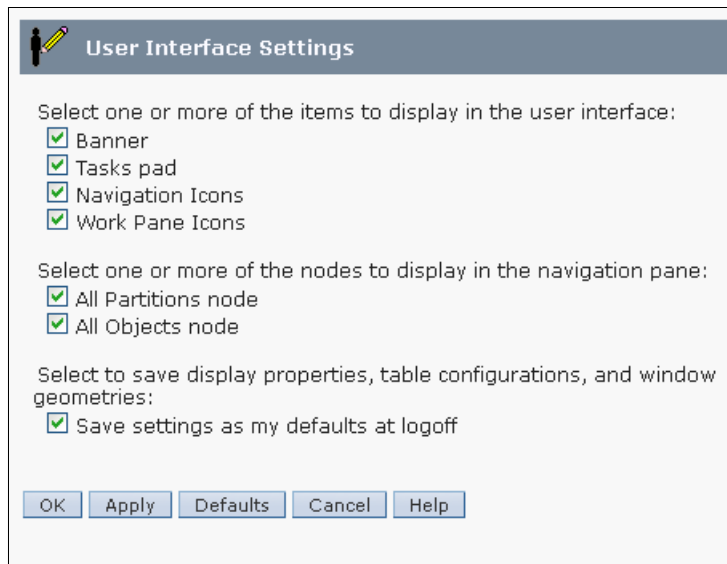


Figure 5-41 Change Default User Interface Settings task

5.2.17 Change User Interface Settings task

Use this task to customize settings that control how the HMC interface is displayed. You can display or hide certain user interface components and icons, display or hide specific navigation nodes, and determine whether to save user interface settings changes, as shown in Figure 5-42. Optionally, you can also open the task by clicking the logged on user name link that is displayed in the task bar below the banner.

User interface changes: User interface changes apply to the currently logged on user ID only. If you want to change the default applying the entire HMC user, see 5.2.16, “Change Default User Interface Settings task” on page 237.



The screenshot shows a dialog box titled "User Interface Settings" with a pencil icon. It contains three sections of settings:

- Select one or more of the items to display in the user interface:**
 - Banner
 - Tasks pad
 - Navigation Icons
 - Work Pane Icons
- Select one or more of the nodes to display in the navigation pane:**
 - All Partitions node
 - All Objects node
- Select to save display properties, table configurations, and window geometries:**
 - Save settings as my defaults at logoff

At the bottom, there are five buttons: OK, Apply, Defaults, Cancel, and Help.

Figure 5-42 Change User Interface Settings task

5.2.18 Change Date and Time task

Use this task to change the time and date of the battery-operated HMC clock, as shown in Figure 5-43. You can also add or remove time servers for the Network Time Protocol (NTP) service, as shown in Figure 5-44 on page 240. The time setting adjusts automatically for daylight savings time in the time zone you select.

You generally use this task in the following situations:

- ▶ The battery is replaced in the HMC.
- ▶ Your system is physically moved to a different time zone.

Change Date and Time

Customize Console Date and Time | NTP Configuration

Battery Operated Hardware Management Console Clock

Use this panel to verify or change the date and time on the Hardware Management Console clock. Changing this setting does not change the date/time settings on the managed system or logical partitions managed by this console.

Select the part of the date or time field that you want to change, or use the custom controls to adjust the date or time. The time setting will automatically be adjusted for daylight savings time in the timezone that you select if the clock is set to UTC.

Clock: UTC

Time: 10:00:00 AM

Date: 5/2/13

Time zone: America/New_York

Refresh

OK Cancel Help

Figure 5-43 Customize Console Date and Time task

Figure 5-44 shows the NTP Configuration tab where you can add or remove time servers.

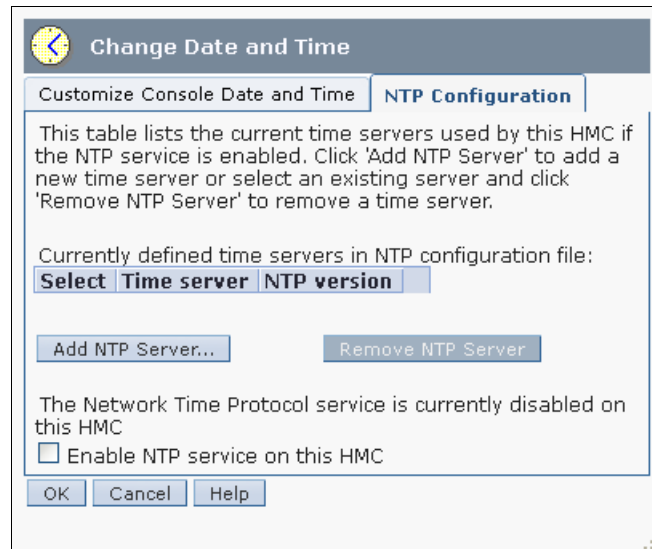


Figure 5-44 NTP Configuration task

5.2.19 Launch Guided Setup wizard

This wizard helps you set up your new system and the HMC. To set up your system and HMC successfully, complete all the tasks in the order that the wizard presents them. After you complete this wizard, you can use the properties for an object to make changes. We explain this task in 3.1.2, “HMC Guided Setup wizard” on page 64.

5.2.20 Launch Remote Hardware Management Console task

Use this task to start a session to another HMC. The remote HMC is an HMC that is on a different subnet from the service processor. Therefore, the service processor cannot be automatically discovered with IP multicast. If you want to connect to the remote HMC, you need the TCP/IP address or host name of the remote HMC.

Remote HMC access: If you access the HMC remotely, this task is not shown in the work pane.

5.2.21 Change User Password task

Use this task to change the HMC access password. After the password is changed, it must be changed on all other systems that access this HMC. We explain this task in 4.2.2, “Changing the user password” on page 114.

5.2.22 Manage User Profiles and Access task

Use this task to manage system users who log on to the HMC. We explain this task in 4.2.1, “Managing user profiles and access” on page 114.

5.2.23 Manage Task and Resource Roles task

Use this task to define and customize managed resource roles and task roles. We explain this task in 4.2.3, “Customizing user task roles and managed resource roles” on page 115.

5.2.24 Manage Users and Tasks task

Use this task to display the list of users that are currently logged on and the list of all tasks that run in this system. You can see information about users currently logged in and running tasks, as shown in Figure 5-45.

The following is the list of users currently logged on. The table below lists all tasks running in the system.

Users Logged On

Select	Session Id	User Name	Logon Time	Running Tasks	Access Location	Notes
<input type="checkbox"/>	38	hscroot	11/8/12 5:18 PM	2	172.16.254.14	Disconnected at 11/8/12 5:21 PM. The user ran the Disconnect task.
<input type="checkbox"/>	39	hscroot	11/8/12 5:21 PM	1	172.16.254.14	This is your session
<input type="checkbox"/>	40	hscroot	11/8/12 5:23 PM	1	172.16.254.14	

Logoff Disconnect

Running Tasks

Select	Task Id	Task Name	Targets	Session Id	Start Time
<input type="radio"/>	339	Change User Password		38	11/8/12 5:20 PM
<input type="radio"/>	340	Properties	lp1	38	11/8/12 5:21 PM
<input type="radio"/>	344	Manage Users and Tasks		39	11/8/12 5:21 PM
<input type="radio"/>	347	View HMC Events		40	11/8/12 5:24 PM

Switch To Terminate

Close Help

Figure 5-45 User and Tasks task

You can also do following tasks:

- ▶ Logoff
You can log the user off the HMC. This task is enabled only when you are assigned a user ID with Access Administrator roles.
- ▶ Disconnect
You can disconnect the user from the HMC. This task is enabled only when you are assigned a user ID with Access Administrator roles.
- ▶ Switch to
You can switch to another task that is running in your session.

- ▶ Terminate

You can end a task that is running in your session. If you are assigned a user ID with Access Administrator roles, you can also end tasks that are in other sessions.

5.2.25 Manage Certificates task

Use this task to manage the certificates that are used on the HMC. It provides the capability of getting information about the certificates that are used on the console. This task allows you to create a certificate for the console, change the property values of the certificate, and work with existing and archived certificates or signing certificates. We explain this task in 4.3, “Certificate Management” on page 118.

5.2.26 Configure Key Distribution Center task

Use this task to configure the key distribution center (KDC) servers that are used by this HMC for Kerberos remote authentication. Kerberos is a network authentication protocol that is designed to provide strong authentication for client/server applications by using secret-key cryptography. Optionally, you can import a service-key file into the HMC. The service-key file contains the host principal that identifies the HMC to the KDC server. Service-key files are also known as *keytabs*.

From this task you can do the following functions:

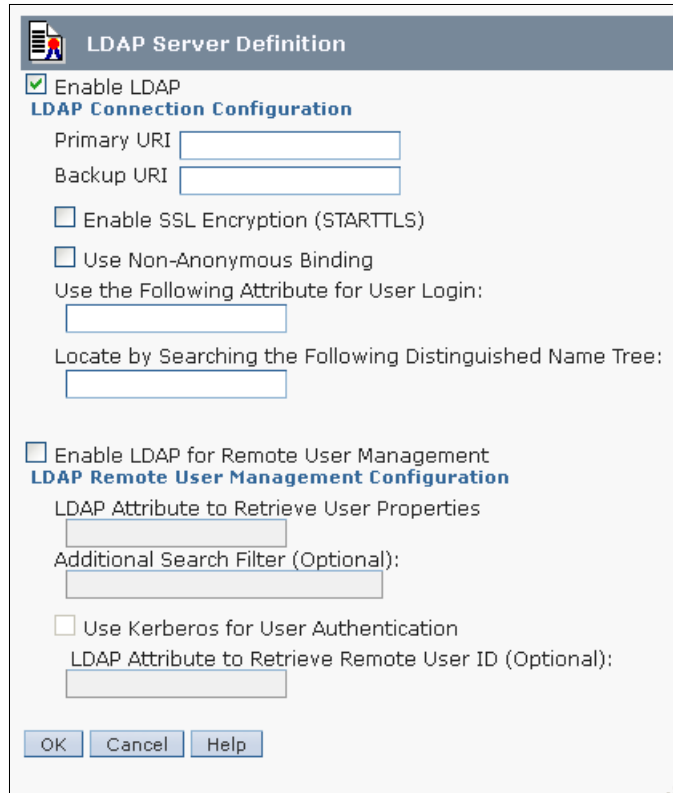
- ▶ View KDC server
- ▶ Modify KDC server
- ▶ Add KDC server
- ▶ Remove KDC server
- ▶ Import service key
- ▶ Remove service key

KDC servers:

- ▶ You must enable the NTP service on the HMC and set the HMC and the KDC servers to synchronize time with the same NTP server.
- ▶ You must set the user profile of each remote user to use Kerberos remote authentication instead of local authentication. A user that is set to use Kerberos remote authentication always uses Kerberos remote authentication, even when the user logs on to the HMC locally. We explain the *managing user profile* task in 4.2.1, “Managing user profiles and access” on page 114.

5.2.27 Configure LDAP task

Use this task to configure Lightweight Directory Access Protocol (LDAP) servers that are used by this HMC for LDAP authentication. You can use this task to enable LDAP authentication on this HMC to view LDAP servers that are used by this HMC for LDAP remote authentication to add LDAP servers. Or, use this task to remove LDAP servers from this HMC, as shown in Figure 5-46.



The screenshot shows a dialog box titled "LDAP Server Definition". It contains the following elements:

- Enable LDAP
- LDAP Connection Configuration**
- Primary URI
- Backup URI
- Enable SSL Encryption (STARTTLS)
- Use Non-Anonymous Binding
- Use the Following Attribute for User Login:
- Locate by Searching the Following Distinguished Name Tree:
- Enable LDAP for Remote User Management
- LDAP Remote User Management Configuration**
- LDAP Attribute to Retrieve User Properties
- Additional Search Filter (Optional):
- Use Kerberos for User Authentication
- LDAP Attribute to Retrieve Remote User ID (Optional):
- Buttons: OK, Cancel, Help

Figure 5-46 Configure LDAP task

LDAP remote authentication: You must set the user profile of each remote user to use LDAP remote authentication instead of local authentication. A user that is set to use LDAP remote authentication always uses LDAP remote authentication, even when the user logs on to the HMC locally. We explain the *managing user profile* task in 4.2.1, "Managing user profiles and access" on page 114.

5.2.28 Remote Command Execution task

Use this task to enable remote command execution by using the SSH facility. If you want to enable command-line access to the HMC through SSH, you must check **Enable remote command execution using the ssh facility**, as shown in Figure 5-47.

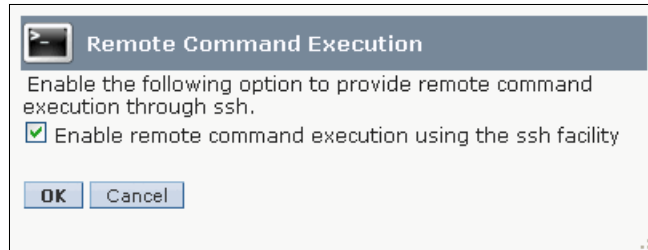


Figure 5-47 Remote Command Execution task

5.2.29 Remote Virtual Terminal task

Use this task to enable *remote virtual terminal* access for remote clients. A remote virtual terminal connection is a terminal connection to a logical partition from another remote HMC. If you want to enable remote virtual terminal access, you must check **Enable remote virtual terminal connections**, as shown in Figure 5-48.

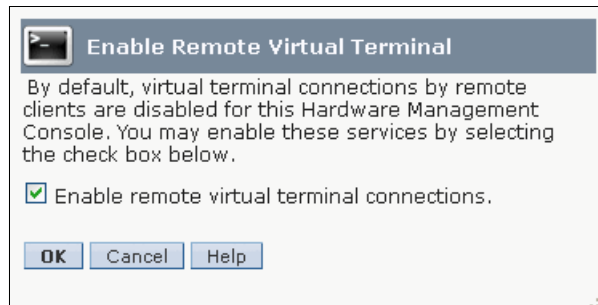


Figure 5-48 Remote Virtual Terminal task

5.2.30 Remote Operation task

Use this task to control whether the HMC can be operated by using a web browser from a remote workstation, as shown in Figure 5-49. By default, remote browser access to the HMC is disabled.




Figure 5-49 Remote Operation task

Remote HMC access: If you access the HMC remotely, you cannot change the status in this task.

5.2.31 Change Language and Locale task

Use this task to set the language and location for the HMC. After you select a language, you can select a locale that is associated with that language, as shown in Figure 5-50 on page 247. The language and locale settings determine the language, the character set, and other settings specific to the country or region (such as formats for date, time, numbers, and monetary units).

Figure 5-50 on page 247 shows the Change Language and Locale task.

Changes to language and location: Changes that are made in this task affect only the language and locale for the HMC itself. If you access the HMC remotely, the language and locale settings on your browser determine the settings that the browser uses to display the HMC interface.

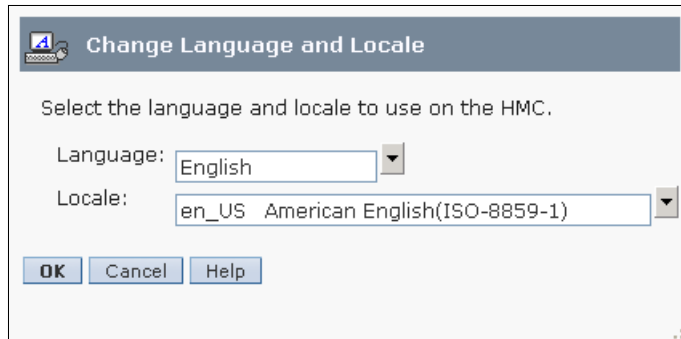


Figure 5-50 Change Language and Locale task

5.2.32 Create Welcome Text task

Use this task to customize the welcome message or to display a warning message that opens before users log on to the HMC. You can customize text and classification, which is shown in Figure 5-51.

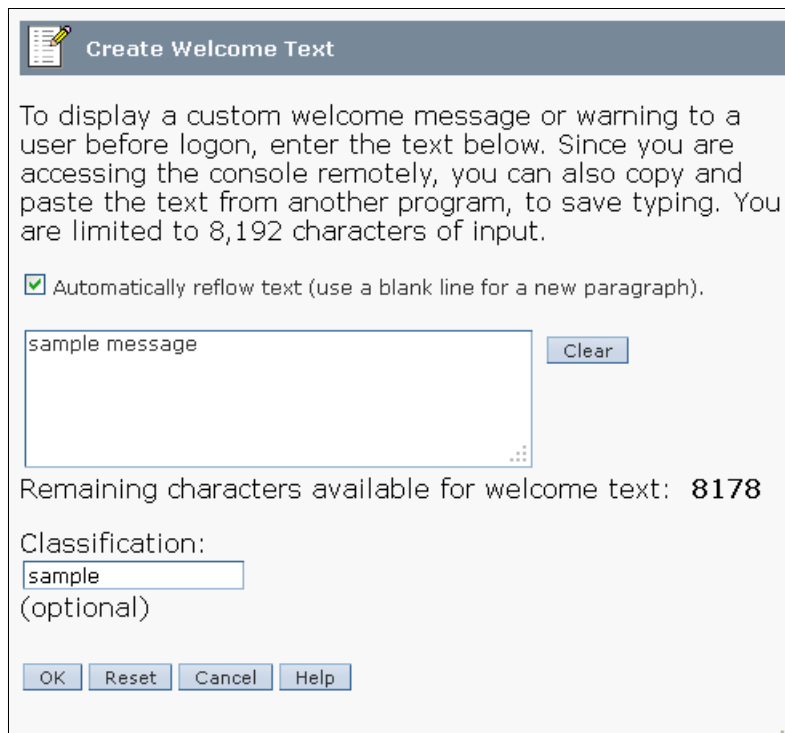


Figure 5-51 Create Welcome Text task

The text that you enter in the message input area for this task displays in the welcome window after you initially access the console, as shown in Figure 5-52.

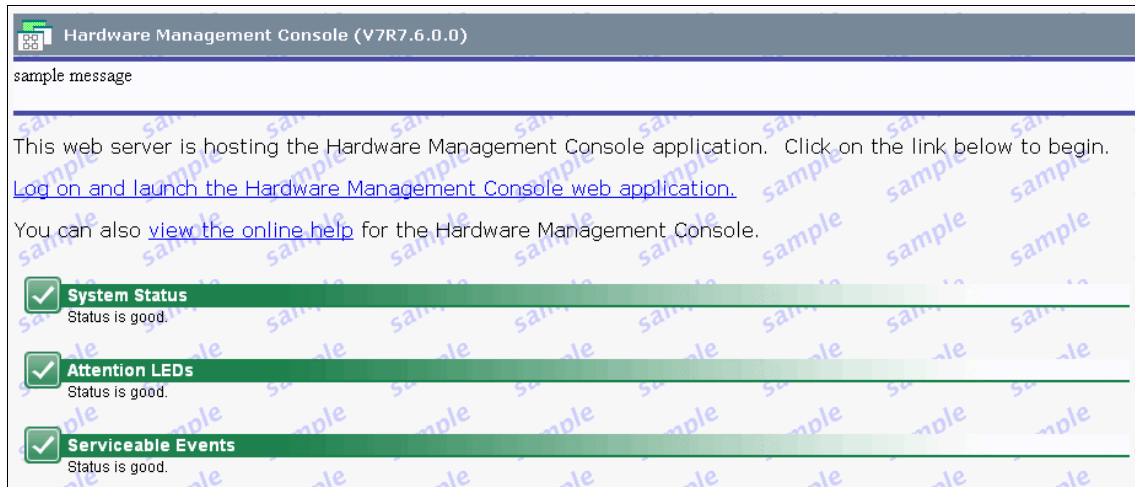


Figure 5-52 Sample of welcome text

5.2.33 Manage Data Replication task

Use this task to enable customized data replication. Customized data replication allows another HMC to obtain customized console data from or send data to this HMC. With the Customizable Data Replication service, you can configure a set of HMCs to replicate any changes automatically to certain types of data so that the configured set of HMCs keep this data synchronized without manual intervention.

Customizable console data:

- ▶ Customizable console data is accepted from other HMCs only after specific HMCs and their associated allowable customizable data types are configured.
- ▶ Before you enable this replication service, you might want to save your original data settings in case you have to restore these settings in the future.

You can configure the following types of data:

- ▶ User profile data
 - User identifications, methods of authentication, user managed resource roles and task roles, logon session properties, and remote access settings.

- ▶ Kerberos configuration data
 - Key Distribution Center (KDC), realm and host name that is used by the HMC for Kerberos authentication.
- ▶ LDAP configuration data
 - LDAP server name and distinguished name tree that is used by the HMC for LDAP authentication.
- ▶ Password policy configuration data
- ▶ Client information data
 - Client information for a server or group of servers that includes administrator, system, and account information about the system that is being installed.
- ▶ Group data
 - All user-defined groups that are defined to the HMC.
- ▶ Modem configuration data
 - Dial type (tone or pulse) and other settings such as whether to wait for a dial tone.
- ▶ Outbound connectivity data
 - Information for dialing out, such as whether to enable the local system as a call home server, or whether to allow dialing to use the local modem, the dial prefix, and phone numbers.

Figure 5-53 shows the types of customizable console data that can be replicated.

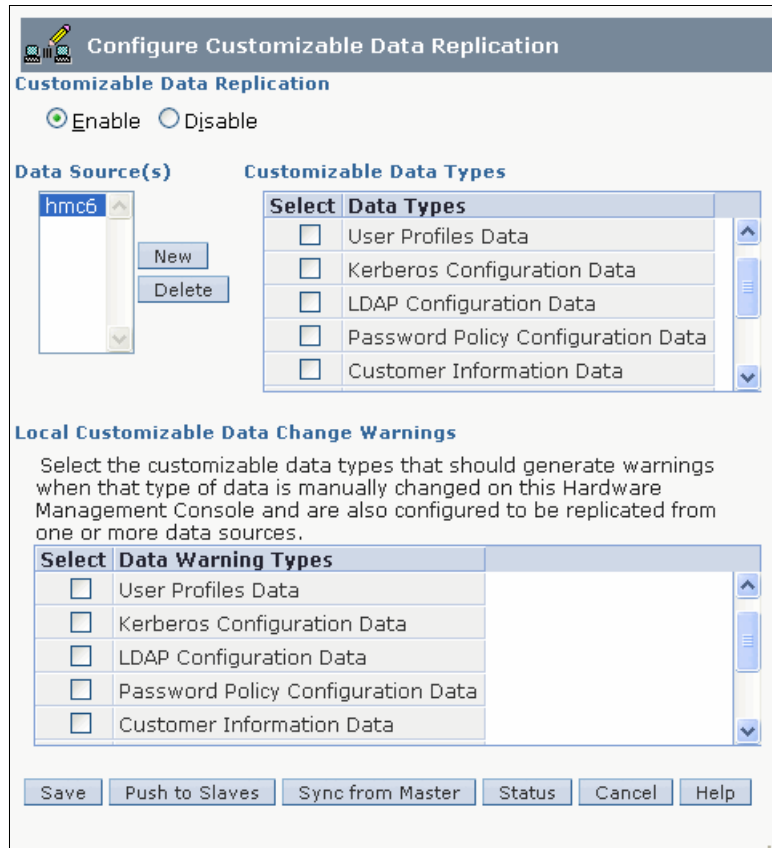


Figure 5-53 Manage Data Replication task

You can enable the Customizable Data Replication service for the following types of operations:

► **Peer-to-Peer replication**

Provides automatic replication of the selected customized data types between peer HMCs. Changes that are made on any of these consoles are replicated to the other consoles.

► **Master-to-Subordinate replication**

Provides automatic replication of the selected customized data types from one or more designated master HMCs to one or more designated subordinate HMCs. Changes that are made on a master console are automatically replicated to the subordinate console.

Configuring Peer-to-Peer replication

Figure 5-54 illustrates the process to configure Peer-to-Peer replication.

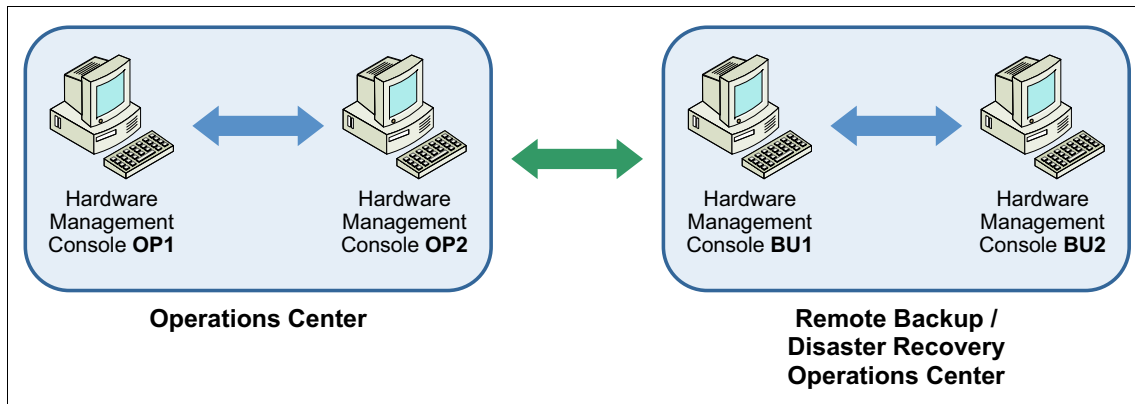


Figure 5-54 Peer-to-Peer replication

To configure Peer-to-Peer replication, follow these steps:

1. Open the Manage Data Replication task.
2. Select **Enable** in the Configure Data Replication panel.
3. The Configure Customizable Data Replication window opens.
4. Click **New** under Data Source. The Configure New Replication Source window opens.
5. Select an HMC to be used as a data source from the Discovered Console Information list, and click **Add**.

Alternatively, you can enter the TCP/IP address of the HMC that you want to use as a data source in the TCP/IP Address Information field, and then click **Find**.
6. The Customizable Data Replication window opens again.
7. Select the types of data that you want to replicate from the Customizable Data Types list, from a peer HMC that is selected currently under Data Source.
8. Click **Save** to close the Customizable Data Replication window.
9. Repeat steps 1 through 8 on each of the HMCs that you want to act as peers with one another.

HMC communication: When communication is established between the HMCs, the requested types of customizable data are replicated automatically from one HMC to the other immediately following the change in the data itself.

Configuring Master-to-Subordinate replication

Figure 5-55 illustrates the process to configure Master-to-Subordinate replication.

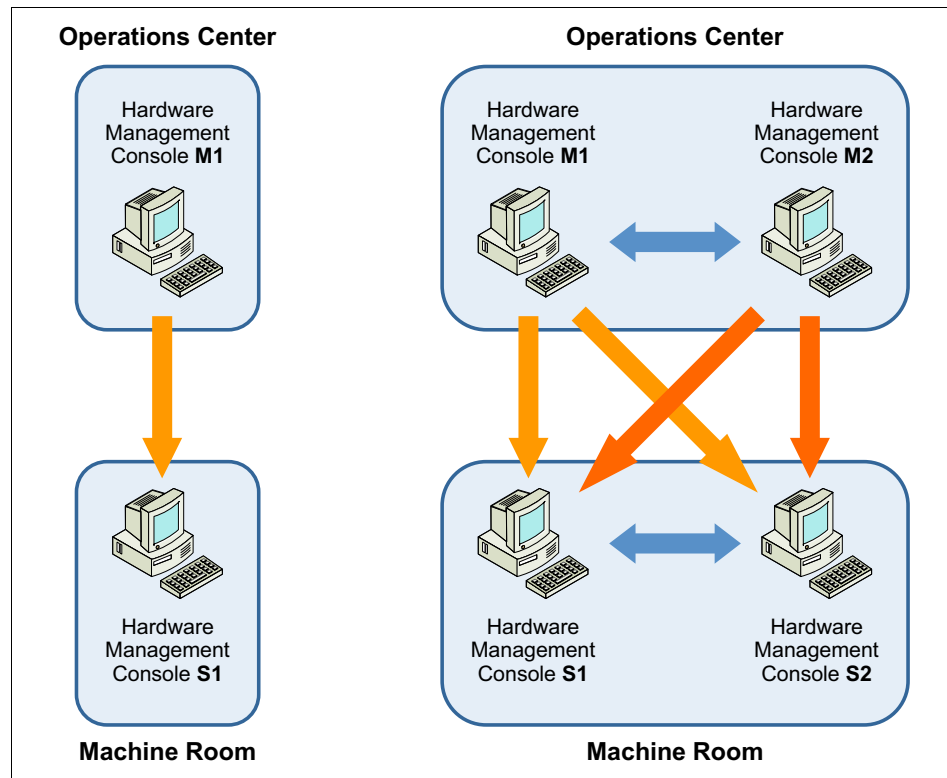


Figure 5-55 Master-to-Subordinate replication

To configure master-to-subordinate replication involves two steps:

1. Setting up a master console.
2. Setting up the subordinate console.

Setting up a master console

To set up a master console:

1. Open the Manage Data Replication task. The Customizable Data Replication window opens.
2. Select **Enable** in the Configure Data Replication panel.
3. Click **Save** to close the Customizable Data Replication window.

Configuring more master consoles: If you want to configure more master consoles, see “Configuring Peer-to-Peer replication” on page 251.

Setting up the subordinate console

To set up the subordinate console:

1. Open the Manage Data Replication task.
2. Select **Enable** in the Configure Data Replication panel.
3. The Customizable Data Replication window opens.
4. Click **New** under Data Source. The Configure New Replication Source window opens.
5. Select an HMC to be used as a data source from the Discovered Console Information list, and click **Add**.

Alternatively, you can enter the TCP/IP address of the HMC that you want to use as a data source in the TCP/IP Address Information field, and then click **Find**.
6. The Customizable Data Replication window opens again.
7. Select the types of data that you want to replicate from the Customizable Data Types list, from a peer HMC selected currently under Data Source.
8. Click **Save** to close the Customizable Data Replication window.
9. Repeat steps 1 through 8 on each of the HMCs that you want to act as peers with one another.

HMC communication: When open communication is established between the HMCs, the requested types of customizable data are replicated automatically from one HMC to the other immediately following the change in the data itself.

5.2.34 Manage Install Resources task

Use this task to add or remove operating environment installation resources for your HMC. You can use the HMC to deploy a system plan that contains information for installing one or more operating environments on one or more logical partitions. To install an operating environment as part of deploying a system plan, the HMC must be able to access and to use an installation resource for that operating environment.

You can create or define an installation resource for the following types of options (See Figure 5-56):

► **Create local installation resource**

You can create and define an installation resource in an HMC local hard disk drive from installation images.

► **Define remote installation resource**

You can define an installation resource from a remote NIM server.

Add Install Resource

Select one of the following options to add a new install resource

Create local install resource

Operating environment name: AIX

Operating environment version: * 7.1

Source is located on CD/DVD

Define remote install resource

Operating environment name: AIX

Operating environment version: * 5.3

NIM Master host name: *

NIM resource group: *

OK Cancel Help

Figure 5-56 Manage Install Resources task

Each installation resource that you define and create for the HMC is available for selection in the Deploy System Plan wizard. We explain System planning tools in 2.1, “System planning tools” on page 14.

5.3 Remotely access the logical partitions from HMC

The HMC provides virtual terminal emulation for AIX, Linux, and Virtual I/O Server logical partitions and virtual 5250 console emulation for IBM i logical partitions. This section describes how to use the virtual terminal.

5.3.1 Remote AIX, Linux, and Virtual I/O Server Terminal

AIX, Linux, and Virtual I/O Server require a console for installation and some operations. The HMC virtual terminal provides virtual terminal console access to every logical partition without a physical device assigned, as shown in Figure 5-57. One virtual terminal is available for each logical partition of the managed system.




Figure 5-57 Virtual terminal window

To open the virtual terminal:

1. Select the check box for the logical partition you want to open the terminal window in the Work pane.
2. Select **Console Window** → **Open Terminal Window**.
3. The virtual terminal window opens.

Logical partitions support only one terminal connection at a time. Therefore, if a partition reserves the terminal connection, the HMC should not attempt to open a terminal connection to the same partition. If the HMC attempts to open a reserved connection, the server firmware returns an error, as shown in Figure 5-58.

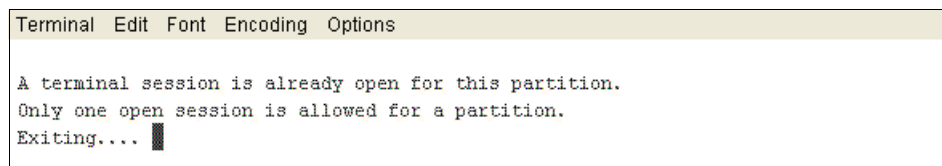


Figure 5-58 Virtual terminal error

5.3.2 Remote IBM 5250 terminal

The IBM i console can be accessed remotely by using an IBM 5250 terminal emulator, from the complete installation of IBM i Navigator V6.1 or later for Windows Operating System. You can start the IBM 5250 emulator by going to **Program Files** → **IBM i System Navigator** → **Emulator** → **Start or Configure Session** and start a *New Session*, as shown in Figure 5-59.

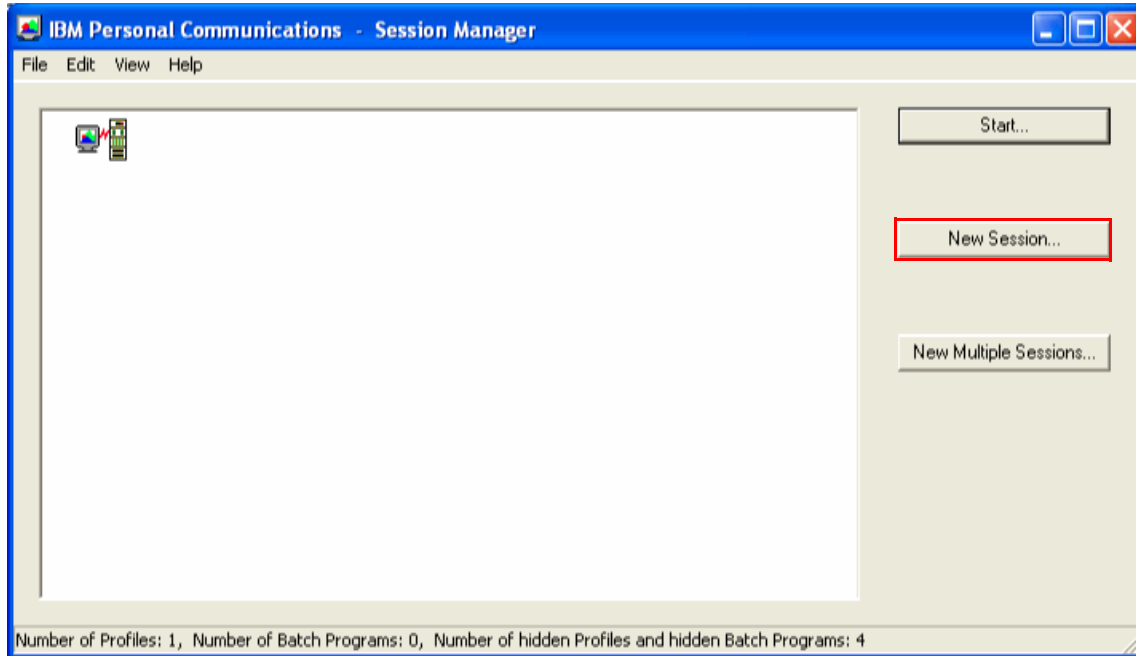


Figure 5-59 Start a new session from IBM Personal Communications menu

In the new window that is displayed, fill the HMC IP address, change the port from 23 to 2300, then finally click Properties. See Figure 5-60.

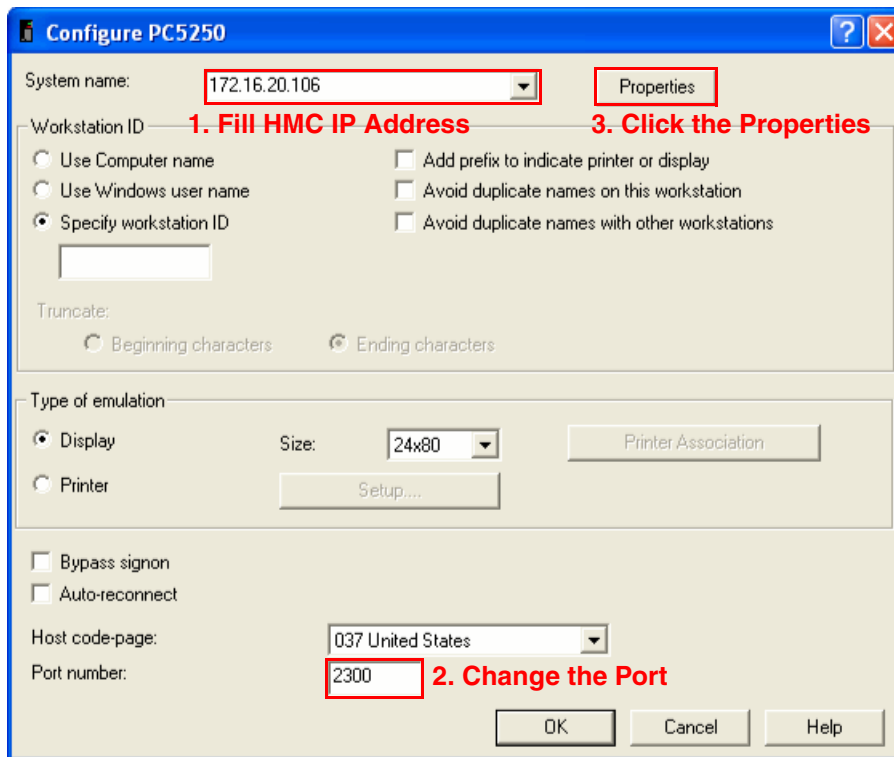


Figure 5-60 Configure IBM 5250 session terminal emulator

Next, the connection window is displayed, and you have to change the *User ID signon information* option to *Use default User ID, prompt as needed*, then use name, as shown in Figure 5-61.

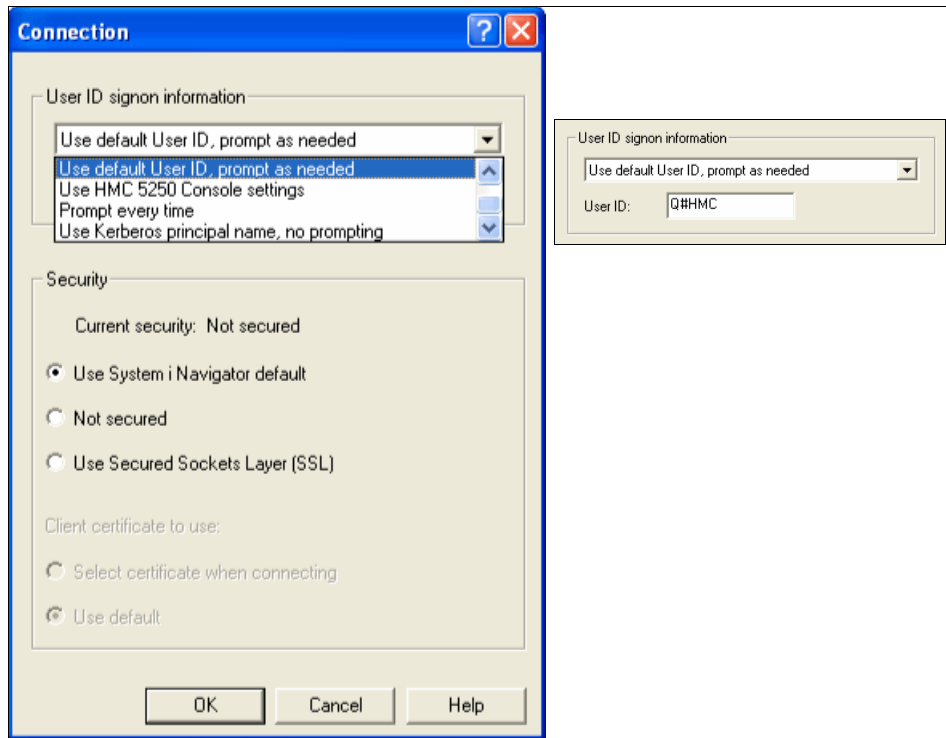


Figure 5-61 Configure User ID signon information window

Press **OK** to complete the user ID signon information, then **OK**, once again to return to the IBM 5250 emulator and choose connect from the *Communication* tab. If the connection is good, it brings you to the specify language option menu. Enter 21 for the American English language option. See Figure 5-62.

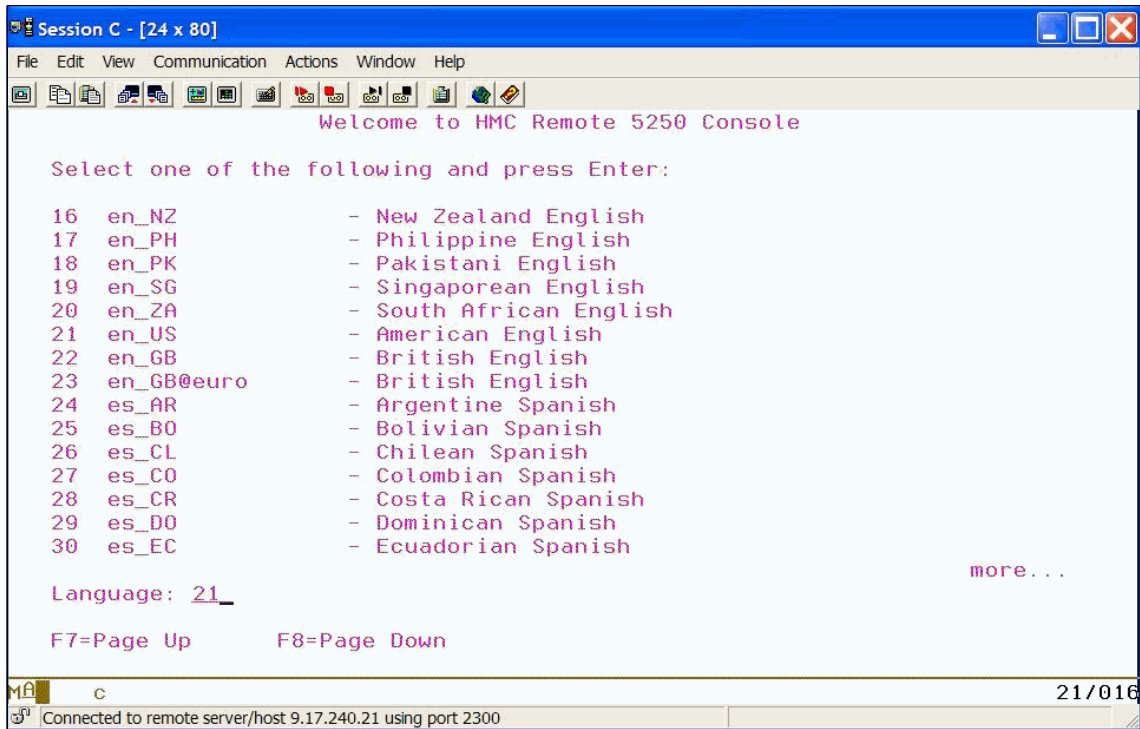


Figure 5-62 Enter option 21 for American English language option

The next window brings you to the managed system that is managed in the HMC. Choose the correct managed system, then continue, as shown in Figure 5-63.

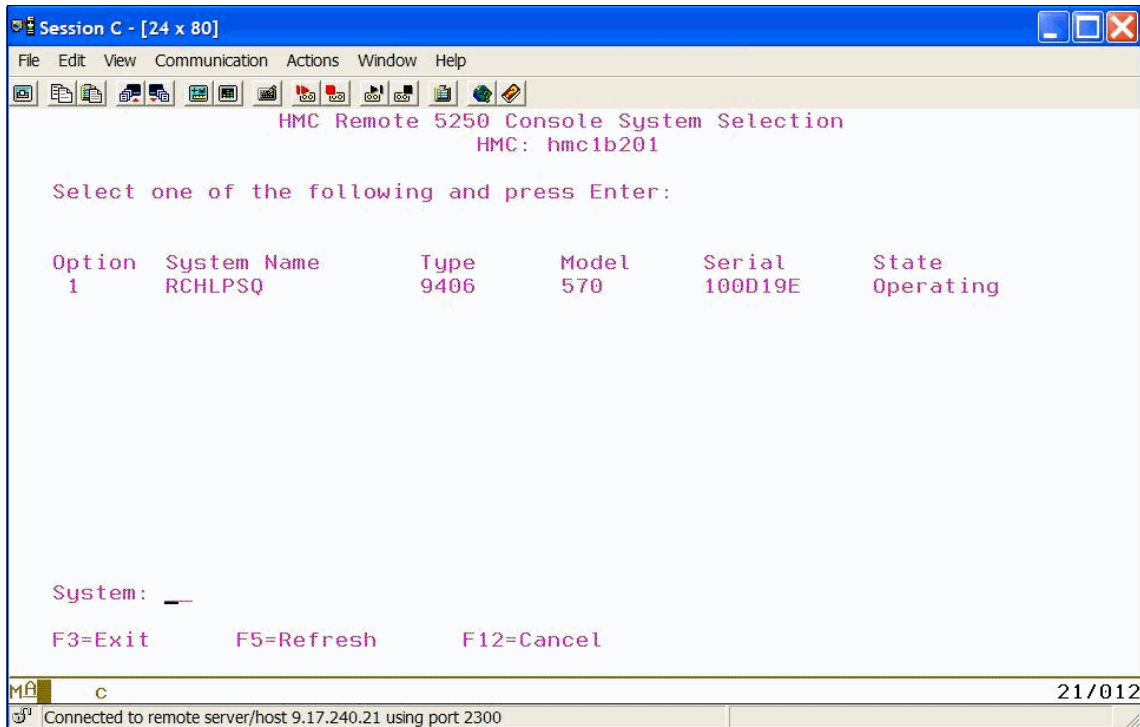


Figure 5-63 Select IBM Power Systems to connect

The next window gives two types of connection: shared connection and dedicated connection.

Dedicated connection The console is available for only one remote connection and is not shared to any connection, including the local 5250 connection from HMC.

Shared connection The console is available for several remote connections, including the local 5250 console connection from HMC. You must set up the shared key access for the first shared connection. The other connection to follow is to provide the same shared key when you created the first shared connection, and having the same window as the first connection.

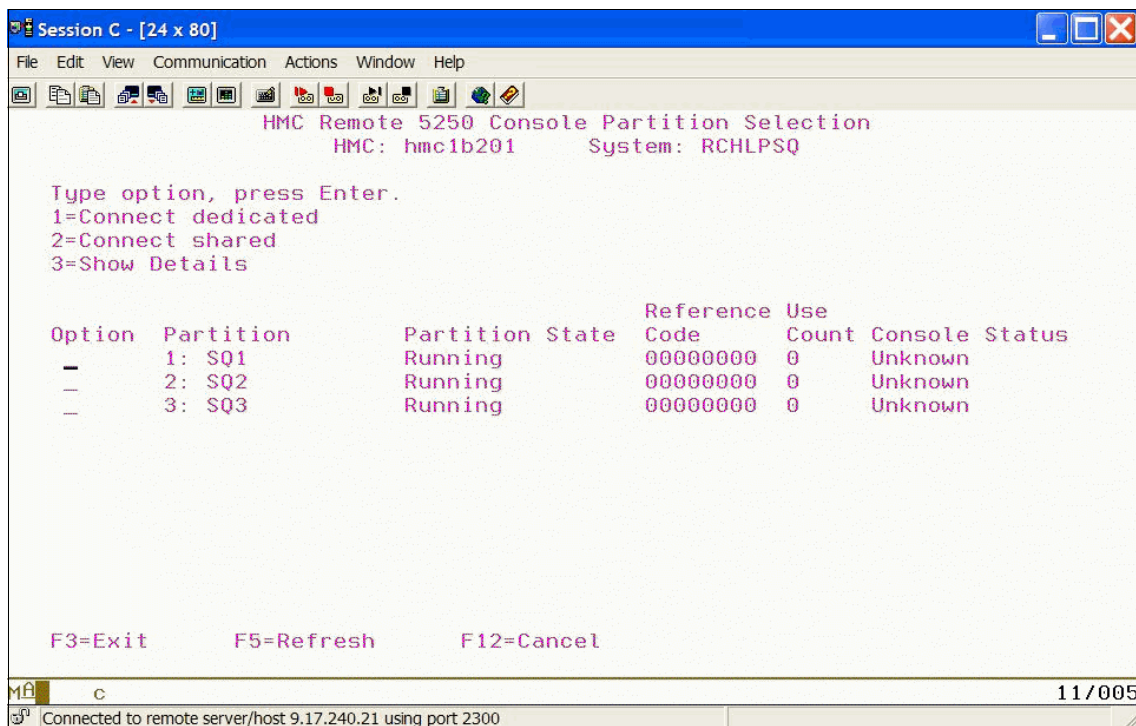


Figure 5-64 Type of connection to IBM i LPAR

Either type of the connection that you choose brings you to Dedicated Service Tools window or the IBM 5250 login panel.

5.4 Managing partition data

This section describes how to manipulate partition profile data on the HMC. It includes information about the following functions:

- ▶ **Restore:** Load profile data locally from the HMC or from removable media.
- ▶ **Initialize:** Initialize all profile images.

Initialize: *Initialize* removes *all* profile data that is saved to the HMC. Do not do this action if you want to preserve *any* of the profile data that is saved on the HMC.

- ▶ **Backup:** Save profile data to the HMC.
- ▶ **Delete:** Remove profile data.

5.4.1 Restore task

The *restore* option allows you to restore profile data from a local backup that is stored on the HMC. To use this option, select **Configuration** → **Manage Partition Data** → **Restore**. Select the backup image that you want to restore, and then select **OK**, as shown in Figure 5-65. The restore can take approximately a minute and a half or longer to complete.

Profile Data Restore: 8233-E8B-SN10DD51P

Select a profile backup file from which to restore the managed system's profile data. Then select a restore option.

Select	File Name	Backup Time
<input type="radio"/>	backupFile	October 31, 2012 11:01:13 AM EDT
<input checked="" type="radio"/>	test_1031	October 31, 2012 11:19:41 AM EDT
<input type="radio"/>	original_1025	October 25, 2012 11:47:45 AM EDT

Restore Options

- Full restore from the selected backup file
- Backup priority -- merge current profile and backup
- Managed system priority -- merge current profile and backup

OK **Cancel** **Help**

Figure 5-65 Profile Data Restore task

You have the following options for restore:

▶ **Full restore from the selected backup file**

The full restore option restores all profile data by using your backup file only. Profile modifications that are done after the selected backup file was created are lost. If you want to preserve any of the existing partition profile data that is on the managed server during the restore process, consider using *Backup priority* or *Managed system priority* instead of a *Full restore from the selected backup file*.

▶ **Backup priority: merge current profile and backup**

The backup priority option merges the stored backup file with recent profile activity, but the backup information takes precedence. If information conflicts, the stored backup data is restored over the recent profile activity.

An example of when this option is useful is if you have a partition where you changed the memory, processors, or adapters, but the previous configuration of the partition performed better. By using this option, you can restore the partition to its previous state from a backup. This process holds true for all partitions included in the backup data.

▶ **Managed system priority: merge current profile and backup**

The managed system priority option merges recent profile activity with the stored backup file, but the recent profile activity takes preference. If information conflicts, the recent profile activity is restored over the stored backup data. This option is useful to restore deleted partitions without affecting the other partitions on the system.

5.4.2 Initialize task

The *initialize* option clears all current profile data on the managed server and effectively removes all partitions from the hypervisor and system image on the HMC. (This option does not remove profile backups. To do that, see 5.4.4, “Delete task” on page 266.)

To use this option, select **Configuration** → **Manage Partition Data** → **Initialize**. If you are certain that you want to remove all the partitions on the managed server, select **Yes**, as shown in Figure 5-66.

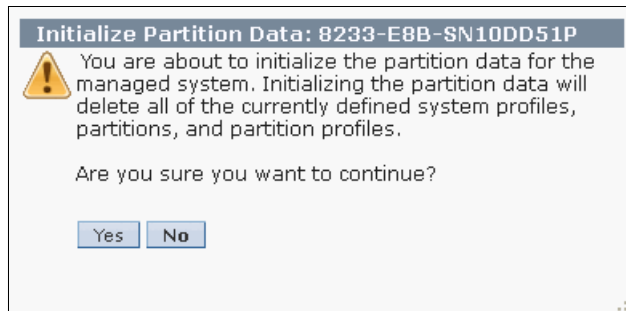


Figure 5-66 Initialize Profile Data task

Must be in Not Activated state to initialize: You cannot initialize profile data if any partition on the managed system is in the *Standby* or *Operating* state. All partitions on the managed system must be in the *Not Activated* state to initialize profile data.

Times vary on how long it takes to initialize profile data. To initialize data can take up to two minutes or longer.

When complete, your managed server is clear of any profile data. From this point, you can create new partitions on the server or restore partition profile data from a previous backup. Read 5.4.1, “Restore task” on page 262 on how to restore profile data.

5.4.3 Backup task

The *backup* option allows for the backup of profile data for all partitions on a managed server to be saved locally on the HMC. To use this option, select **Configuration** → **Manage Partition Data** → **Backup**. Enter a name for the backup, and then select **OK** to continue, as shown in Figure 5-67.

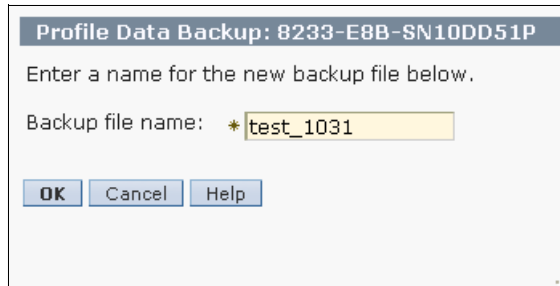


Figure 5-67 Profile Data Backup task

The backup can take a minute or longer, depending upon how many partitions are on the system.

When the backup is complete, a message is displayed (Figure 5-68).




Figure 5-68 Backup profile data results window

5.4.4 Delete task

The *delete* option allows you to delete a single backup image of the managed server partition profile data. To use this option, select **Configuration** → **Manage Partition Data** → **Delete**. Select the partition profile image that you want to remove from the list and select **OK**, as shown in Figure 5-69.

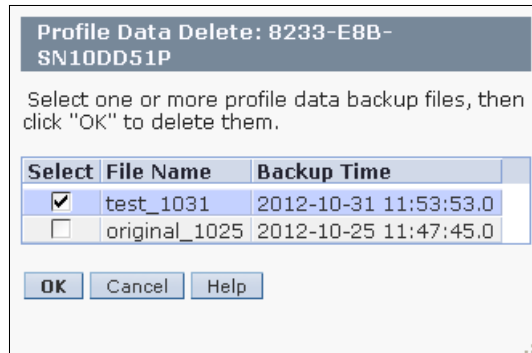


Figure 5-69 Profile Data Delete task

A status window opens after you select **OK**. The HMC returns you to the main partition view when the deletion is complete.



Service support

This chapter describes the service support function on the HMC, including service management, HMC software maintenance, and Advanced System Management Interface (ASMI).

6.1 Service Management

The Service Management functionality can be found in the left pane of the HMC home page.

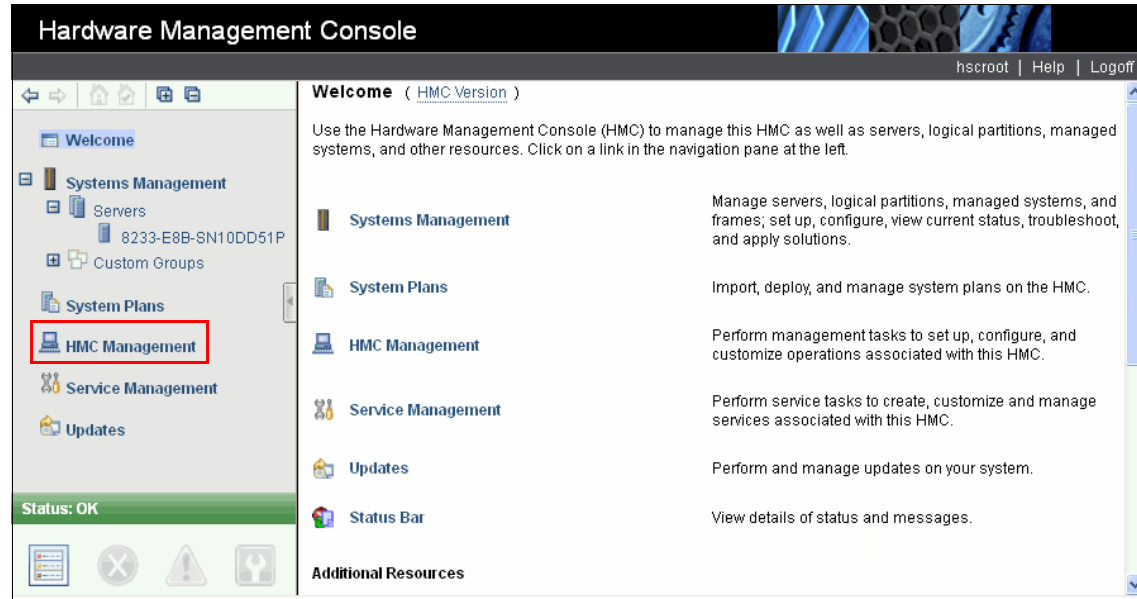


Figure 6-1 Access Service Management from the HMC main menu

The main view of the Service Management area is divided into two sections, as shown in Figure 6-2.

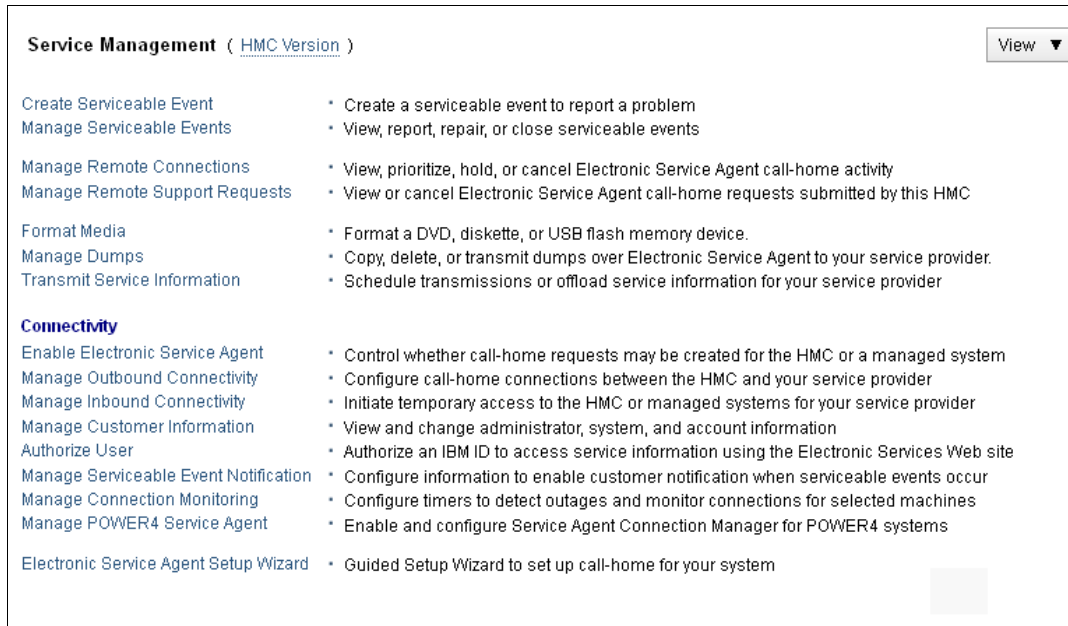


Figure 6-2 Service Management: main view

The first Service Management section describes the *management tasks* that can be performed on the HMC. The following list provides HMC management tasks:

- ▶ Handling service events
- ▶ Managing remote connections
- ▶ Formatting removable media
- ▶ Managing service dumps
- ▶ Transmitting service data

The second Service Management section covers *connectivity* with the HMC. The following list provides HMC connectivity tasks:

- ▶ Enabling electronic service agent
- ▶ Handling outbound connectivity
- ▶ Permitting inbound connectivity
- ▶ Specifying client information
- ▶ Registering eService authorized user
- ▶ Setting contact information for serviceable events
- ▶ Managing connection monitoring
- ▶ Handling previous service agent connectivity
- ▶ Guided setup wizard to set up *call home* for your system

6.1.1 Management tasks

The options that can be performed in the top half of the service management area mostly pertain to service events, formatting and using removable media, and sending in service reports to IBM. There are also troubleshooting tools available in this area of the HMC that pertain to troubleshooting not only with managed servers but also with the HMC itself.

Service events

There are three options for service events on the HMC:

- ▶ Create event
- ▶ Manage event
- ▶ Load event

Create event

Select **Create Serviceable Event** to manually report a hardware failure on a managed server or on the HMC itself (as shown in Figure 6-3).

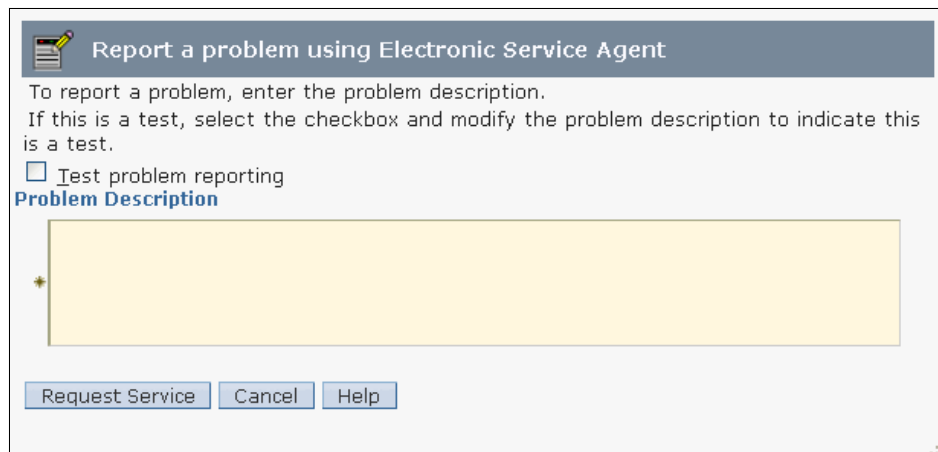


Figure 6-3 Service Management: creating a serviceable event

Under Problem Description, provide as much information as possible about the problem that you encountered, including the hardware that is involved and references to any error logs or reports that are associated with the event.

When completed, select **Request Service**. If your connectivity to IBM is set up properly as described in “Manage Outbound Connectivity task” on page 290, the error report is sent to IBM.

Manage events

Select **Manage Serviceable Events** to open the window that is shown in Figure 6-4.

Manage Serviceable Events

Use this window to specify selection criteria for the serviceable events you wish to view or manage. Only events that meet all the criteria that you specify will be displayed.

Event criteria

Serviceable event status: * Open

Problem number: * ALL

Error criteria

Reporting MTMS: * ALL

Failing MTMS: * ALL

Reference code: * ALL

Number of days to view: * ALL

Field-Replaceable Unit (FRU) criteria

Part number: * ALL

Location code: * ALL

OK Cancel Help

Figure 6-4 Service Management: manage serviceable events

Use the pull-down menus to alter the filters to display the current service events that are recorded on the HMC. When your selections are complete, select **OK**. Then, you have a results window similar to Figure 6-5.

Manage Serviceable Events - Serviceable Event Overview

Selected ▾

This list shows all serviceable events that match your selection criteria. Each event is grouped with all errors that are associated with that event. Use the menu bar above to perform actions on the serviceable event.

Compact table view Full table view

--- Select Action --- ▾

Select ^	Problem # ^	PMH # ^	Reference code ^	Status ^	Last reported time ^	Failing MTMS ^
<input type="checkbox"/>	6		11001510	Open	Oct 26, 2012 9:56:10 AM	8233-E8B/10DD51P
Total: 1 Filtered: 1 Selected: 0						

Figure 6-5 Service Management: serviceable event overview

The codes that are contained in the *Reference code* column are live links. You can select these codes to get further information pertaining to the service event in question. As shown in figure Figure 6-6, by selecting the reference code *11001510*, a window is displayed that shows more information about this service event.

[Subscribe to this information](#)

POWER7 information

11001510

11001510

Explanation

Detected AC loss

Response

Before replacing any parts, verify that the AC input voltage is correct. Follow the instructions for the items listed in the FRU List.

Problem determination

No additional problem determination.

Failing Item

- [PWRSPPLY](#)
- [TWRCARD](#)

Figure 6-6 Service Management; event description

Similarly, as shown in Figure 6-7, you can select a service event, and then click **Selected** to choose one of the following menu options:

- ▶ **View Details:** Get reference details about the service event.
- ▶ **Repair:** Connect to ResourceLink to attempt to fix the problem that is associated with the service event.
- ▶ **Call Home:** Manually report the service event to IBM.
- ▶ **Manage Problem Data:** Send specific files that are associated with a service event to IBM or off load data to removable storage.
- ▶ **Close Event:** Remove the service event from the list.



Figure 6-7 Service Management: reporting options

By having a service event that is selected and clicking **View Details**, a window opens, as shown in Figure 6-8.

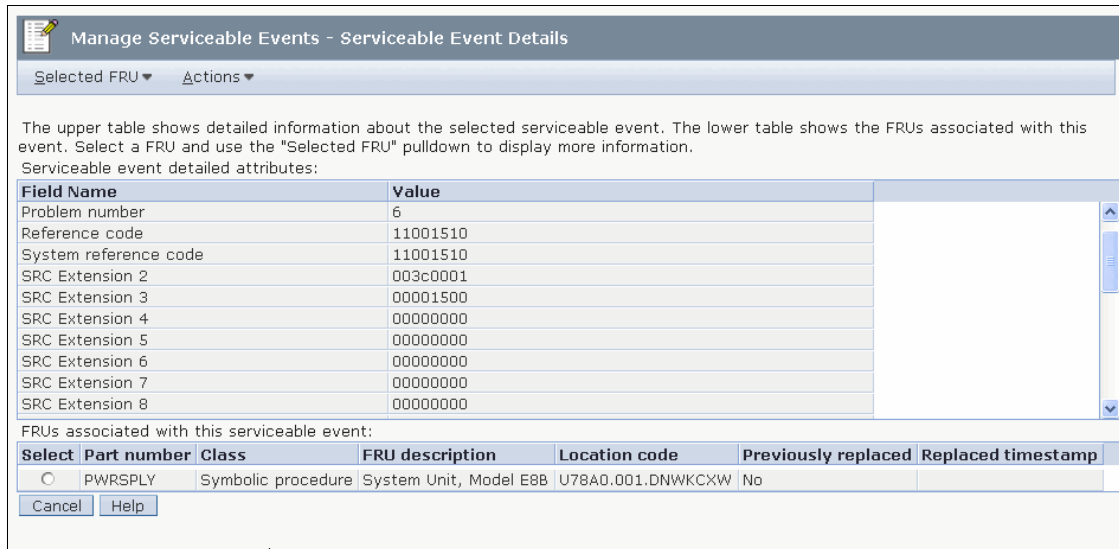


Figure 6-8 Service Management: View details about service event

In this view, you can see the service reference code (SRC) extensions that are associated with the service event that you selected. Also, you can take actions that are associated with the hardware by clicking **Selected FRU** or **Actions**.

You can return to the Serviceable Event Overview panel by closing the window.

You can select a service event and then select **Repair** to begin repair actions that are associated with the service event, as shown in Figure 6-9.

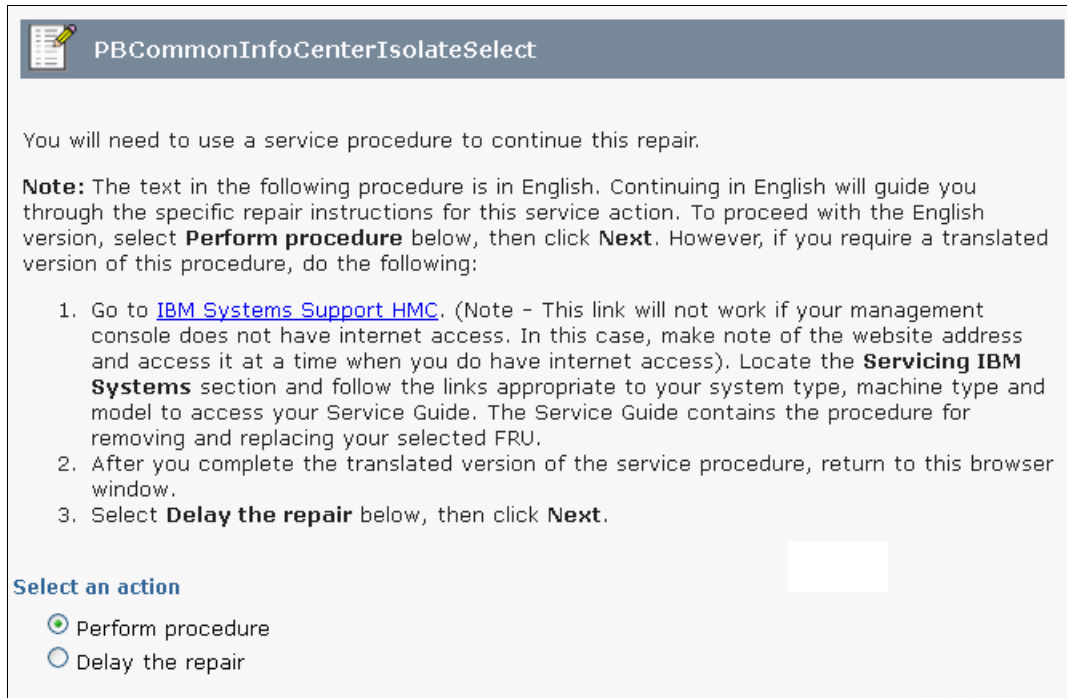


Figure 6-9 Service Management, repair action on service event

If you select **Perform Procedure**, you can get step by step guidance in resolving the service event or you can delay the repair. If you select **Delay the repair**, you are returned to the Serviceable Event Overview window, as shown in Figure 6-5 on page 272.

When you select a service event and click **Manage Problem Data**, the window that is shown in Figure 6-10 opens.

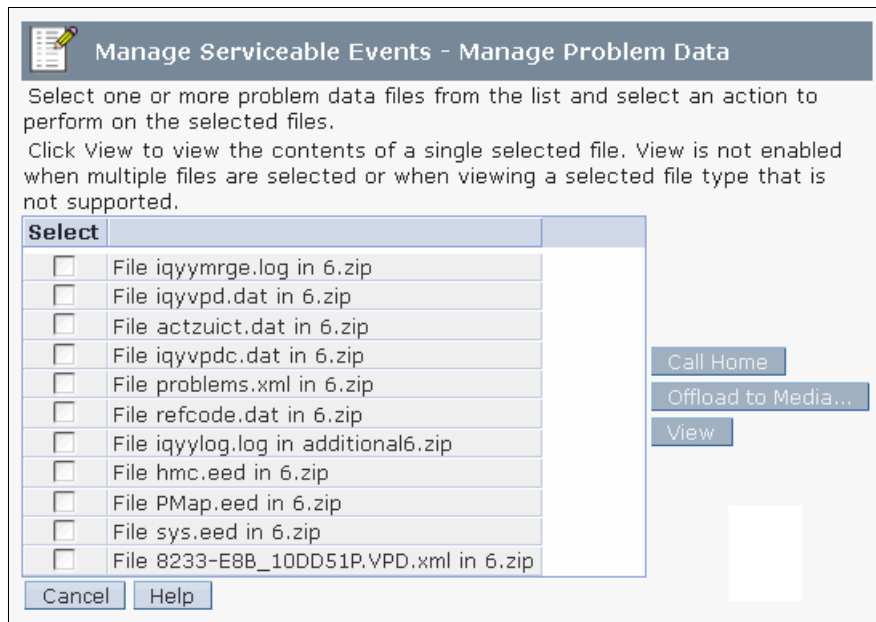


Figure 6-10 Service Management: manage problem data

You can select specific files that are associated with the service event with the following possible actions:

- ▶ Send the file to IBM by clicking **Call Home**.
- ▶ Save the associated files to removable media by clicking **Offload to Media**.
- ▶ Examine the contents of the associated file by clicking **View**.

Finally, if you select a service event and click **Close Events**, a window opens as shown in Figure 6-11.

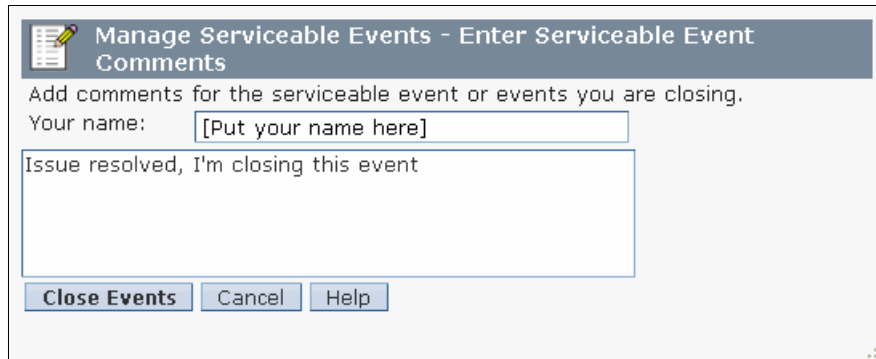


Figure 6-11 Service Management: close a service event

Here, you provide your name and the reason for closing the associated service event.

Close Events task: After you perform the Close Events task, all other options except Manage Problem Data are disabled for the service event in question on the Serviceable Event Overview window.

Remote access

The remote access options of the Service Management area of the HMC allow for these functions:

- ▶ Manage remote connections by manually configuring the serviceable event queue for transmission to IBM.
- ▶ Manage remote support requests by managing the local queue of serviceable events on your HMC.

Figure 6-12 shows the remote access options of the Service Management area.

Manage Remote Connections	• View, prioritize, hold, or cancel Electronic Service Agent call-home activity
Manage Remote Support Requests	• View or cancel Electronic Service Agent call-home requests submitted by this HMC

Figure 6-12 Service management: remote access options

Manage Remote Connections option

If the call home server service of the HMC is enabled, use the Manage Remote Connections option to manage the console's remote connections manually.

The console manages its remote connections automatically. It puts requests on a queue and processes them in the order in which they are received. However, you can use the Manage Remote Connections option to manage the queue manually.

Figure 6-13 shows the Manage Remote Connections window.

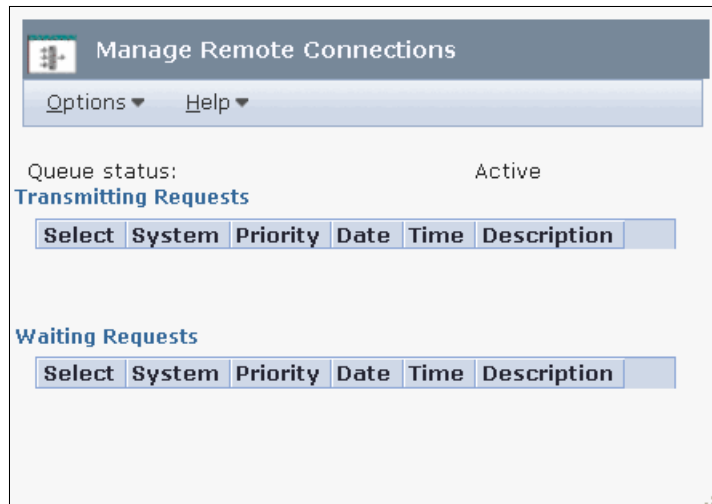


Figure 6-13 Manage Remote Connections window

Use the Options menu in this window to do the following functions:

- ▶ Stop transmissions
- ▶ Move priority requests ahead of others
- ▶ Delete requests

Manage Remote Support Requests option

You can use the Manage Remote Support Requests option to view or manage call home requests that are submitted by the HMC that are either being processed or that are waiting to be processed.

Figure 6-14 shows the Manage Remote Support Requests window.



Figure 6-14 Manage Remote Support Requests window

Use the Options menu in this window to do the following functions:

- ▶ View all call home servers
- ▶ Cancel selected requests
- ▶ Cancel all active requests
- ▶ Cancel all waiting requests

Managing HMC service data

The HMC service data area of the HMC allows you to do the following functions:

- ▶ Format media for use with various HMC functions.
- ▶ Manipulate managed server dump information.
- ▶ Schedule transmittal of vital product data (VPD), Capacity on Demand (CoD), and serviceable event information about the HMC.

Figure 6-15 shows the HMC data options area of the Service Management area.

Format Media	• Format a DVD, diskette, or USB flash memory device.
Manage Dumps	• Copy, delete, or transmit dumps over Electronic Service Agent to your service provider.
Transmit Service Information	• Schedule transmissions or offload service information for your service provider

Figure 6-15 Service Management: HMC data options

Format Media option

This option is to format your removable media, a DVD-RAM, or a USB flash memory device:

1. Select **Format Media**. The window that is shown in Figure 6-16 opens.

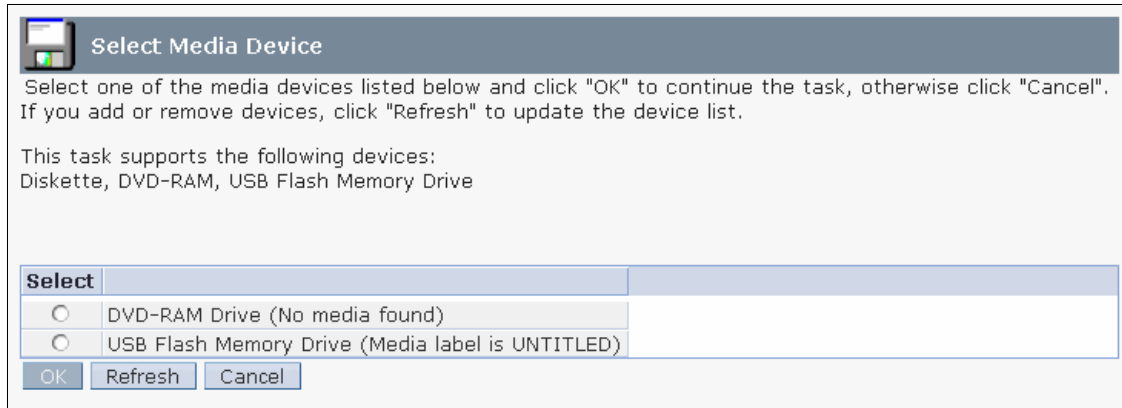


Figure 6-16 Service Management: Format Media

Tip: If the USB option is not displayed in the **Format Media** window, it might be not formatted correctly or not installed properly.

2. Select the media device that you want to format, and select **OK**.
3. Then, either insert or attach the appropriate media device to be formatted.

Manage Dumps option

Before you can use the Manage Dumps option, you need create a dump from the server view:

1. Select **Systems Management** → **Servers** and then select the name of the server. On the lower panel, select **Serviceability** → **Hardware** → **Manage Dumps** (Figure 6-17).




Figure 6-17 Service Management: Initiate dump from the server

2. In the Manage Dumps window, select **Action** → **Initiate System Dump**, as shown in Figure 6-18.

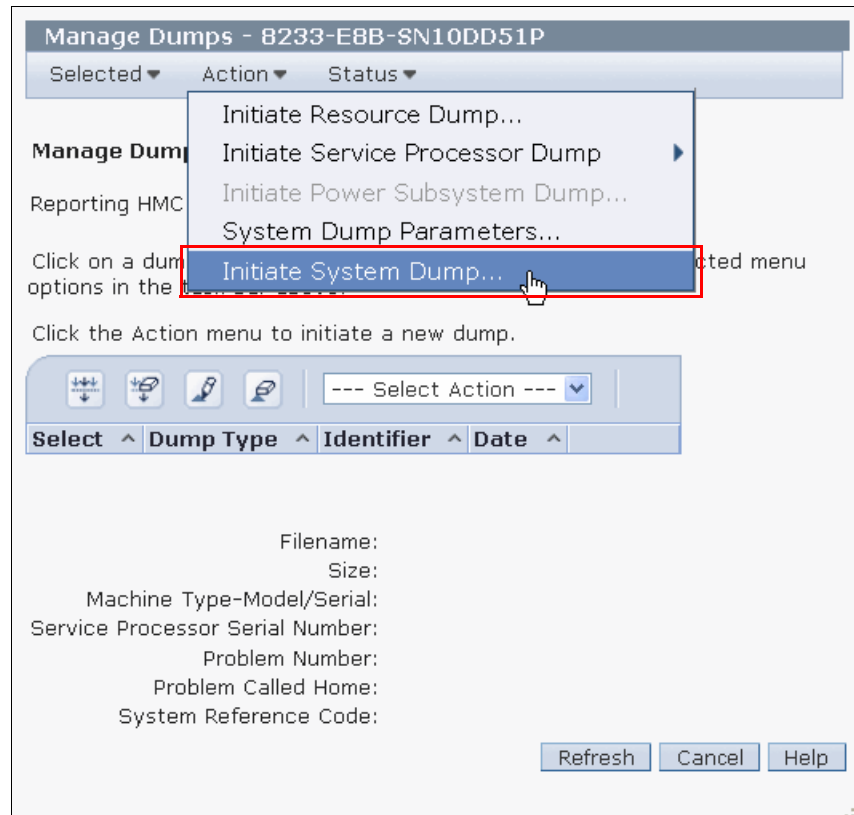


Figure 6-18 Service Management: Initiate system dump

Warning: Do a system dump only for the problem determination procedure, because it forces a shutdown on every logical partition (LPAR) in the selected managed system.

3. Specify the system where you want to initiate the dump (Figure 6-19). Verify the selected server, and then select **OK**.

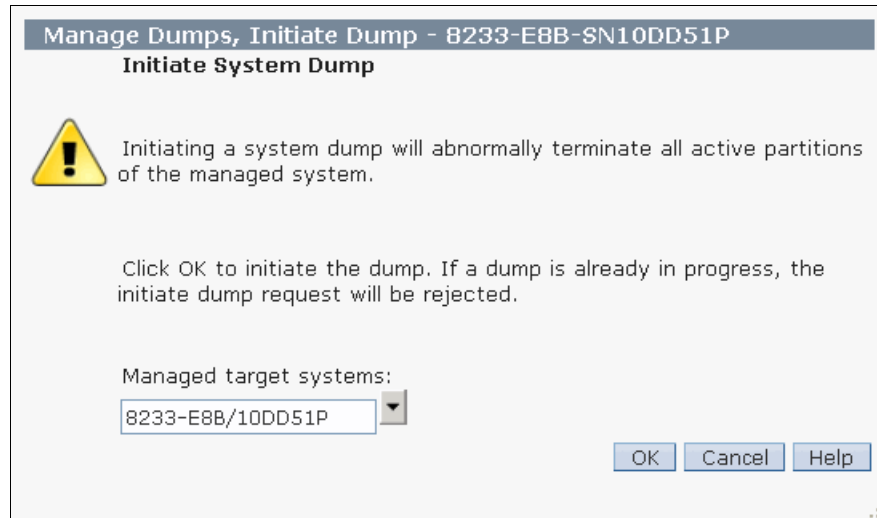


Figure 6-19 Service Management: Initiate system dump for managed system

4. If there are no errors that are associated with the system dump, you get a status message as shown in Figure 6-20. Click **OK**.

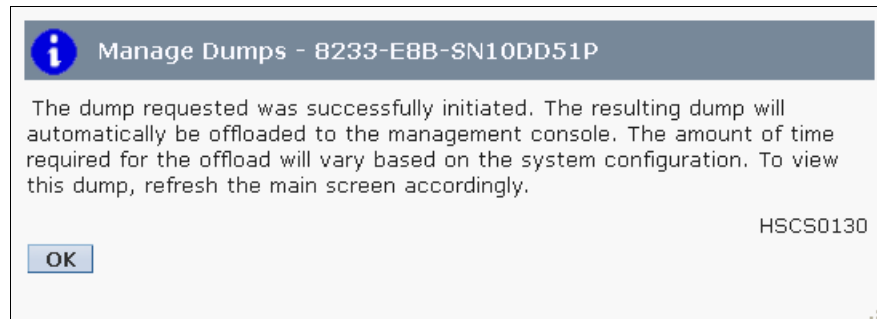


Figure 6-20 Service Management: System dump complete

5. Now that there is a dump available for manipulation in Service Management, select **Service Management** → **Manage Dumps**. The system dump displays in the results window, as shown in Figure 6-21 on page 285.

6. Select the dump that you want to manipulate. Then, click **Selected**. From this menu, you can do the following functions:
 - **Copy Dump to Media:** Move dump information from the HMC to removable media such as a DVD.
 - **Copy Dump to Remote System:** Specify a remote FTP server, ID, and password to which to transmit your system dump.
 - **Call Home Dump:** Sends your system dump data to IBM.
 - **Delete Dump:** Removes dump data from the HMC.

Manage Dumps - 8233-E8B-SN10DD51P

Selected ▾ Action ▾ Status ▾

- Copy Dump to Media...
- Copy Dump to Remote System...
- Call Home Dump...
- Delete Dump...

Click on a dump to select it, then choose a task from the Selected menu options in the task bar above.

Click the Action menu to initiate a new dump.

Select ^	Dump Type ^	Identifier ^	Date ^
<input checked="" type="radio"/>	SYSDUMP	00000001	Oct 26, 2012 6:19:45 PM
<input type="radio"/>	FSPDUMP	1A000000	Oct 26, 2012 6:20:41 PM

Filename: SYSDUMP.10DD51P.00000001.20121026181945.raw364268144.gz
 Size: 14.77 MB
 Machine Type-Model/Serial: 8233-E8B/10DD51P
 Service Processor Serial Number: Not supported
 Problem Number: Not assigned
 Problem Called Home: No
 System Reference Code: A1003000

Refresh Cancel Help

Figure 6-21 Service Management: Dump results

Transmit Service Information option

This area of Service Management is used to transfer service event information from the HMC to IBM. As shown in Figure 6-22, there are three tabs that are associated with service data transmission.

Transmit Service Information

Transmit FTP

You can transmit information to your service provider immediately or you can schedule the transmission.

Operational test (Heartbeat):

Interval (days): 7

Time: 1:25:41 AM

To run operational test immediately, click Run. Run

Hardware service information (VPD) transmission:

Enable

Interval (days): 7

Time: 2:05:42 AM

To transmit the hardware service information immediately, click Send. Send

Software service information transmission:

Enable

Interval (days): 7

Time: 2:54:02 AM

To transmit the software service information immediately, click Send. Send

To manage the software service information transmission by partition, click Manage. Manage

Performance management transmission:

Enable

Interval (days): 1

Time: 4:25:07 AM

To transmit the performance management information immediately, click Send. Send

OK Cancel Help

Figure 6-22 Service Management: Transmit

In the Transmit tab, you can do the following functions:

- ▶ Control the frequency of data transmissions to IBM
- ▶ Send service data immediately
- ▶ Separate schedules for service information and performance management data transmission

Schedules: Schedules of service information and performance management data might be different. You can also send one to IBM without having to send the other.

When you made your selections for how often the data is to be sent or when you made an immediate transfer, select **OK**.

The File Transfer Protocol (FTP) tab allows for control of where to send FTP data when it is sent (Figure 6-23). If the HMC is behind a firewall, the FTP tab allows you to specify the appropriate settings to transmit service data beyond the firewall.

Transmit Service Information

Transmit FTP

Provide configuration data to allow the use of FTP to offload service information.

FTP Server

Enable FTP offload of service information

Name: testcase.boulder.ibm.com Port: 21

Directory: /hardware/toibm/edddata/

User name: anonymous

Password:

Passive: on

If your network includes a company firewall, you will need to specify configuration information about the firewall in order for you to use an FTP site to offload service information.

FTP Firewall

Enable firewall configuration settings

Authentication format: 1 - USER user@real.host.name Port: 21

Host name:

User name:

Password:

Exclusion list:

Passive: on

FTP Test/Reset

To perform a test FTP with your FTP settings, click Test.

To reset all your FTP settings to their original default values, click Reset.

Figure 6-23 Service Management: FTP settings

At the bottom of the FTP tab are the options to test the FTP connection and to reset all of the FTP settings to their original defaults.

Test your settings: It is highly recommended that you test your FTP and firewall settings after set up to ensure that your data is getting to IBM. Testing is important if your system is enabled for a type of CoD that requires monthly reporting, such as that described in “On/Off Capacity on Demand” on page 164.

The **Send** button transmits related data to IBM (Figure 6-23 on page 288). You can transmit the following information:

- ▶ **Hardware Management Console log:** The full log of serviceable event data on the HMC.
- ▶ **Problem determination data:** Information that is specific to serviceable events logged by the HMC.
- ▶ **Managed systems VPD data collection:** The full system inventory of managed servers that are attached to the HMC.

6.1.2 Connectivity

To get to the connectivity options, select **Service Management** in the HMC workplace window. The following options are available:

- ▶ **Enable Electronic Service Agent:** The call home function allows the HMC to dial in to the IBM network through a modem or the Internet to report:
 - Serviceable events
 - CoD usage (On/Off, Reserve Capacity, and Utility Capacity)
 - Hardware failures
- ▶ **Outbound Connectivity:** This window allows for configuration of the HMC modem or for configuration of Ethernet connectivity to the outside Ethernet.
- ▶ **Inbound Connectivity:** Allow your service provider temporary access to your HMC or partitions of a managed system.
- ▶ **Customer Information:** Specify administrator, system, and account information.
- ▶ **Authorize User:** Register a client user ID with the eService website.
- ▶ **Serviceable Event Notification:** Define information to enable client notification when service events occur.
- ▶ **Connection Monitoring:** Configure timers to detect outages and monitor connections for selected computers.
- ▶ **POWER4 Service Agent:** Activate Service Agent Connection Manager for POWER4 systems.
- ▶ **Electronic Service Agent Setup wizard:** Allows you to set up the electronic service agent through a guided setup wizard.

Enable Electronic Service Agent option

To activate regular system status reporting through the call home feature, select **Service Management** → **Enable Electronic Service Agent**. Select any systems on which you want to affect the call home feature, and then select either **Enable** or **Disable** (Figure 6-24). Click **OK** to save your selections or click **Cancel** to negate your selections.

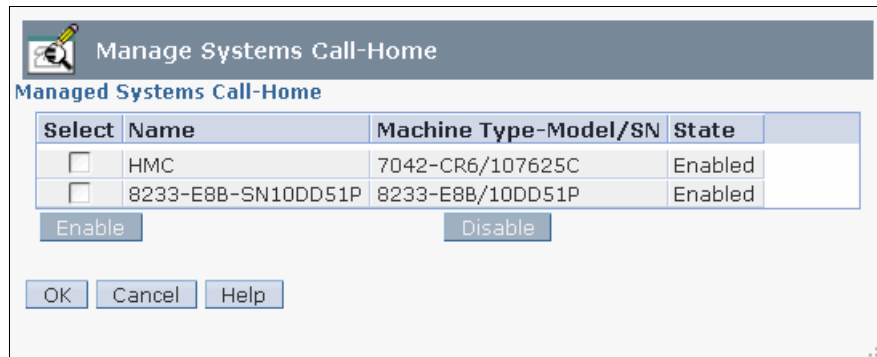


Figure 6-24 Service Management: Enable the call home feature

Along with enabling your HMC or managed servers for the call home feature, configure and test your outbound connectivity to ensure that your reports can get to IBM. Refer to “Manage Outbound Connectivity task” on page 290 for more information.

Manage Outbound Connectivity task

You can achieve outbound connectivity either through a modem on the HMC or through Internet connectivity. First, select **Service Management** → **Manage Outbound Connectivity**.

Modem configuration

To set your modem configuration:

1. On the Local Modem tab, select **Allow local modem dialing for service**, and then select **Modem Configuration** (Figure 6-25).

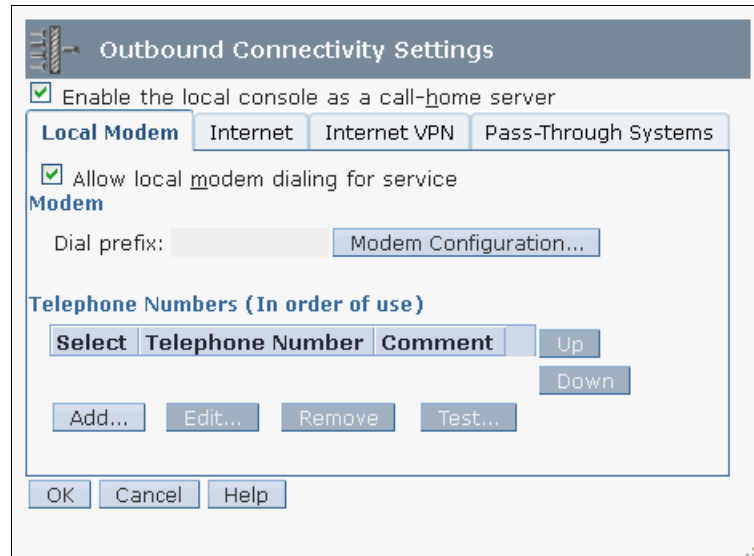


Figure 6-25 Manage Outbound Connectivity option: Local Modem tab

2. Set the modem to tone or pulse dialing and set a dial prefix, if required, as shown in Figure 6-26. Then, select **OK**.

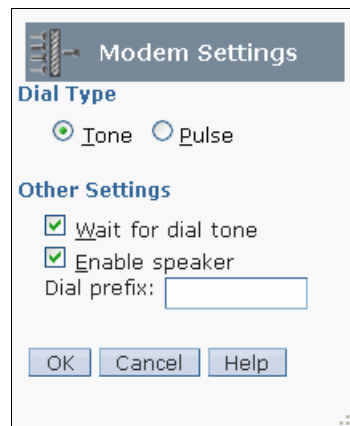
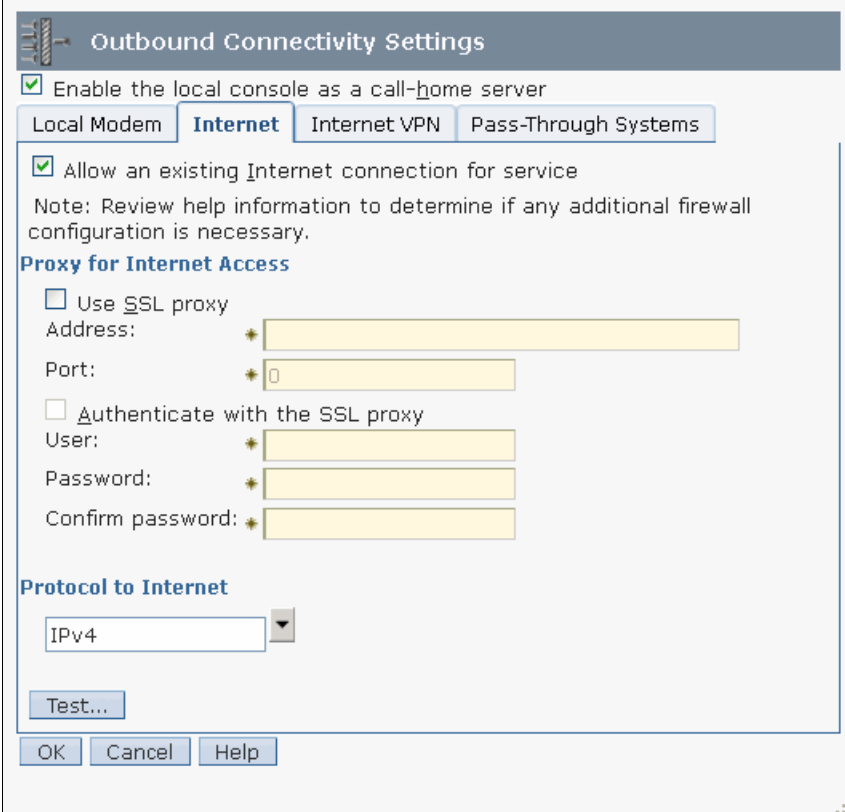


Figure 6-26 Modem configuration

Internet configuration

To configure your HMC for Internet configuration:

1. Select the Internet tab, as shown in Figure 6-27.
2. Select **Allow an existing Internet connection for service**. If your HMC is behind a firewall, you must select **Use SSL proxy** and provide the address and port for your proxy in the address and port fields.
3. When complete, you can test the configuration by clicking **Test**. If the test completes successfully, select **OK** to save your settings.



The screenshot shows the 'Outbound Connectivity Settings' dialog box with the 'Internet' tab selected. The 'Enable the local console as a call-home server' checkbox is checked. Below the tabs, the 'Allow an existing Internet connection for service' checkbox is checked, with a note about firewall configuration. The 'Proxy for Internet Access' section has the 'Use SSL proxy' checkbox checked, with fields for 'Address', 'Port', 'User', 'Password', and 'Confirm password'. The 'Authenticate with the SSL proxy' checkbox is unchecked. The 'Protocol to Internet' dropdown is set to 'IPv4'. A 'Test...' button is located below the proxy settings. At the bottom of the dialog are 'OK', 'Cancel', and 'Help' buttons.

Figure 6-27 Manage Outbound Connectivity option: Internet access

Internet VPN tab

Use the Internet VPN tab, which is shown in Figure 6-28 to allow the HMC to use a virtual private network (VPN) over an existing internet connection for service. This option is highly recommended if it is available within your IT infrastructure. By using a VPN for remote connectivity, you are encrypting the information as it is transmitted from your HMC to IBM. This process helps to ensure that your systems data is kept private and helps keep your HMC secure.

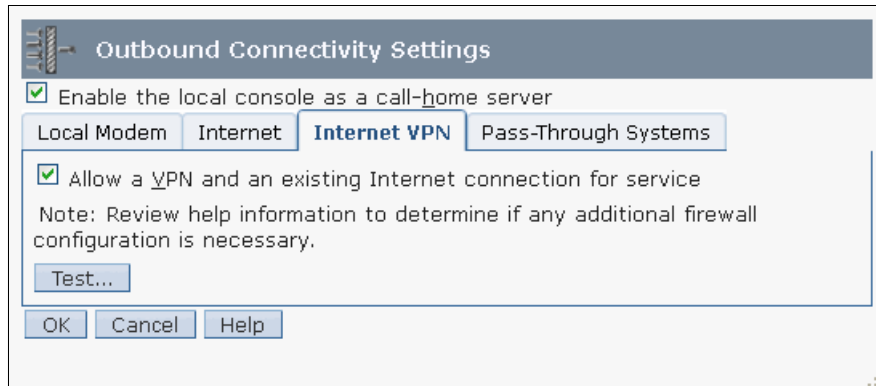


Figure 6-28 Manage Outbound Connectivity option: Internet VPN

Pass-Through Systems tab

Use the Pass-Through Systems tab to allow other managed servers, systems, and other devices on the same network as the HMC to use it as an access point to the external Internet. This configuration is useful if you set up a VPN connection as described in "Internet VPN tab". This tab is also useful when you want to allow all devices that receive a DHCP lease from the HMC to use the HMC as a point of access to the Internet.

It is recommended, though not required, to have separate network adapters on the HMC for a private (the HMC and all managed servers) and open (the HMC, servers, systems, and the external internet) network. By having separate networks for HMC management and Internet connectivity, you are helping secure the data that is transferred between the HMC and managed servers. For information about how to set up open and private networks on the HMC, read 4.1.2, "Configuring the HMC network setting" on page 98.

Using the Pass-Through Systems option allows for a single network that is both private and open. It is highly recommended that if you allow the HMC as a pass-through point that you also configure a VPN connection on the HMC as shown in "Internet VPN tab". By having the pass-through point through a VPN, you are encrypting the data and helping to keep it secure.

To use the HMC as a pass-through point:

1. Select **Allow pass-through systems for service** (Figure 6-29). Until you select this option, you are not able to allow systems for pass-through nor are you able to use the Edit, Remove, and Test options until you add a server.
2. To add servers for pass-through service, select **Add**.

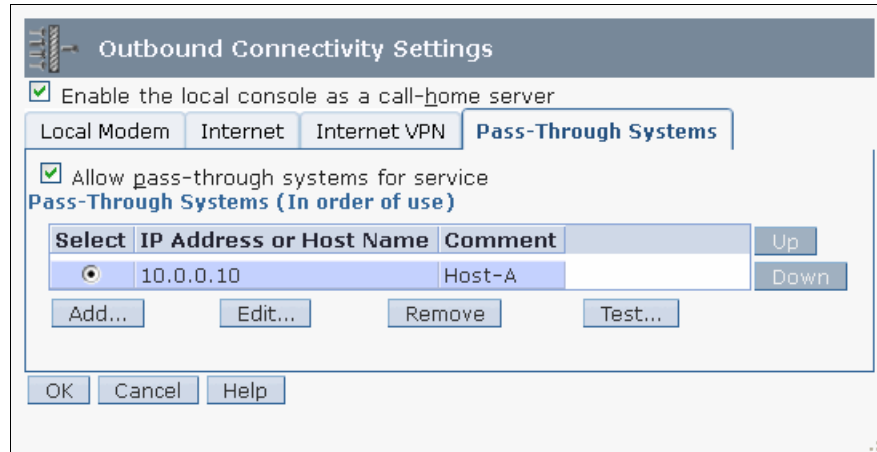


Figure 6-29 Manage Outbound Connectivity option: Pass-Through Systems tab

3. In the “IP address or host name” field, enter either the IP address or the fully qualified host name for the server that you want to add, and a comment for the server, then select Add. See Figure 6-30.

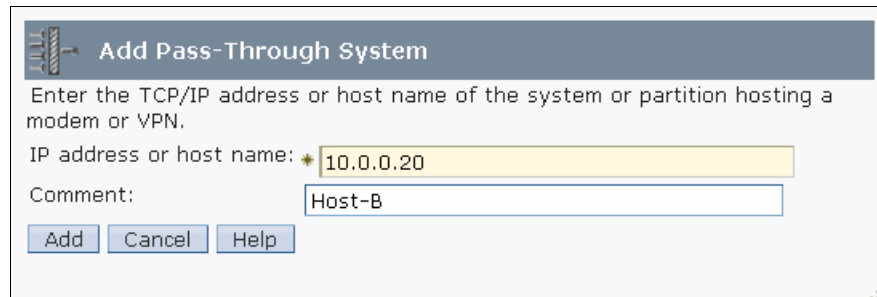


Figure 6-30 Add servers for pass-through access

4. If your selection is successful, you receive a results window similar to what is shown in Figure 6-31. The system that you entered is listed in the area labeled Pass-Through Systems.

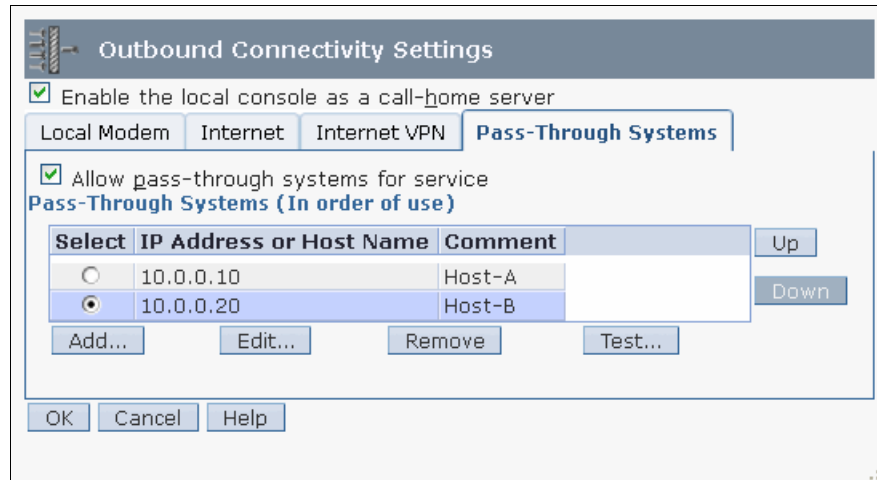


Figure 6-31 Add servers for pass-through access: Results

5. Now with a server or servers that are entered for pass-through, you can use the Edit, Remove, and Test options that are available on this window.

Manage Inbound Connectivity task

The Manage Inbound Connectivity task allows the HMC to receive a connection from an outside source through the Internet or over a modem. First, select **Service Management** → **Manage Inbound Connectivity**.

In the Manage Inbound Connectivity window, you use the Remote Service tab to allow an Internet connection through Point-to-Point Protocol (PPP) or VPN (Figure 6-32). Under Connection Type are the access types. You can allow local console and managed server partition access. To save your settings on this tab, select **OK**.

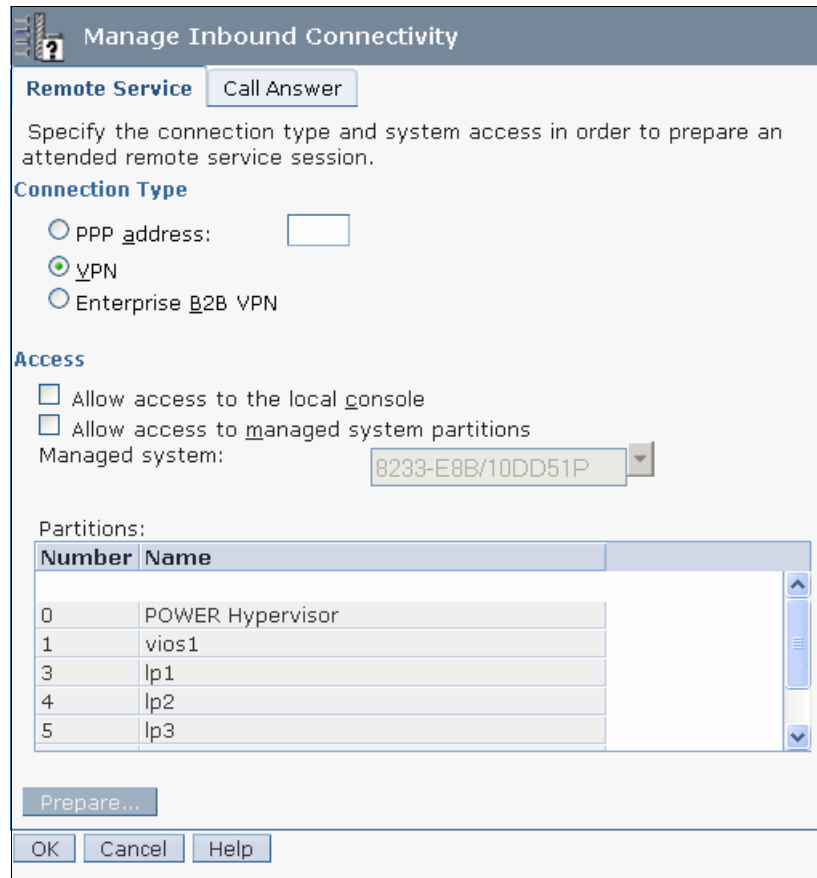


Figure 6-32 Manage Inbound Connectivity window: Remote Service tab

You use the Call Answer tab to allow a modem that is inside or connected to the HMC to answer incoming calls on the phone line to which it is connected. If you select **Allow local modem call answering** and then select **OK**, the HMC answers automatically on the phone line to which it is connected. The HMC then negotiates a connection with the other end if it is a modem. See Figure 6-33.

Security guideline: As a general security guideline, turn off this option.

Refrain from activating this setting unless you have approval from your local IT infrastructure's management and administrators. Many IT departments have rules and restrictions that govern the use of modems over phone lines in their environment. Therefore, allowing the modem on your HMC to answer phone calls might be a security violation in your environment.

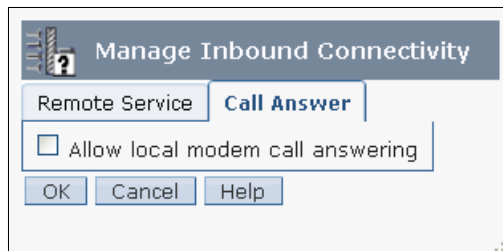


Figure 6-33 *Manage Inbound Connectivity* task, *Call Answer* tab

Manage Customer Information task

With the Manage Customer Information task, enter contact information that pertains to the HMC ownership. On the Administrator tab, you can enter a primary and secondary phone number, and the owning email address and fax number (Figure 6-34).

Manage Customer Information

Administrator System Account

Contact Information

Company name: *

Administrator name: *

Email address: *

Phone number: *

Alternate phone number:

Fax number:

Alternate fax number:

Mailing Address

Street address: *

Street address 2:

City or locality: *

Country or region: * <Select one>

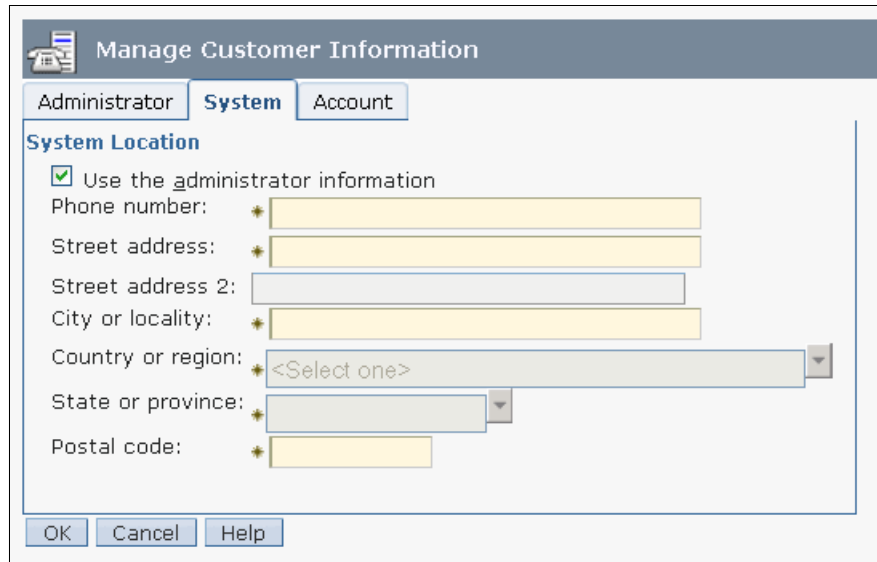
State or province: *

Postal code: *

OK Cancel Help

Figure 6-34 Manage Customer Information task, Administrator tab

On the System tab, you can enter the postal address at the HMC location (Figure 6-35). This information can be used to send you fixes and patches through mail and also can provide an address for service technicians for when the HMC must be serviced.



The screenshot shows a dialog box titled "Manage Customer Information" with three tabs: "Administrator", "System", and "Account". The "System" tab is selected. Under the "System Location" section, there is a checked checkbox labeled "Use the administrator information". Below this are several input fields, each with a red asterisk indicating a required field: "Phone number:", "Street address:", "Street address 2:", "City or locality:", "Country or region:" (a dropdown menu showing "<Select one>"), "State or province:" (a dropdown menu), and "Postal code:". At the bottom of the dialog are three buttons: "OK", "Cancel", and "Help".

Figure 6-35 Manage Customer Information task, System tab

Finally, on the Account tab, you enter the fields for the account that owns the HMC (Figure 6-36). If you do not know your customer number, enterprise number, and so forth, contact your sales representative who can provide this information to you.

The information about this window can be useful for IBM representatives to gather inventory and account data that is associated with your HMC and your enterprise.

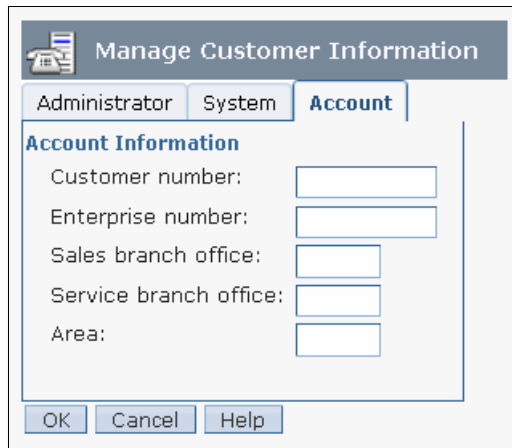
The image shows a screenshot of a software dialog box titled "Manage Customer Information". At the top, there are three tabs: "Administrator", "System", and "Account", with "Account" being the active tab. Below the tabs, the section is titled "Account Information". It contains five input fields, each with a label to its left: "Customer number:", "Enterprise number:", "Sales branch office:", "Service branch office:", and "Area:". At the bottom of the dialog box, there are three buttons: "OK", "Cancel", and "Help".

Figure 6-36 Manage Customer Information task, Account Information tab

Manage eService Registration task

With the Manage eService Registration task, add the HMC and managed servers to your IBM Electronic Services profile. By enabling eService registration, you can do the following functions:

- ▶ View the latest IBM Electronic Services news
- ▶ Customize the web page with links that apply to your systems
- ▶ View reports that are created from the Electronic Service Agent information that your system sent to IBM
- ▶ Submit a service request for hardware and software
- ▶ Search for information to solve your system problems
- ▶ Find services available in your country

Before you proceed with registering your HMC, visit the following website:

<http://www.ibm.com/support/electronic>

If you do not have an IBM ID, you must create one. Register as shown in Figure 6-37, at the following website:

<http://www.ibm.com/registration>

Already have an IBM ID?: If you already have an ID, then skip to “Registering eService from the HMC” on page 302.

Create an IBM ID, as shown in Figure 6-37. When you complete all the fields on this window, select **Continue**.

The screenshot shows the IBM registration page. At the top, there is the IBM logo and a navigation menu with links for Home, Solutions, Services, Products, Support & downloads, and My IBM. A search bar is located on the right. Below the navigation menu, there is a sidebar with links for My IBM profile, My IBM registration, Help and FAQ, and Help desk. The main content area is titled "My IBM registration" and "Step 1 of 2". It contains a form with the following fields and text:

- Asterisks (*) indicate fields required to complete this transaction.
- Preferred language for profiling: English
- IBM has sold its PC business to Lenovo Group Ltd. To facilitate your ability to browse for information on PC products and services, your ID and password will provide you access to both the IBM and Lenovo web sites. IBM is not responsible for the privacy practices or the content of the Lenovo web site. [Learn more](#) about IBM & Lenovo.
- Remember, you can't change your IBM ID once you've signed up.
- To learn what is acceptable as a password, see [guidelines for IBM IDs and passwords](#).
- * IBM ID:
- [Why do I have to provide an email address as my IBM ID?](#)
- * Password:
(Minimum 8 characters)
- * Verify password:
- Please enter a security question that only you can answer. Then, enter the answer to the question. Occasionally, you may be asked to answer this question to confirm your identity. Enter a question that is simple to answer and is easy to remember.
- * Security question:
- * Answer to security question:

Figure 6-37 Electronic services, registration

Registering eService from the HMC

When you have an IBM ID, select **Service Management** → **Manage eService Registration**. Then, follow these steps:

1. Enter the email address that you used to create your IBM ID and then select **OK** (Figure 6-38).

Manage eService Registration

IBM provides personalized Web functions that use information collected by IBM Electronic Service Agent. To use these functions, such as to download fixes directly to your HMC or servre firmware, you must have an e-mail address registered on the IBM Registration website at <https://www.ibm.com/account/profile>. If you already have an e-mail address at that site, enter it below. Else, create an account profile first.

Now, enter the e-mail address below to associate this system with the IBM Website.

Web authorization

e-mail ID 1 *

e-mail ID 2 (optional)

You may use the following website to view any or all systems that you have registered to IBM with the account profile above:
<http://www.ibm.com/support/electronic>.

Figure 6-38 Enter email addresses to associate the HMC with eService

2. When you register your HMC successfully, revisit the website and log in to access the eService and **Sign in** (Figure 6-39):

<http://www.ibm.com/support/electronic>

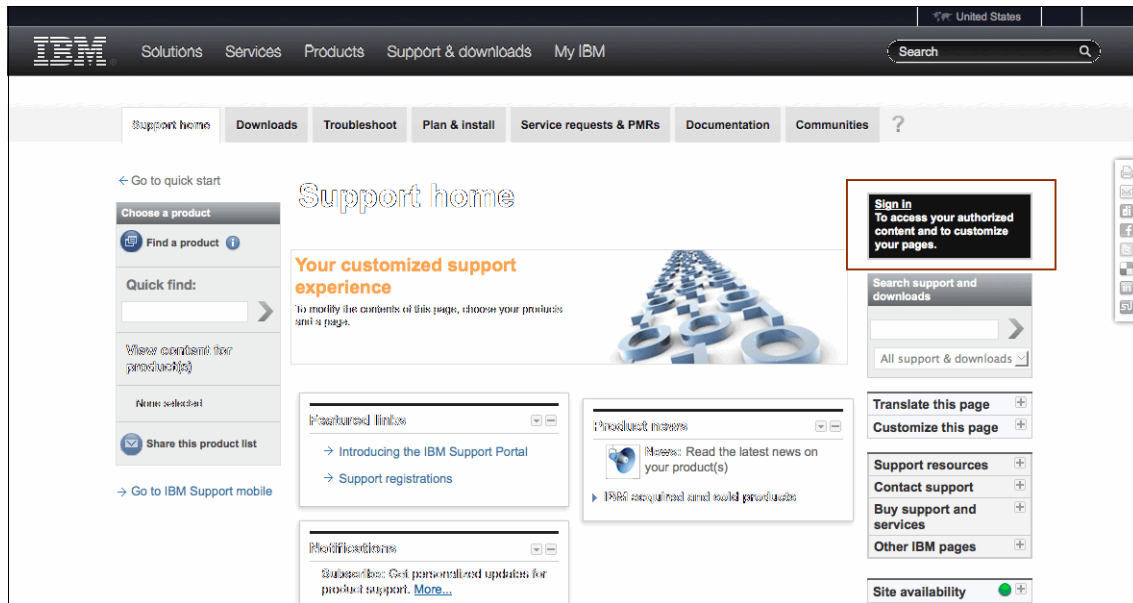


Figure 6-39 IBM Electronic Support website; Sign in with your IBM ID

- Then, select **My Systems**, as shown in Figure 6-40, to view the servers that are associated with your IBM ID on the Electronic Services website.

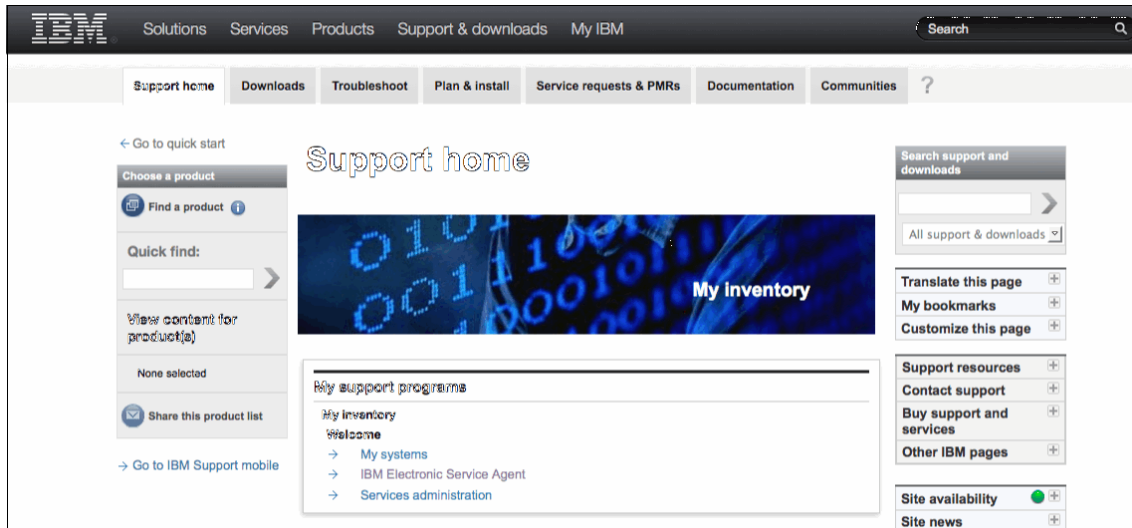
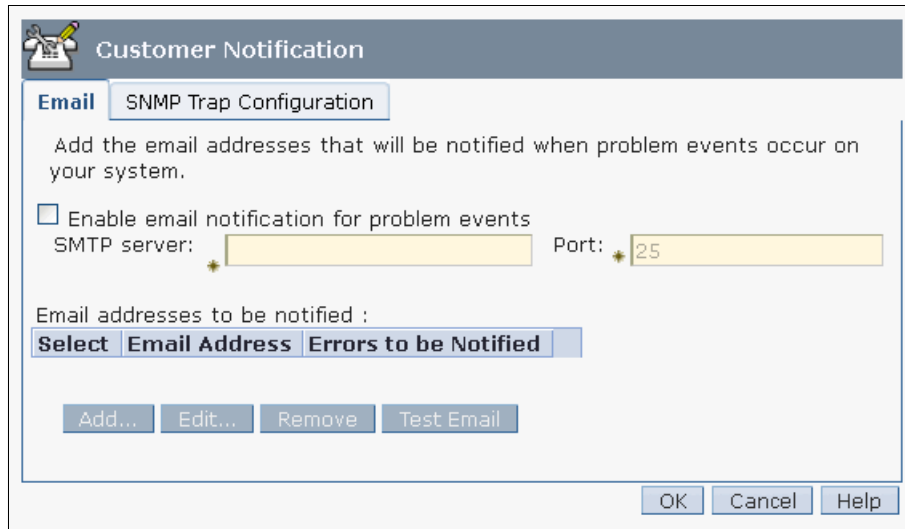


Figure 6-40 Electronic services My systems portal

Manage Serviceable Event Notification task

This option allows you to set up email addresses to be contacted for hardware notices or all serviceable events.

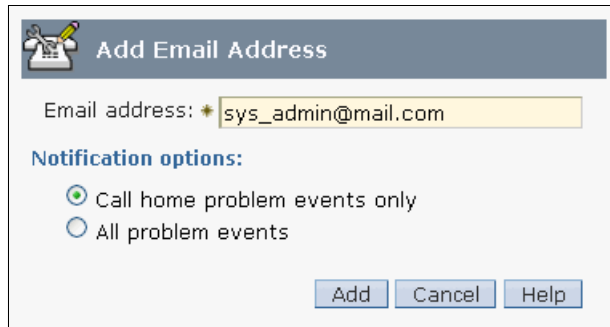
To set up email addresses to be contacted for serviceable events, select **Service Management** → **Manage Serviceable Event Notification**. On the Email Address tab, click **Add** to enter an email address (Figure 6-41).



The screenshot shows a dialog box titled "Customer Notification" with a sub-tab "Email" and "SNMP Trap Configuration". The main text reads: "Add the email addresses that will be notified when problem events occur on your system." Below this is a checkbox labeled "Enable email notification for problem events". Underneath the checkbox are two input fields: "SMTP server:" followed by a yellow text box, and "Port:" followed by a yellow text box containing the number "25". Below the input fields is the text "Email addresses to be notified :". Underneath this text is a table with three columns: "Select", "Email Address", and "Errors to be Notified". Below the table are four buttons: "Add...", "Edit...", "Remove", and "Test Email". At the bottom right of the dialog box are three buttons: "OK", "Cancel", and "Help".

Figure 6-41 Manage Serviceable Event Notification task

You can enter an email address to be contacted and then select to choose to have this email address contacted for either all problem events or just call home hardware events (Figure 6-42). You can enter multiple email addresses one at a time, and select **Add**. When you add all the email addresses that you want contacted for serviceable events, select **Cancel** to close this window and to return to the main HMC view.



Add Email Address

Email address: * sys_admin@mail.com

Notification options:

- Call home problem events only
- All problem events

Add Cancel Help

Figure 6-42 Add Email Address panel for notification

Manage Connection Monitoring task

Connection monitoring generates serviceable events when communication problems are detected between the HMC and managed systems. If you disable connection monitoring, no serviceable events are generated for networking problems between the selected machine and this HMC.

To enable connection monitoring with a managed server, select **Service Management** → **Manage Connection Monitoring**. The Manage Connection Monitoring window opens, as shown in Figure 6-43.

Manage Connection Monitoring

Connection Monitoring Timer Settings

Number of disconnected minutes considered an outage: 15

Number of connected minutes considered a recovery: 2

Number of minutes between outages considered a new incident: 20

Select one or more machines. Then click Enable or Disable to enable or disable customer notification of connection monitoring errors from this HMC to the selected machines. Enabling this capability allows connection monitoring errors to be passed to Service Agent for notification. However, you must still configure Service Agent to handle those notifications.

Select	Machine Name	State	Machine Type-Model/SN
<input type="checkbox"/>	8233-E8B-SN10DD51P	Enabled	8233-E8B/10DD51P

Enable Disable

OK Cancel Help

Figure 6-43 Manage Connection Monitoring task

Here, you can manipulate the following components:

- ▶ **Number of disconnected minutes considered an outage:** Set the number of minutes of disconnect that are considered an outage. If connectivity is restored before this threshold is met, the managed server is considered recovered and no serviceable event is reported.
- ▶ **Number of connected minutes considered a recovery:** Specify the number of minutes of required connectivity that is required to put a managed server in a recovered state. This number is directly associated with the disconnected minutes threshold, in that this value is monitored when the outage threshold is met.
- ▶ **Number of minutes between outages considered a new incident:** Specify the amount of time that is required between outages that are required to report a new outage report.

Select the managed server that you want and then select either **Enable** or **Disable** to manipulate connection monitoring. To save your settings, select **OK**.

Manage POWER4 Service Agent task

If your IT environment has POWER4 servers that are attached to an HMC, you can use the Manage POWER4 Service Agent option to have your V7 HMC act as a focal point for Service Agent reporting.

To use this option, select **Service Management** → **Manage POWER4 Service Agent**. Then, in the window that opens, select **Enable Service Agent Connection Manager**, as shown in Figure 6-44.

Next, you can manipulate settings for the following features:

- ▶ **Secure Mode:** If cleared, data transmission is unencrypted. If selected, then data between the HMC and the POWER4 HMC is encrypted through the HTTPS protocol.
- ▶ **URL for configuration download.**
- ▶ **Password for configuration updates:** It is highly recommended that you change the password from its default value for security purposes.

When you finish with your selections, select **OK**.

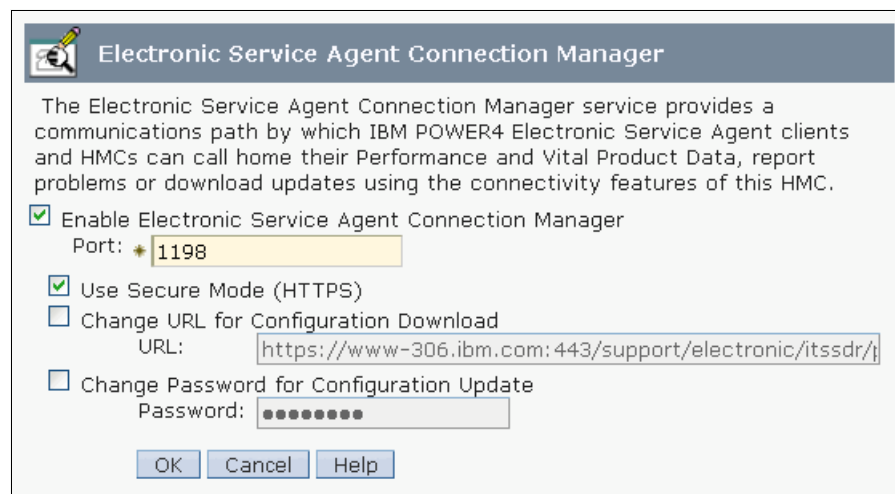


Figure 6-44 Manage POWER4 Service Agent task

6.2 Software maintenance

There are various options that are available for maintaining both Hardware Management Console (HMC) and managed system firmware levels. We show you, through examples, the main firmware update options.

We describe the different methods of updating the HMC to a new software level and installing individual fix packs. We also cover the temporary and permanent side of the firmware on a POWER7. Finally, we show the various options available to POWER7 system's firmware, either through the HMC or through an AIX service partition.

6.2.1 HMC Data backup

Before you begin any firmware upgrade, it is important that you maintain a current HMC Data backup or Critical Console Data (CCD) backup. This backup can be useful in recovering the HMC in the event of the loss of a disk drive.

When you move to a new version level of HMC or use a Recovery CD to update the HMC, you must create an HMC Data backup immediately following the installation. If you update HMC code between releases by using the Corrective Service files downloadable from the web and then create new HMC Data backups after the update, you can use those HMC Data backups and the last-used recovery CD to rebuild the HMC to the level in use when the disk drive was lost.

Another example where an HMC Data backup would be useful is when replacing a service processor or BPC on a POWER6 or POWER7 processor-based server. You have to make a fresh HMC Data backup *before* starting the replacement to preserve the DHCP lease file on the HMC that lists the starting flexible service processor (FSP) and BPC IP addresses. If for some reason things do not work after you replace the FSP or BPC, you can use the backup to restore the original information. If the replacement is successful, a new IP address is assigned to the new component, and the lease file is updated. A new HMC Data backup is created that captures the freshly updated DHCP lease file.

With the HMC, you can back up the following important data:

- ▶ User-preference files
- ▶ User information
- ▶ HMC platform-configuration files
- ▶ HMC log files
- ▶ HMC updates through Install Corrective Service

Use the archived data only with a reinstallation of the HMC from the product CDs.

Attention: You cannot restore the HMC Data backup on different versions of HMC software.

Manual backup of HMC Data

To back up the HMC, you must be a member of one of the following roles:

- ▶ Super administrator
- ▶ Operator
- ▶ Service representative

You must format the DVD in the DVD-RAM format before you can save data to the DVD. To format a DVD, select **HMC Management** → **Format Media** from the HMC workplace window (Figure 6-45).

To back up HMC Data, click **HMC Management** → **Backup HMC Data**.

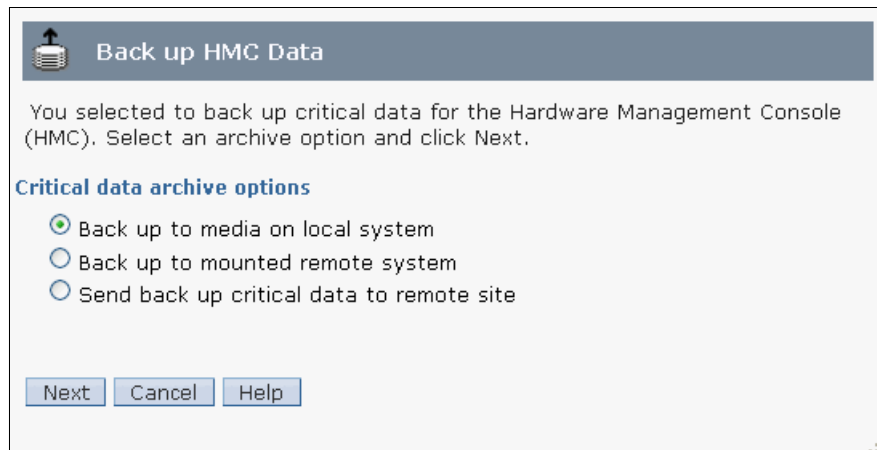


Figure 6-45 Back up HMC Data menu

Then, select an archive option. You can back up to a local media (DVD or USB flash memory device) on the HMC, back up to a remote system mounted to the HMC file system, or a remote site through FTP. After you select an option, click **Next** and follow the instructions to back up the data.

Scheduled HMC Data backup

Back up the HMC Data up at least once a week. Also keep two copies of the HMC Data backup: one copy from the upgrade or any changes to the HMC, and one backup HMC Data to store in a safe place. For information about how to make a backup, see 6.2.1, “HMC Data backup” on page 309.

To schedule a CCD backup, select **HMC Management** → **Schedule Operations** from the HMC workplace window. Then, follow these steps:

1. In the Customize Scheduled Operations window, select **Options** → **New** (Figure 6-46).

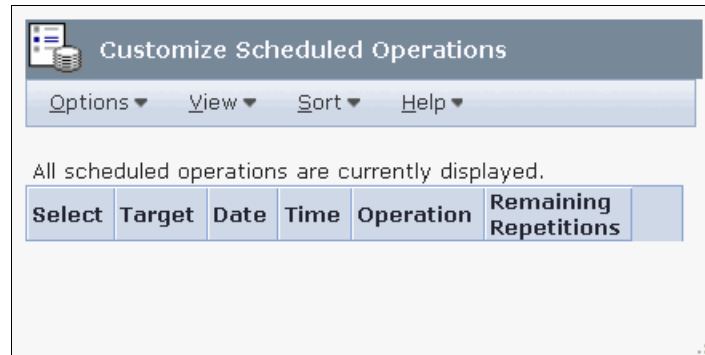


Figure 6-46 Customize Scheduled Operations window

2. In the Add a Scheduled Operation window, select **OK**.

3. On the Date and Time tab, select the date and time, and time window for the first backup, as shown in Figure 6-47. The scheduled operation starts at that time window.

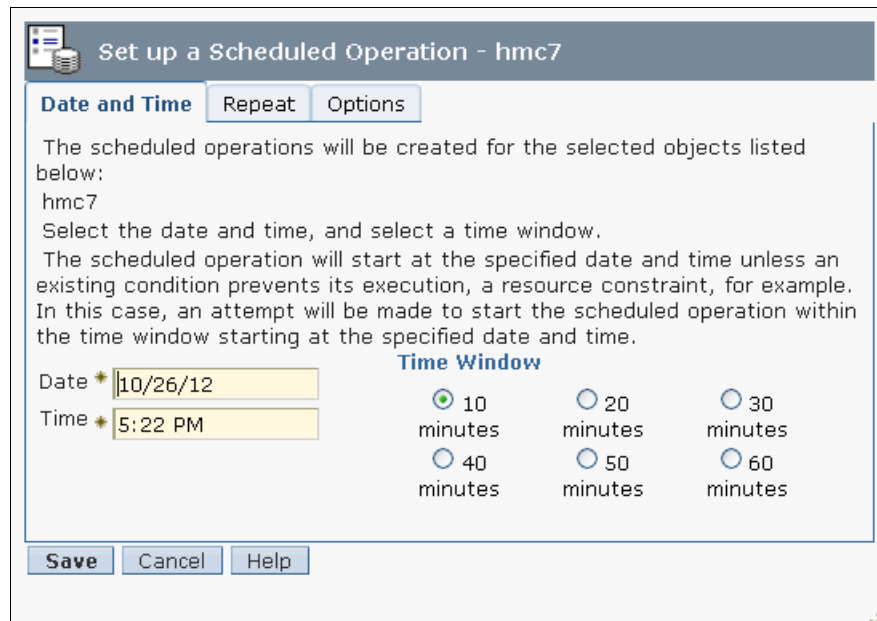


Figure 6-47 Set up a Scheduled Operation window

4. On the Repeat tab, select the repeat options for the backup (Figure 6-48).

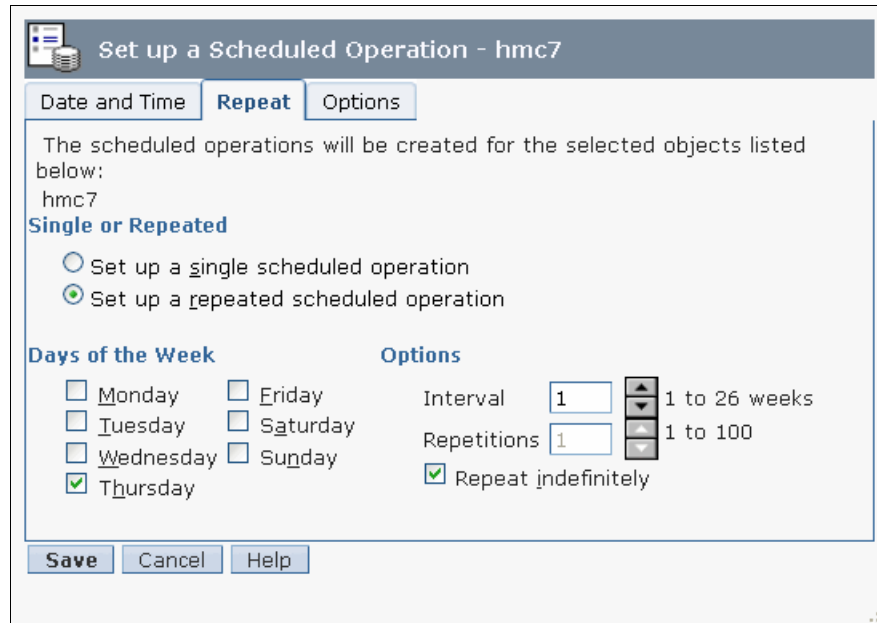


Figure 6-48 Scheduled backup HMC Data Repeat Option

5. On the Options tab, select the options that are available to back up media (Figure 6-49). The most common option is Local DVD media.

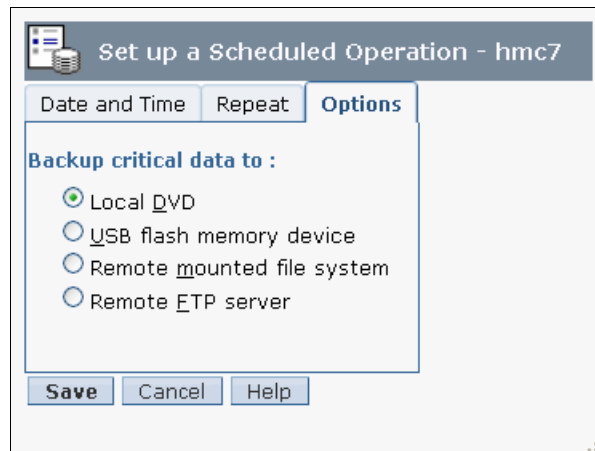


Figure 6-49 Scheduled HMC Data backup storage Options tab

6. After setting all the options, select **Save**. Then, select **OK**.
7. After you save the scheduled operations, you can view the operations by selecting **HMC management** → **Schedule Operations**. Select the CCD scheduled operation and then select **View**.

6.2.2 Restoring HMC Data

There are different ways of restoring HMC Data, depending on the option that you use to back up the data. These options are now described.

Restoring data from Removable Media

Restore HMC Data from the menu that is displayed at the end of the HMC reinstallation. You can choose between two removable media: a DVD backup, or USB flash memory device backup. To restore from DVD, you have to insert the DVD that contains the archived HMC data and select DVD option. To restore from a USB flash memory device, insert the USB flash memory device in one of the HMC USB ports, and select restore from USB flash memory device. On the first start of the newly installed HMC, the data is restored automatically. This option also works after the installation of the HMC is done. Follow these instructions.

Restoring data from FTP, SFTP, NFS, or removable media

If the critical console data was archived remotely either on an FTP server or remote file system, follow these steps:

1. Manually reconfigure network settings to enable access to the remote server after the HMC is installed. For information about configuring network settings, see 4.1.2, “Configuring the HMC network setting” on page 98.
2. In the HMC workplace window, select **HMC Management** → **Restore HMC Data**. Then, select the type of restoration and click **Next**.

3. Follow the directions to restore the HMC Data. The data restores automatically from the remote server when the system restarts. See Figure 6-50.

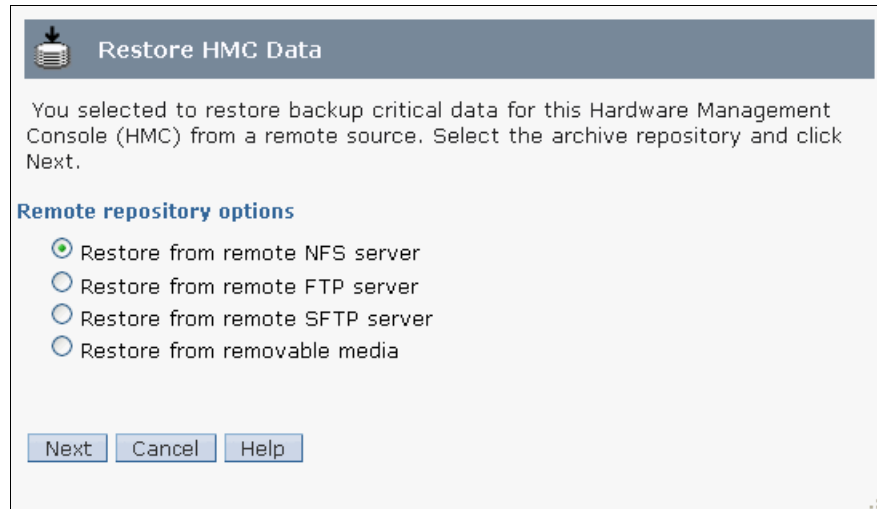


Figure 6-50 Restore HMC Data menu from HMC management

6.2.3 HMC software maintenance

The HMC is independent from the server. The server and all partitions can remain active while maintenance is done on the HMC, allowing you to easily keep your HMC at the latest maintenance level.

The HMC software level must be maintained the same as managed system firmware. HMC firmware is packaged as a full recovery CD set or as a corrective service pack or fix image. The HMC recovery CDs are bootable images and can be used to perform a complete recovery of the HMC (scratch installation) or an update to an existing HMC version.

A corrective fix updates the minor version level of code on the HMC. The HMC update packages are available on CDs or as downloadable, compressed files. The downloadable, compressed files have different naming formats depending on whether they are individual fixes or complete update packages:

- ▶ MHxxxx.zip - individual HMC fixes
Where xxxx is the HMC fix number

- ▶ HMC_Update_VxRyMz_n.zip - HMC update packages

Where *x* is the version number, *y* is the release number, *z* is the modification number, and *n* is the image number (if there are multiple images)

How to determine the HMC software version

The level of machine code on the HMC determines the available features, including concurrent server firmware maintenance and enhancements to upgrading to a new release.

To determine the HMC software version, click **Updates** in the HMC workplace window. In the work area, view and record the information that displays under the HMC Code Level heading, including the HMC version, release, maintenance level, build level, and base versions. See Figure 6-51.

The screenshot displays the Hardware Management Console (HMC) interface. The top navigation bar includes 'Welcome', 'Systems Management', 'Servers', 'System Plans', 'HMC Management', 'Service Management', and 'Updates'. The 'Updates' section is active, showing 'HMC Code Level' information. A red box highlights the following details: Version: 7, Release: 7.6.0, Service Pack: 0, Build Level: 20121005.1, and Base Version: V7R7.6.0. Other details include Serial Number: 107625C, Model Type: 7042CR6, and BIOS: D6E157AUS-1.15. Below this, the 'System Code Levels' section features a table with columns for Name, Status, Platform IPL Level, Activated Level, EC Number, and Deferred Level. The table contains one entry for the server 8233-E8B-SN10DD51P, which is in an 'Operating' status. The table also shows a filter, tasks, and views dropdown, and a footer indicating 'Total: 1 Filtered: 1 Selected: 0'.

Figure 6-51 Shows the version number of HMC and managed systems

6.2.4 Which firmware or fix level is correct for your system

One of the most important tasks is to determine the correct level of firmware or fix level for your system. IBM has an online tool that is called the *Fix Level Recommendation Tool* (FLRT). The Fix Level Recommendation Tool assists system administrators in formulating a maintenance plan for IBM Power Systems servers.

For each hardware model that you select, the tool displays fix level information in a report for the following components:

- ▶ HMC
- ▶ System Firmware (SF)
- ▶ AIX 5L
- ▶ Virtual I/O Server (VIOS)
- ▶ High Availability Cluster Multi-Processing (HACMP)

- ▶ General Parallel File System (IBM GPFS™)
- ▶ Cluster Systems Management (CSM)
- ▶ PowerHA IBM SystemMirror®
- ▶ IBM Software Component: Information Management, IBM Lotus®, IBM Rational®, IBM Tivoli® software, and IBM WebSphere®

To connect to the Fix Level Recommendation Tool, see the following website:

<https://www.ibm.com/support/customer/fixlevelrecommendationtool/>

Figure 6-52 shows the Fix Level Recommendation Tool website main page.

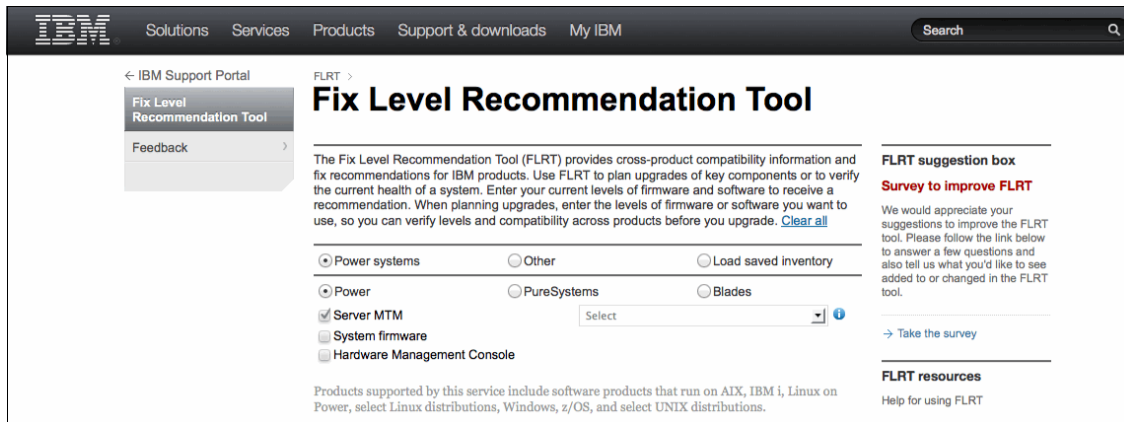


Figure 6-52 Fix Level Recommendation Tool website

Only the products that you select on the Fix Level Recommendation Tool entry page are listed on the *inventory page*. The Fix Level Recommendation Tool is most useful for querying the combination of two or more products to ensure that they are at the recommended level and that any interdependencies are met. For example, specific system firmware levels are required for some HMC releases, and Virtual I/O Server virtualizes only a system with AIX 5L Version 5.3 and higher or IBM i V6R1 and higher.

The report that the Fix Level Recommendation Tool produces includes two sections:

- ▶ Your selected level
- ▶ The recommended minimum fix level

Take the following steps to run an FLRT report:

Selecting products for an FLRT report

On the Fix Level Recommendation Tool main page, select the products for which you want to check recommended levels and click **Submit**. The Fix Level Recommendation Tool displays the information that is shown in Figure 6-53.

Power systems Other Load saved inventory

Power PureSystems Blades

Server MTM 8233-E8B (Power 750 Express)

System firmware AL730_095

Hardware Management Console V7 R750 SP1

LPAR_01

Partition name:

Partition type:

Product search:

▶ **Disk systems**

▼ **Virtualization software**

PowerVM Virtual I/O Server 2.2.1.3

▼ **Operating systems**

AIX 6100-07-05

▶ **Cluster software**

▶ **Information management**

▶ **Lotus**

▶ **Rational**

▶ **Tivoli software**

▶ **WebSphere**

Name of report (optional):

Figure 6-53 Fix Level Recommendation Tool product options window

FLRT product selection page

The FLRT landing page initially displays a list of operating systems. Other products become available after an operating system is selected.

Operating system family

From this section, select the operating system running on your server, or in one of the LPARs on your server. After you choose an operating system, FLRT displays selections for Platform and Server.

Platform

From this section, choose either Power Systems or BladeCenter (POWER based processors). FLRT supports blade servers that are based on POWER technology.

Server

Use the drop-down menu in this section to select the machine type and model (MTM) for your server and the clock speed (GHz). If only one clock speed is available for the selected server, FLRT displays that clock speed.

Enter the current details from your system. To find the current fix level of the HMC and managed system, refer to “How to determine the HMC software version” on page 316. Select **Submit**.

The Fix Level Recommendation Tool displays the report, as shown in Figure 6-54 on page 320.

FLRT >

Fix Level Recommendation Tool

The following consolidated information is for guidance purposes only. This information was obtained from generally available product support documentation. These combinations of product levels are supported by IBM.

Date: 2012.10.29

Model: IBM Power 750 Express (8233-E8B)
Click [here](#) for the latest device firmware for this model.

[+ Expand all](#) [- Collapse all](#)

System results

		Input level	Recommended update	Recommended upgrade
✓	System firmware Release date EoSPS	AL730_095 2012.08.23 2013.05.31	None	None
✓	HMC Software Release date EoSPS	V7 R750 SP1 2012.08.17 2014.05.31	MH01325 2012.09.07	None

LPAR_01 - AIX

		Input level	Recommended update	Recommended upgrade
✓	VIOS Release date EoSPS	2.2.1.3 2011.12.14 Not Announced	2.2.1.4 2012.05.23 Not Announced	None
✓	AIX Release date EoSPS	6100-07-05 2012.07.18 2014.10.17	6100-07-05 2012.07.18 2014.10.17	None

The recommendations database was last updated on 2012.10.26.

Save inventory

Figure 6-54 Fix Level Recommendation Tool recommendation window





Report heading

The heading of the report lists the date and name of the report (if you entered a name). It also lists the server that you selected and the clock speed. The heading also includes a link to a page that shows you the latest firmware for the devices that are supported by your machine. You can download device firmware from that same page.

Detailed results

This section displays product compatibility, detailed results, and fix recommendations. FLRT uses several icons to indicate the type of information that is displayed in the report, as shown in Table 6-1. More information is supplied, depending on the type of result indicated.

Table 6-1 Several icons to indicate the type of information in the FLRT report

	Ok Green check mark: Displayed when the level that you input is supported for longer than six months and the input level is the currently recommended update or latest level.
	Blue “i” (information) circle: Displayed when the input level is supported for longer than six months but there is also a recommended update available.
	Yellow caution triangle: Displayed when the input level has less than six months that are left for service and there is also a recommended upgrade available. If the input level no longer has updates available, the yellow caution triangle is displayed, if there is an upgrade option.
	Alert Red stop sign: Displayed for incompatibilities. If end of service pack support (EoSPP) is reached and there is no upgrade recommendation, the red stop sign is also displayed.

Obtaining HMC updates and recovery software

You can order Recovery CDs or download packages that contain the files that you have to burn your own Recovery CD. The files that you use to create CDs have a .iso file extension. The CDs created from these packages are bootable. You can download updates to the HMC code and emergency fixes, and you can order CDs containing the updates and fixes. The CDs containing updates and fixes are *not* bootable.

Important: If you are not sure what code level is correct for your machine, read 6.2.4, “Which firmware or fix level is correct for your system” on page 316.

Use the following URL to download the latest HMC software:

<http://www.ibm.com/support/fixcentral/options>

Fix Central

Fix Central provides fixes and updates for your system's software, hardware, and operating system.

For additional information, click on the following link.
[Getting started with Fix Central](#)

Select product Find product

Select the product below.

When using the keyboard to navigate the page, use the **Alt** and **down arrow** keys to navigate the selection lists.

Product Group
Systems **← 1. Select IBM System**

Select from Systems
Power **← 2. Select Power System Server**

Product
Firmware, SDMC and HMC **← 3. Select Firmware HMC**

Machine type-model
Select one **← 4. Select your machine type-model**

Continue

Figure 6-55 Displays the HMC software that is available on the website

Complete the required fields, as shown in Figure 6-55. After completing the selections, click **Continue**.

The Fix Central website shows you available firmware options. Choose HMC Firmware, as shown in Figure 6-56.

Firmware and HMC

Select fix type

Current selections

→ Your Machine Type-Model is 8233-E8B

Available options

- All firmware components. Obtain system firmware, device firmware, SDMC, and HMC updates. The power subsystem firmware will be included if applicable.
- System firmware. Obtain system firmware only. The power subsystem firmware will be included if applicable.
- Device firmware. Obtain device firmware only. Available for adapters hard disks and media devices.
- SDMC Code. Obtain SDMC Update Images, Service packs, and Interim fixes.
- HMC Firmware. Obtain HMC Recovery images, service packs, and specific fixes.

Figure 6-56 Choose HMC Firmware for select fix type option

After you press **Continue**, Fix Central asks about the HMC release level. In this option, choose V7R760.

Firmware and HMC

Select HMC release level

Current selections

- Your Machine Type-Model is 8233-E8B
- "HMC Firmware only" has been selected

HMC firmware release levels

The HMC firmware level must be compatible with all the systems it s managing. Click [here](#) to view the HMC to system firmware supported combinations table.

V7R7.6.0

Continue **Back**

Figure 6-57 HMC firmware release level menu

Select the version of the software that you want to download. then press **Continue** (Figure 6-58). The next window opens the license agreement acceptance.

Firmware and HMC

Select HMC fixes

Current selections

- Your Machine Type-Model is 8233-E8B
- "HMC Firmware only" has been selected
- You have requested Release Level V7R7.6.0

HMC Package Selection

Select the HMC Recovery Firmware, Service Pack, and Specific Fixes to include in your package.
Click [here](#) for the HMC / System Firmware supported combinations table.

Select HMC Recovery V7R7.6.0M0

<input checked="" type="checkbox"/>	HMC Recovery V7R7.6.0M0 Recovery Image	Released 18 Oct 2012	MH01326	Description
<input checked="" type="checkbox"/>	HMC Recovery V7R7.6.0M0 Release Update Package	Released 18 Oct 2012	MH01327	Description
HMC specific fixes				
<input checked="" type="checkbox"/>	Mandatory efix for HMC V7R760	Released 18 Oct 2012	MH01328	Description

Figure 6-58 HMC software and update files available on the IBM website

Save the file on your computer and burn it on a DVD or order a CD from IBM, as shown in Figure 6-59.

IBM ID needed: You need an IBM ID. Select **IBM ID** and follow the instructions to register. After you register, log in to the website and complete the necessary information.

Firmware and HMC

Package download

Obtain package

Download using Download Director (requires Java installation) [What is this?](#)

Download using Bulk FTP

Download ISO image to burn to media

Order fixes on media

Include informational files, and files required by the HMC, SDMC and Systems Director for installation

Package list

HMC Recovery V7R7.6.0M0	MH01326	3003469824 bytes
HMC V7R7.6.0M0	MH01327	2455799808 bytes
Mandatory efix for HMC V7R760	MH01328	1091108864 bytes
<i>Total</i>		6550378496 bytes

Figure 6-59 Downloading option HMC software and fix images

Obtaining and applying HMC code from an FTP server

If your HMC has a VPN connection to the Internet, you might choose to do the HMC update directly from the IBM support FTP server.

Important: Some of the HMC update packages are large (over 2 GB) and can take time to download.

To do the HMC update directly from an FTP server, follow these steps:

1. First, back up the HMC Data as described in 6.2.1, “HMC Data backup” on page 309.
2. Then, in the HMC workplace window, click **Updates** → **Update HMC**. The Install Corrective Service window opens (Figure 6-60).
3. On the Installation and Configuration Options window (Figure 6-60), fill the required fields.

Install HMC Corrective Service Wizard

Installation and Configuration Options

Remote server was selected as the source where the HMC corrective service packages can be located.

Choose the remote server type.

FTP
 NFS

Type the information required to access the remote FTP server.
Click Next to continue.

Remote Server: * ftp.software.ibm.com
User ID: * anonymous
Password: *
Remote directory: /software/server/hmc/fixes/

* indicates a required parameter

< Back Next > Finish Cancel Help

Figure 6-60 Install Corrective Service window

- **Remote site:** ftp.software.ibm.com
- **Patch file:** /software/server/hmc/fixes/

Name of patch file changes: The name of the patch file changes with each new update. Refer to 6.2.4, “Which firmware or fix level is correct for your system” on page 316.

- **User ID:** anonymous
 - **Password:** Your email address
4. Click **Next**. You are given a set of available updates.
 5. Follow the instructions to install the update.
 6. Shut down and restart the HMC for the update to take effect.

Applying HMC code from CD or DVD

To apply HMC code from a CD or DVD, follow these steps:

1. You can either order a DVD with HMC updates from IBM, or download .iso or a compressed file from the IBM software support site. For more information, see “Obtaining HMC updates and recovery software” on page 321.
2. Insert the CD or DVD in the HMC.
3. In the HMC workplace window, click **Updates** → **Update HMC**. The Install Corrective Service window opens (Figure 6-60 on page 327).
4. Select **Apply corrective service from removable media (CD/DVD or USB flash memory device)** and click **OK**.
5. Follow the instructions to install the update. If you have more than one CD or DVD, then follow the procedure from step 2 for the second CD or DVD.
6. Shut down and restart the HMC for the update to take effect.
7. To verify that the HMC machine code update installed successfully, see “How to determine the HMC software version” on page 316.
8. If the level of code that is displayed is not the level that you installed, do the following steps:
 - a. Try the machine code update again. If you created a CD or DVD for this procedure, use a new media.
 - b. If the problem persists, contact your next level of support.

Upgrading the HMC machine code

Upgrade restrictions: You cannot use this procedure to upgrade from a POWER4 HMC to a POWER5 HMC. You must do a full installation. To upgrade from Version 6 to Version 7, see “Upgrading HMC from Version 6 to Version 7” on page 333.

To upgrade the HMC machine code, follow these steps:

1. Determine the HMC machine code level that is required for your system. See 6.2.4, “Which firmware or fix level is correct for your system” on page 316.
2. Obtain the recovery image. See “Obtaining HMC updates and recovery software” on page 321.
3. Back up the profile data of the managed system. In the HMC workplace window, select **System Management** → **Servers**. Then, select the server and ensure that the state is *Operating* or *Standby*.

Under Tasks, select **Configuration** → **Manage Partition Data** → **Backup**. Enter a backup file name and record this information. Then, click **OK**.

Repeat these steps for each managed system.

4. Backup critical console data as described in 6.2.1, “HMC Data backup” on page 309.

Back up the HMC Data: It is absolutely necessary to back up the HMC Data.

5. Before you upgrade to a new version of HMC software, as a precautionary measure, record HMC configuration information as follows:
 - a. In the HMC workplace window, select **HMC Management**. Then, in the tasks list, select **Schedule Operations**. The Scheduled Operations window displays with a list of all managed systems.
 - b. Select the HMC that you plan to upgrade and click **OK**. All scheduled operations for the HMC display.

Tip: If you do not have any scheduled operations, skip to step 6.

- c. Select **Sort** → **By Object**. Select each object and record the following details:
 - Object Name
 - Scheduled date

- Operation Time (displayed in 24-hour format)
 - Repetitive
- If Yes, select **View** → **Schedule Details**. Then, record the interval information and close the scheduled operations window. Repeat for each scheduled operation.
- d. Close the Customize Scheduled Operations window.
6. Record remote command status:
- a. In the navigation area, select **HMC Management**. Then, in the tasks list, click **Remote Command Execution**.
 - b. Record whether the Enable remote command execution using the **ssh** facility check box is selected.
 - c. Click **Cancel**.

Saving upgrade data

You can save the current HMC configuration in a designated disk partition on the HMC. Save upgrade data only immediately before upgrading your HMC software to a new release. This action allows you to restore HMC configuration settings after upgrading.

Only one level of backup data is allowed: Each time that you save upgrade data, the previous level is overwritten.

HMC Version 7 also gives an option to save upgrade data on hard disk, DVD, or USB flash memory device. It is strongly suggested to save a copy on a USB flash memory device. See Figure 6-61.

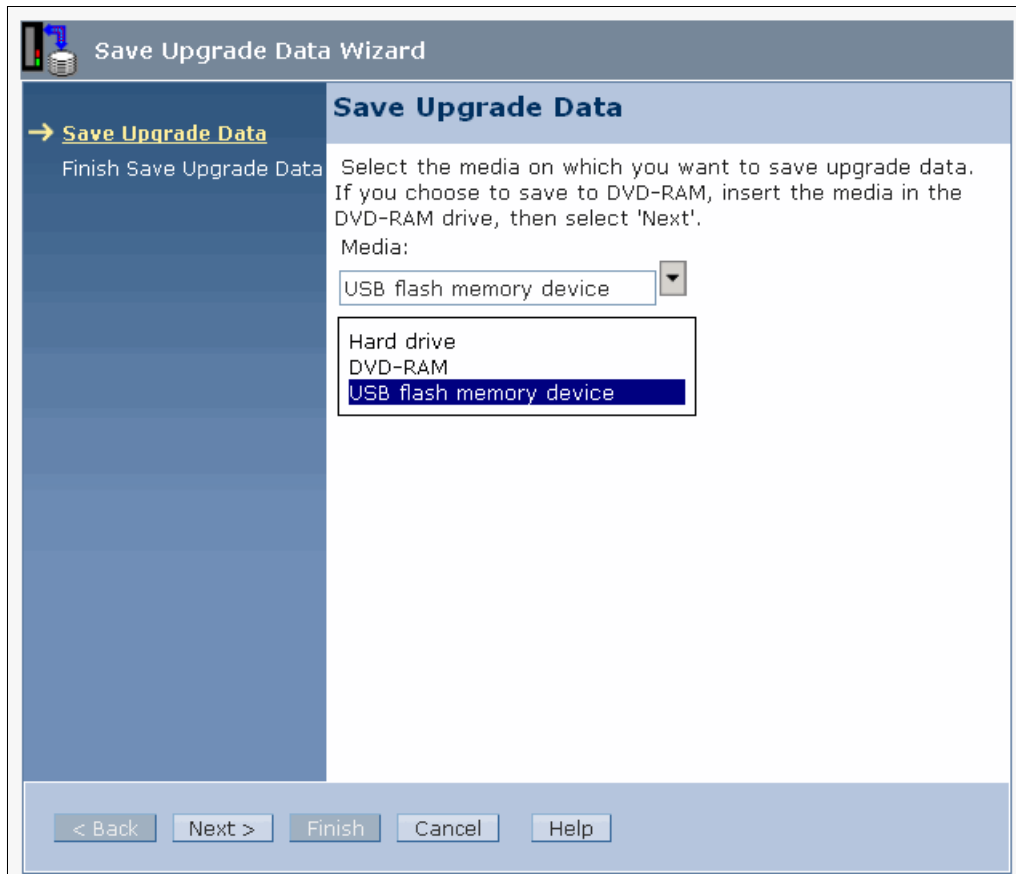


Figure 6-61 Save Upgrade Data wizard

To save upgrade data:

1. In the HMC workplace window, select **HMC Management**. Then, in the tasks list, select **Save Upgrade Data**. Select **Hard drive**.
2. Repeat the process of number 1, but using USB flash memory device option, then click **Next**.
3. Click **Finish**.
4. Wait for the task to complete. If the Save Upgrade Data task fails, contact your next level of support before proceeding.

Important: If the save upgrade data task fails, do not continue the upgrade process.

5. Click **OK**. Then, click **Close**.

To upgrade the HMC software:

1. Restart the system with the Recovery DVD-RAM in the DVD-RAM drive by inserting the HMC Recovery DVD-RAM into the DVD-RAM drive.
2. In the navigation area, select **HMC Management** → **Shutdown or Restart**. Then, select **Restart the HMC** and click **OK**.
3. The HMC restarts and boots from the bootable recovery DVD. The window shows the following options:
 - Install
 - UpgradeSelect **Upgrade** and click **Next**.
4. When the warning displays, choose from the following options:
 - If you saved upgrade data during the previous task, continue with the next step.
 - If you did not save upgrade data previously in this procedure, you must save the upgrade data now before you continue.
5. Select **Upgrade** from media and click **Next**. Confirm the settings and click **Finish**.
6. Follow the prompts as they display.

If window blank, press Spacebar: If the window goes blank, press the Spacebar to view the information. The first DVD can take approximately 20 minutes to install.

7. Select option **1. Install additional software from media** and press Enter. Press any key to confirm the installation. The HMC displays status messages as it installs the packages.
8. When the second media installation is complete, remove the media from the drive and close the media drawer.
9. Select option **2. Finish the installation** and press Enter. The HMC completes the booting process.
10. At the login prompt, log in using your user ID and password.

11. Accept the License Agreement for Machine Code twice. The HMC code installation is complete.
12. Verify that the HMC machine code upgrade installed successfully. See “How to determine the HMC software version” on page 316.

You completed upgrading the HMC machine code procedure.

Upgrading HMC from Version 6 to Version 7

This section shows you how to upgrade your HMC Version 6 to HMC Version 7 while you maintain your configuration data.

Important: You must be at a minimum of Version 6 to upgrade to the POWER6 HMC machine code level, which is Version 7 Release 4.

To upgrade from Version 6 to Version 7, follow these steps:

1. Determine the HMC machine code level that is required for your system. See 6.2.4, “Which firmware or fix level is correct for your system” on page 316.
2. Obtain the recovery CD or DVD as described in “Obtaining HMC updates and recovery software” on page 321.

3. Back up the profile data of the managed system. In the HMC workplace window, select **System Management** → **Servers**. Select the server and ensure that the state is *Operating* or *Standby*.

Under Tasks, select **Configuration** → **Manage Partition Data** → **Backup**. Enter a backup file name and record this information. Then, click **OK**.

Repeat these steps for each managed system.

4. Backup Critical Console Data (CCD), HMC Data in HMC 7 Version, as described in 6.2.1, “HMC Data backup” on page 309.

It is necessary to back up the CCD.

5. Before you upgrade to a new version of HMC software, as a precautionary measure, record HMC configuration information as follows:
 - a. In the Navigation area, select **HMC Management**. Then, in the tasks list, select **Schedule Operations**. The Scheduled Operations window displays with a list of all managed systems.
 - b. Select the HMC that you plan to upgrade and click **OK**. All scheduled operations for the HMC are displayed.

Tip: If you do not have any scheduled operations, skip to step 6.

- c. Select **Sort** → **By Object**.
 - d. Select each object and record the following details:
 - Object Name
 - Scheduled date
 - Operation Time (displayed in 24-hour format)
 - RepetitiveIf Yes, select **View** → **Schedule Details**. Then, record the interval information. Close the scheduled operations window. Repeat for each scheduled operation.
 - e. Close the Customize Scheduled Operations window.
6. Record remote command status:
 - a. In the navigation area, select **HMC Management**. Then, in the tasks list click **Remote Command Execution**.
 - b. Record whether the Enable remote command execution using the `ssh` facility check box is selected.
 - c. Click **Cancel**.
 7. Save the upgrade data as described in “Saving upgrade data” on page 330.

Important: If this step is not followed properly, you lose all your partition information.

8. Upgrade the HMC Software from Version 6 to Version 7.

Note: You can upgrade only your HMC from Version 6 to Version 7. If you have an HMC Version 6, you need to upgrade it to Version 7 first. You need to have a recovery DVD from step 2 in this procedure.

- a. Insert the Version 7 recovery DVD in the DVD drive.
- b. In the navigation area, select **HMC Management** → **Shutdown or Restart**. Then, select **Restart the HMC** and click **OK**.
- c. The HMC restarts and boots from the bootable recovery DVD. The window shows the following options:
 - Install
 - Upgrade
- d. Select **Upgrade** and click **Next**.

- e. When the warning displays, choose from the following options:
 - If you saved upgrade data during the previous task, continue with the next step.
 - If you did not save upgrade data previously in this procedure, you must save the upgrade data now before you continue. Refer to previous step.
- f. Select **Upgrade** from media and click **Next**. Confirm the settings and click **Finish**. Follow the prompts as they display.

If window blank, press Spacebar: If the window goes blank, press the Spacebar to view the information. The first DVD can take approximately 20 minutes to install.

- g. Select option **1. Install additional software from media** and press Enter. Press any key to confirm the installation. The HMC displays status messages as it installs the packages. When the second media installation is complete, remove the media from the drive and close the media drawer.
- h. Select Option **2. Finish the installation** and press Enter. The HMC completes the booting process.
- i. At the login prompt, log in using your user ID and password.
- j. Accept the License Agreement for Machine Code twice. The HMC code installation is complete.
- k. Verify that the HMC machine code upgrade installed successfully. See “How to determine the HMC software version” on page 316.

6.2.5 Managed system firmware updates

In this section, we describe different options that are available to install system firmware. The system firmware is also referred to as *licensed internal code*. It is on the service processor.

Important: The HMC machine code needs to be equal to or greater than the managed system firmware level. Also, if an HMC manages multiple servers at different firmware release levels, the HMC machine code level must be equal to or higher than the system firmware level on the server that is at the latest release level.

Firmware overview

Depending on your system model and service environment, you can download, install, and manage your server firmware updates by using different methods.

The default firmware update policy for a partitioned system is through the HMC. If you do not have HMC attached to your system, refer to your operating system documentation to upload the code by using the operating system.

System firmware is delivered as a *Release Level* or a *Service Pack*. Release Levels support the general availability (GA) of new function or features and new machine types or models. Upgrading to a higher Release Level can be disruptive to client operations. Thus, IBM intends to introduce no more than two new Release Levels per year. These Release Levels are supported by Service Packs. Service Packs are intended to contain only firmware fixes and are not intended to introduce new functionality. A Service Pack is an update to an existing Release Level.

Upgrading and updating your firmware: Installing a Release Level is also referred to as *upgrading* your firmware. Installing a Service Pack is referred to as *updating* your firmware.

The file naming convention for System Firmware is as follows:

▶ POWER5

01SFxxx_yyy_zzz

where

- xxx is the release level
- yyy is the service pack level
- zzz is the last disruptive service pack level

So, for example, System Firmware 01SF240_320, as displayed on the Firmware Download page, is Release Level 240, Service Pack 320.

▶ POWER6

EMxxx_yyy_zzz

where

- xxx is the release level
- yyy is the service pack level
- zzz is the last disruptive service pack level

So, for example, System Firmware 01EM310_026, as displayed on the Firmware Download page, is Release Level 310, Service Pack 026.

▶ POWER7

AMxxx_yyy_zzz

where

- xxx is the release level
 - yyy is the service pack level
 - zzz is the last disruptive service pack level
- ▶ So, for example, System Firmware 01QAM730_095, as displayed on the Firmware Download page, is Release Level 730, Service Pack 095.

The Service Pack maintains two copies of the server firmware. One copy is held in the t-side repository (temporary) and the other copy is held in the p-side repository (permanent):

- ▶ **Temporary side:** Apply new firmware updates to the t-side first and test before they are permanently applied. When you install server firmware updates on the t-side, the existing contents of the t-side should be permanently installed on the p-side first.

We recommend that under normal operations the managed system run on the t-side version of the system firmware.

- ▶ **Permanent side:** The permanent side holds the last firmware release that was running on the temporary side. You know that this firmware was running for a while on the temporary side and is stable. This method is also a good way to hold a backup firmware on the system. If for any reason your temporary firmware gets corrupted, you can start from the permanent side and recover your system.

Before you update your system firmware, move current firmware that is on the temporary side to the permanent side.

We recommend that under normal operations the managed system runs on the t-side version of the system firmware.

When you install changes to your firmware, you have three options:

- ▶ **Concurrent installation and activate:** Fixes can be applied without interrupting running partitions and restarting managed system.
- ▶ **Concurrent installation with deferred disruptive activate:** Fixes can be applied as delayed and activated the next time that the managed system is restarted.
- ▶ **Disruptive installation with activate:** Fixes can be applied only by turning off the managed system.

You want to choose the option that fits the status of the server that you are updating. For example, you do not want to use a disruptive installation option on a production server. However, on a test server, this option might not be an issue.

Check compatibility: Always check the compatibility between HMC software and managed system firmware at the following URL:
<http://www.ibm.com/support/fixcentral/firmware/supportedCombinations>

Obtaining system firmware

This section describes how to view or to download the firmware fix. Download the fix to your computer with an Internet connection and then create a fix CD that you apply on the server. If necessary, contact service and support to order the fix on CD.

You can download fixes from the following URL:

<http://www.ibm.com/support/fixcentral/options>

Repeat the same process of “Obtaining HMC updates and recovery software” on page 321 until you arrive at the HMC firmware type, choose **System firmware** (Figure 6-62).

Firmware and HMC
Select fix type

Current selections
→ Your Machine Type-Model is 8233-E8B

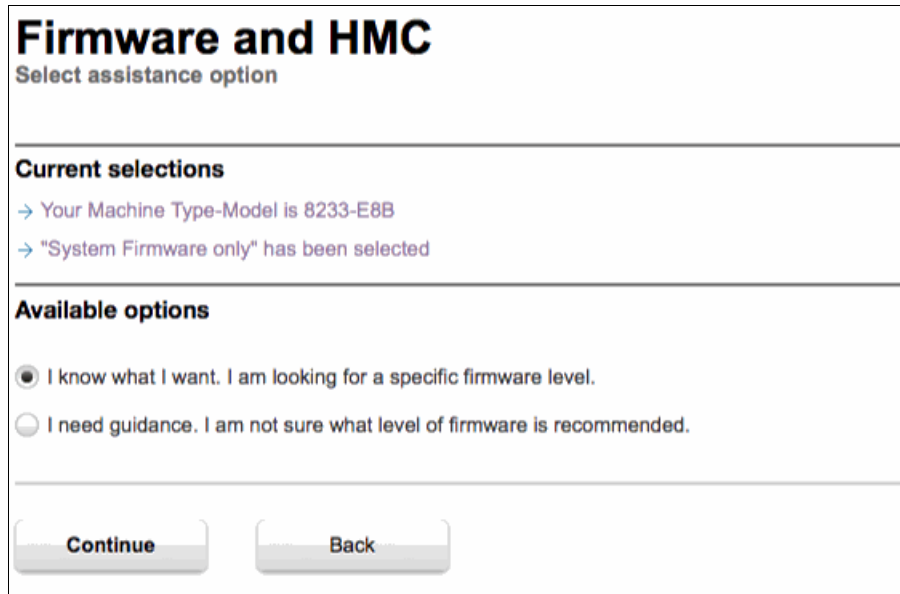
Available options

- All firmware components. Obtain system firmware, device firmware, SDMC, and HMC updates. The power subsystem firmware will be included if applicable.
- System firmware. Obtain system firmware only. The power subsystem firmware will be included if applicable.
- Device firmware. Obtain device firmware only. Available for adapters hard disks and media devices.
- SDMC Code. Obtain SDMC Update Images, Service packs, and Interim fixes.
- HMC Firmware. Obtain HMC Recovery images, service packs, and specific fixes.

Continue **Back**

Figure 6-62 Firmware and HMC select fix type: System firmware

Decide what version of the firmware is correct for your system, as described in 6.2.4, “Which firmware or fix level is correct for your system” on page 316. Use the Fix Level Recommendation Tool to decide the level of firmware that you require for your system and check the compatibility website on page 338, then select **I know what I want** (Figure 6-63). The Fix Central website gives guidance if you are unsure which firmware is recommended.



Firmware and HMC
Select assistance option

Current selections

- Your Machine Type-Model is 8233-E8B
- "System Firmware only" has been selected

Available options

- I know what I want. I am looking for a specific firmware level.
- I need guidance. I am not sure what level of firmware is recommended.

Continue **Back**

Figure 6-63 Firmware and HMC assistance option

Select the version of the system firmware that is applicable to your system. Then, accept the user license agreement and click **Continue** at the bottom of the window to download directly to your workstation or choose FTP (Figure 6-64).

Firmware and HMC

Package download

Obtain package

Download using Download Director (requires Java installation) [What is this?](#)

Download using Bulk FTP

Download ISO image to burn to media

Include informational files, and files required by the HMC, SDMC and Systems Director for installation

Package list

POWER7 System Firmware AL730_099	01AL730_099_035	34963657 bytes
<i>Total</i>		34963657 bytes

Figure 6-64 System firmware download selection

6.3 Advanced System Management Interface

In this chapter, we describe how to set up and use the Advanced System Management Interface (ASMI). The ASMI provides a terminal interface through a standard web browser to the service processor that allows you to do general and administrator level service tasks. The ASMI allows you to do service functions and various system management functions.

6.3.1 Connecting to ASMI

There are three different methods to gain access to the ASMI:

- ▶ Access through the HMC
- ▶ Access through a web browser
- ▶ Access through ASCII terminal

Connection to ASMI by using the HMC

If you have an HMC attached to your managed system, connecting the ASMI by using the HMC is the simplest way to connect. If this is a new system, see 3.1.3, “Connecting managed systems to the HMC” on page 89 for information about how to connect your managed system to the HMC.

To connect to ASMI by using the HMC:

1. In the HMC workplace window, select **System Management** → **Servers**.
2. In the contents area, select the server to which you want to connect the ASMI (Figure 6-65).
3. From the lower panel of the HMC menu, select **Operations** → **Advanced System Management (ASM)**.




Figure 6-65 Launch ASMI from HMC

Connecting to ASMI through a web browser

The web interface to the ASMI is accessible through Microsoft Internet Explorer 7.0, Netscape 9.0.0.4, or Opera 9.24 and Mozilla Firefox 2.0.0.11 running on a PC or notebook that is connected to the service processor. The web interface is available during all phases of system operation, including the initial program load (IPL) and run time. However, some of the menu options in the web interface are unavailable during IPL or run time to prevent usage or ownership conflicts if the system resources are in use during that phase.

To set up the web browser for direct or remote access to the ASMI, complete the following tasks:

1. Connect the power cord from the server to a power source, and wait for the control panel to display *01*.
2. Select a PC or a notebook that has Microsoft Internet Explorer 7.0, Netscape 9.0.0.4, or Opera 9.24 and Mozilla Firefox 2.0.0.11 to connect to your server. You can use this PC or notebook temporarily or permanently to access ASMI.
3. Connect an Ethernet cable from the PC or notebook to the Ethernet port labeled HMC1 on the back of the managed system. If HMC1 is occupied, connect an Ethernet cable from the PC or notebook to the Ethernet port

labeled HMC2 on the back of the managed system. You can use cross-over cable or standard Ethernet cable, both are supported.

4. Configure the Ethernet interface on the PC or notebook to an IP address and subnet mask within the same subnet as the server so that your PC or notebook can communicate with the server. Use Table 6-2 to help you determine these values.

Table 6-2 Default IP address for server connectors HMC1 and HMC2

POWER7 System	Server connector	Subnet Mask	IP address of service processor	Example of your PC IP address
Service processor A	HMC1	255.255.255.0	169.254.2.147	169.254.2.140
	HMC2	255.255.255.0	169.254.3.147	169.254.3.140
Service processor B (if Installed)	HMC1	255.255.255.0	169.254.2.147	169.254.2.140
	HMC2	255.255.255.0	169.254.3.147	169.254.3.140

If you are not sure how to configure your PCs IP settings, then consult your network administrator.

5. Use Table 6-2 to determine the IP address of the Ethernet port to which your PC or notebook is connected, and enter the IP address in the address field of the web browser of your PC or notebook.

For example, if you connected your PC or notebook to HMC1, enter `https://169.254.2.147` in the web browser of your PC or notebook.

Accessing the ASMI by using an ASCII terminal

The ASCII interface to the ASMI provides a subset of the web interface functions. The ASCII terminal is available only when the system is in the platform standby state. It is not available during the IPL or run time. The ASMI on an ASCII terminal is not available during the other phases of system operation, including the IPL and run time.

To set up the ASCII terminal for direct or remote access to the ASMI, complete the following tasks:

1. Use a null modem cable to connect the ASCII terminal to system connector S1 on the back of the server or to system port S1 on the control panel by using an RJ-45 connector.

System port connections: Both system port 1 connections are not available simultaneously; when one is connected, the other is deactivated.

2. Connect the power cord from the server to a power source.
3. Wait for the control panel to display *01*.
4. Ensure that your ASCII terminal is set to the following general attributes.
These attributes are the default settings for the diagnostic programs:
Line Speed-19200,word length-8,parity- none,stop bit-1
5. Press a key on the ASCII terminal to allow the service processor to confirm the presence of the ASCII terminal.
You get the ASMI login window.

6.3.2 Log in to ASMI

To connect successfully to the ASMI, the ASMI requires password authentication.

- ▶ The ASMI provides a Secure Sockets Layer (SSL) web connection to the service processor. To establish an SSL connection, open your browser by using `https://[Your flexible processor IP address]`.
- ▶ The browser-based ASMI is available during all phases of the system operation, including IPL and run time. Some menu options are not available during the system IPL or run time to prevent usage or ownership conflicts if corresponding resources are in use during that phase.
- ▶ The ASMI that is accessed on a terminal is available only if the system is at platform standby.

After you connect to the ASMI as described in 6.3.1, “Connecting to ASMI” on page 341, the login display opens. Enter one of the default user IDs and passwords, as shown in Table 6-3.

Table 6-3 Default login user ID and password

User ID	Default password	Authority level
general	general	general user
admin	admin	administrator
celogin	contact IBM for password	authorized service provider
celogin1	not set, default user disabled	authorized service provider
celogin2	not set, default user disabled	authorized service provider
dev	contact IBM for password	developer user, service only

celogin: celogin1 and celogin2 can be enabled on POWER6 System and later.

When you log in to ASMI, you are asked to change the default password. You are not allowed to proceed unless you change the password.

ASMI login restrictions

The following restrictions apply to ASMI users:

- ▶ Only three users can log in at any one time.
- ▶ If you are logged in and inactive for 15 minutes, your session expires and you have to log in again.
- ▶ If you make five invalid login attempts, your user ID is locked out for five minutes.

The ASMI window that is shown in Figure 6-66 opens after a successful login.

Copyright © 2002, 2012
IBM Corporation.
All rights reserved.

Log out User ID: admin 8233-E8B-SN10DD51P AL730_095

Expand all menus

- Power/Restart Control
- System Service Aids
- System Information
- System Configuration
- Network Services
- Performance Setup
- On Demand Utilities
- Concurrent Maintenance
- Login Profile

Welcome

Machine type-model: 8233-E8B
Serial number: 10DD51P
Date: 2012-10-30
Time: 13:56:57 UTC
Service Processor: Primary (Location: U78A0.001.DNWKCXW-P1)

Current users

User ID	Location
admin	192.168.128.1
admin	192.168.128.1

User Status

User ID	Status
dev	Disabled
celogin	Enabled
celogin1	Enabled
celogin2	Disabled

Figure 6-66 Advanced System Management main menu

6.3.3 Power and restart control

You can use the power and restart control feature to control the system power manually and automatically. In this section, we also describe different options that are available to turn on the system. See Figure 6-67.

We describe the following options in detail:

- ▶ Fast and slow boot
- ▶ Temporary and permanent boot side
- ▶ Normal and permanent operating mode

The screenshot shows the IBM Advanced System Management (ASM) web interface. At the top, the IBM logo and 'Advanced System Management' are displayed. The user is logged in as 'admin' with User ID 'admin', and the system ID is '8233-E8B-SN10DD51P'. The page title is 'Power On/Off System'. The current system power state is 'On', the current firmware boot side is 'Temporary', and the current system server firmware state is 'Running'. The system diagnostic level for the next boot is 'Normal'. The firmware boot side for the next boot is set to 'Temporary'. The system operating mode is set to 'Manual'. The server firmware start policy is set to 'Running (Auto-Start Always)'. The system power off policy is set to 'Stay on'. The default partition environment is set to 'AIX'. There are buttons for 'Save settings' and 'Save settings and power off'.

Figure 6-67 System power on and off options

Power On/Off System option

When you select this option, the right side of the menu gives you the current system power state, current firmware boot side, and system server firmware state. You can select the following boot options:

- ▶ **System boot speed:** Select the speed for the next boot (*Fast* or *Slow*). A fast boot results in some diagnostic tests being skipped, and shorter memory tests being run during the boot. A slow boot goes through all diagnostic tests

and memory tests. Normally, you take this option if you are experiencing some system errors or when you made some system changes, for example CPU and memory upgrades.

- ▶ **Firmware boot side:** Select the side from which the firmware boots: permanent or temporary. When you upgrade your system firmware, typically, firmware updates are tested on the temporary side before they are applied to the permanent side. Therefore, the temporary side should always have the latest firmware. The permanent side has the previous revision.
- ▶ **System operating mode:** Select the operating mode (*Manual* or *Normal*). Manual mode overrides various automatic power-on functions, such as auto-power restart, and enables the power button, which allows you to select power options from the control panel. You can also set this option from the control panel.
- ▶ **Boot to system server firmware:** Select the state for the server firmware: Standby or Running. When the server is in the server firmware standby state, partitions can be set up and activated. The running option restarts your partitions automatically.
- ▶ **System power off policy:** Select the system power off policy. The system power off policy is a system parameter that controls the behavior of the system when the last partition (or the only partition in the case of a system that is not managed by an HMC) is powered off. The choices are:
 - Power off: When the last partition is powered down, the system turns off.
 - Stay on: When the last system is powered down, the system stays on.
 - Automatic: Is the default setting. If the system is not partitioned, the system is turned off. If the system is partitioned, it stays on.

Make your selections and select **Save settings and power on**.

Auto Power Restart function

You can set your system to restart automatically. This function is useful when power is restored after an unexpected power line disturbance causes the system to shut down unexpectedly. Select either **Enable** or **Disable**. By default, the auto power restart value is set to *Disable*. In many cases, you might not want the system to restart automatically, unless you are reasonably certain that the power problem is resolved.

Immediate Power Off function

You can power off the system quickly by using the Immediate Power Off function. Typically, this option is used when an emergency power off is needed. The operating system is not notified before the system is powered off.

Attention: To avoid experiencing data loss and a longer IPL the next time the system or logical partitions are booted, shut down the operating system before doing an immediate power off.

System Reboot function

You can reboot the system quickly by using the reboot function. The operating system is not notified before the system is rebooted.

Rebooting: Rebooting the system shuts down all partitions immediately. To avoid experiencing data loss and a longer IPL the next time the system or logical partitions are booted, shut down the operating system before doing a reboot.

6.3.4 System Service Aids menu

Figure 6-68 shows the System Service Aids menu. From this menu, you can do the following functions:

- ▶ Display system error, event logs.
- ▶ Initiate a system dump.
- ▶ Initiate a service processor dump.
- ▶ Reset the service processor.
- ▶ Reset your system to the factory-shipped configuration settings.

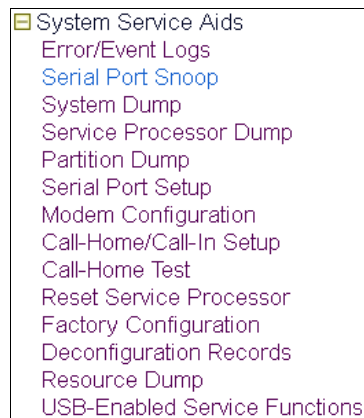


Figure 6-68 System Service Aids menu

The following features are not available when your system is connected to the HMC. These features are part of the Service Management on the HMC:

- ▶ Serial port snoop

- ▶ Partition dump
- ▶ Serial port setup
- ▶ Modem configuration
- ▶ Call home/call in setup
- ▶ Call home test

Error/Event Logs option

From the System Service Aids menu, select **Error/Event Logs**. You can view error and event logs that are generated by various service processor firmware components. The content of these logs can be useful in solving hardware or server firmware issues. You see a selection panel, as shown in Figure 6-69.

Error/Event Logs					
Serviceable/Customer attention events					
√	Log ID	Time	Failing subsystem	Severity	SRC
<input type="checkbox"/>	5303931D	2012-10-26 21:20:54	System Hypervisor Firmware	Predictive Error	B7005191
<input type="checkbox"/>	5303931C	2012-10-26 21:20:50	System Hypervisor Firmware	Predictive Error	B7005191
<input type="checkbox"/>	53039317	2012-10-26 21:20:35	System Hypervisor Firmware	Predictive Error	B7005191
<input type="checkbox"/>	53039316	2012-10-26 21:20:31	System Hypervisor Firmware	Predictive Error	B7005191

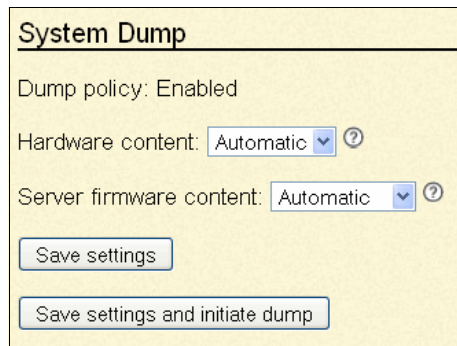
Figure 6-69 Error and event logs

Select the event log that you want to view and scroll to the bottom of the window to select **show details**. The details provide the description of the system reference code (SRC).

System Dump procedure

Use the System Dump procedure only under the direction of your service provider. You can initiate a system dump to capture overall system information, system processor state, hardware scan rings, caches, and other information. This information can be used to resolve a hardware or server firmware issue. A system dump can also be initiated automatically after a system malfunction, such as a check stop or hang.

Select **System Dump** to open the window that is shown in Figure 6-70.



The screenshot shows a window titled "System Dump" with a yellow background. The text "Dump policy: Enabled" is displayed. Below it are two dropdown menus: "Hardware content: Automatic" and "Server firmware content: Automatic", each with a question mark icon to its right. At the bottom of the window are two buttons: "Save settings" and "Save settings and initiate dump".

Figure 6-70 Capturing overall system information with system dump procedure

From this window, you set the following information:

- ▶ **Dump policy:** Select the policy to determine when system dump data is collected. If you select **Enable**, the service processor (SP) collects system dump data only when the SP determines it is necessary, typically only when a specific failure is not identified.

If you select **Disable**, the SP never collects any dump data, unless explicitly requested by the user.

The default is *Enabled*.

- ▶ **Hardware content:** Select the policy to determine how much hardware data is collected for a system dump. If you select **Automatic** (default), the SP collects the hardware data that it determines is necessary, depending on the particular failure.

If you select **Maximum**, the SP collects the maximum amount of hardware data. If you choose this selection, the collection of hardware data can be quite time consuming, especially for systems with many processors.

- ▶ **Server firmware content:** Select the policy to determine how much server firmware data is collected for a system dump. If you select **Automatic**, the SP collects the minimum amount of data necessary to debug server firmware failures. Automatic is the default policy. In some cases, your support engineer might want you to override the default policy.

If you select **Physical I/O**, the SP collects the minimum firmware data plus the firmware data that is associated with physical I/O operations.

If you select **Virtual I/O**, the SP collects the minimum firmware data plus the firmware data that is associated with I/O operations that do not involve physical I/O devices.

If you select **High performance switch HPS Cluster**, the SP collects the minimum firmware data plus the firmware data that is associated with high performance switch operations between this server and other servers in the cluster.

If you select **HCA I/O**, the SP collects the minimum firmware data plus the firmware data associated with the host channel adapter I/O operations.

If you select **Maximum**, the SP collects the maximum amount of server firmware data.

Make your selections and click **Save settings**.

Service Processor Dump option

You use the Service Process Dump option to enable or disable the service processor dump function. The default value is *Enabled*. A service processor dump captures error data after a service processor failure, or upon user request. User request for service processor dump is not available when this policy is set to disabled.

The save settings and initiate dump button is visible only when an SP dump is allowed (that is, when SP dumps are enabled and the previous SP dump data is retrieved). Press this button to initiate an SP dump.

Reset Service Processor option

Typically, rebooting of the SP is done only when instructed by IBM service personnel (Figure 6-71).

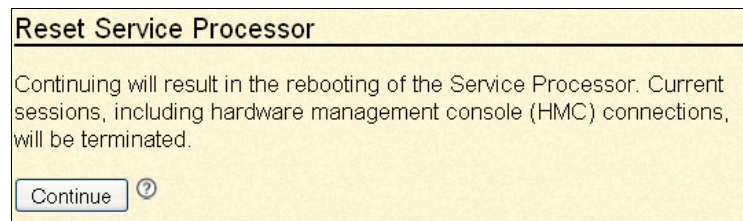


Figure 6-71 Reset service processor

This function is not available if your system is turned on. Clicking **Continue** causes the service processor to reboot. Because the service processor reboots, your ASMI session is dropped, and you have to reconnect your session to continue.

Factory Configuration option

Use this procedure only under the direction of your IBM service personnel (Figure 6-72).

In critical systems situations, you can restore your system to the factory default settings. Doing so results in the loss of all system settings (such as the HMC access and ASMI passwords, time of day, network configuration, and hardware deconfiguration policies) that you have to set again through the service processor interfaces. Also, you lose the system error logs and partition-related information.

Manually record all settings: Before you continue with this operation, ensure that you have manually recorded all settings that have to be preserved.

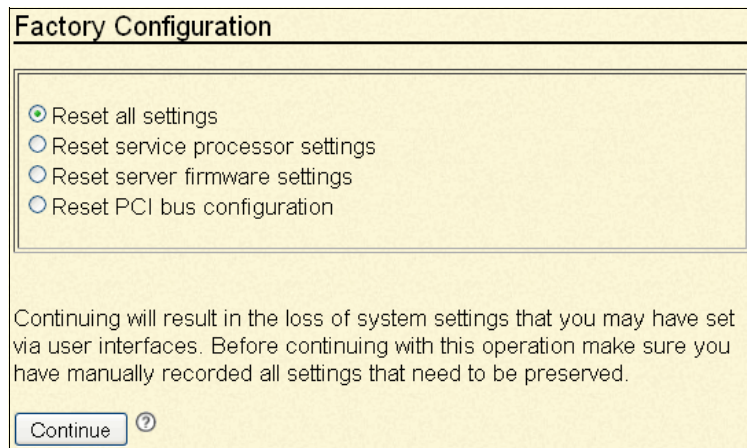


Figure 6-72 Factory configuration reset

In this window, you have the following options:

- ▶ **Reset all settings:** Resets everything. It is a combination of all the others. To complete this operation, the system is powered on and then off, and the service processor is reset.
- ▶ **Reset service processor settings:** Resets the settings of the service processor that include passwords, network addresses, time of day, hardware configuration policies, and so forth. Any sessions that are currently active in the network interfaces are disconnected, and the service processor is reset.
- ▶ **Reset server firmware settings:** Resets the firmware settings only. Partition data is lost.

- ▶ **Reset Peripheral Component Interconnect (PCI) bus configuration:**
Resets the PCI bus and the firmware settings. To complete this operation, the system is turned on and then off.

Make the appropriate selection and then select **Continue**.

6.3.5 System Information menu

The System Information menu gives you the following options (Figure 6-73):

- ▶ Display vital product data.
- ▶ Perform a system power control network (SPCN) power control network trace and display the results.
- ▶ Display the previous boot indicator.
- ▶ Display the progress indicator history.
- ▶ Display the real-time progress Indicator


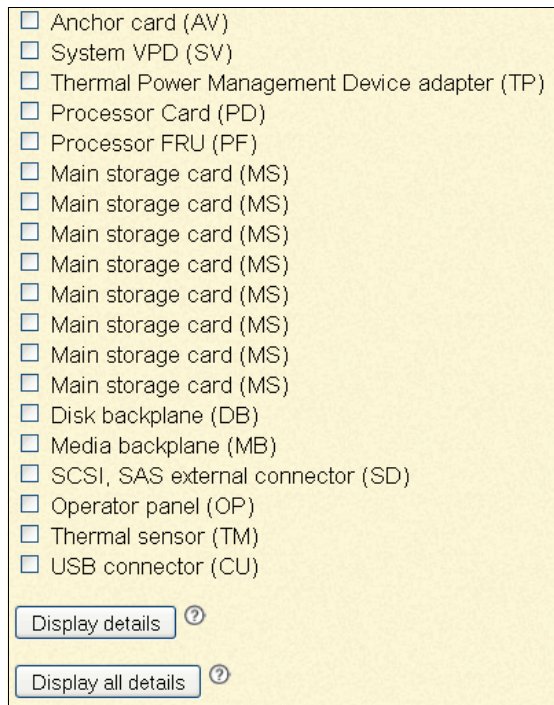


Figure 6-73 System Information menu

Vital Product Data option

Select **Vital Product Data** to view manufacturer's vital product data (VPD) that is stored from the system boot before the one in progress now (Figure 6-74).



- Anchor card (AV)
- System VPD (SV)
- Thermal Power Management Device adapter (TP)
- Processor Card (PD)
- Processor FRU (PF)
- Main storage card (MS)
- Main storage card (MS)
- Main storage card (MS)
- Main storage card (MS)
- Main storage card (MS)
- Main storage card (MS)
- Main storage card (MS)
- Main storage card (MS)
- Main storage card (MS)
- Main storage card (MS)
- Disk backplane (DB)
- Media backplane (MB)
- SCSI, SAS external connector (SD)
- Operator panel (OP)
- Thermal sensor (TM)
- USB connector (CU)

Display details ?

Display all details ?

Figure 6-74 Display details of VPD

If you want to view only selected manufacturer's VPD, such as serial numbers and part numbers, select the feature that you want to view and select **Display details**.

To view details of all the features, select **Display all details** to open the display that is shown in Figure 6-75.

Vital Product Data						
Build name: fips730/b0810a_1260.730						
System brand: S0						
System serial number: 10DD51P						
Machine type-model: 8233-E8B						
FRU ID	Part number	Serial number	FRU number	CCIN	RID	Location code
EV	74Y3758	YL11HA17107L	74Y3757	2A5C	0x1e00	U78A0.001.DNWKCXW
EI	74Y3758	YL11HA17107L	74Y3757	2A5C	0xa200	U78A0.001.DNWKCXW
EF	74Y3758	YL11HA17107L	74Y3757	2A5C	0xa300	U78A0.001.DNWKCXW
BP	74Y3758	YL11HA17107L	74Y3757	2A5C	0x800	U78A0.001.DNWKCXW-P1
CU	74Y3758	YL11HA17107L	74Y3757	2A5C	0x2900	U78A0.001.DNWKCXW-P1-T3
CU	74Y3758	YL11HA17107L	74Y3757	2A5C	0x2901	U78A0.001.DNWKCXW-P1-T4

Figure 6-75 Displays all VPD detail

Power control network trace

You can perform an SPCN trace and display the results. This information is gathered to provide extra debug information when you work with your hardware service provider.

Note: Producing a trace can take an extended amount of time that is based on your system type and configuration. This process is a normal delay because of the amount of time the system requires to query the data.

Previous boot progress indicator

You can display the previous boot progress indicator that was displayed in the control panel during the previous failed boot by selecting this option. During a successful boot, the previous progress indicator is cleared. If this option is selected after a successful boot, nothing is displayed.

The progress indicator information is stored in nonvolatile memory. If the system is powered off using the power-on button on the control panel, this information is retained. If the alternating current (ac) power is disconnected from the system, this information is lost.

Progress Indicator History option

With this option, you can review the progress of codes that displays in the control panel during the previous boot. The codes display in reverse chronological order, as shown in Figure 6-76. (The first entry that is seen is the most recent entry.) This information is gathered to provide extra debug information when you work with your hardware service provider.

Progress Indicator History		
Listed in reverse chronological order.		
<input checked="" type="checkbox"/>	Progress Indicator	Time
<input type="checkbox"/>	RUNTIME	2012-10-26 18:24:03
<input type="checkbox"/>	STANDBY	2012-10-26 18:24:01
<input type="checkbox"/>	C7004091	2012-10-26 18:24:00
<input type="checkbox"/>	C7004091	2012-10-26 18:24:00
<input type="checkbox"/>	C7004087	2012-10-26 18:24:00
<input type="checkbox"/>	C7004080	2012-10-26 18:24:00

Figure 6-76 Progress Indicator History option

Select the code that you want to display and select show details.

You can view the progress and error codes that currently display on the control panel. Viewing progress and error codes is useful when you diagnose boot-related issues. To perform this operation, your authority level must be one of the following possibilities:

- ▶ General
- ▶ Administrator
- ▶ Authorized service provider

Select this option to open the window that is shown in Figure 6-77. This window shows the real-time progress of the system and displays what you have on the system display.

01	M	V=N
HMC=1		T

Figure 6-77 Real-time Progress Indicator

6.3.6 System Configuration menu

Figure 6-78 shows the expanded System Configuration menu. Using this menu, you can do the following functions:

- ▶ Change the system name.
- ▶ Configure I/O enclosure.
- ▶ Change the time of day.
- ▶ Establish the firmware update policy.
- ▶ Establish the detailed PCI error injection policies.
- ▶ Change the interposer plug count.
- ▶ Enable I/O adapter enlarged capacity.
- ▶ View Hardware Management Console connections.
- ▶ Change floating point unit commutation test values.

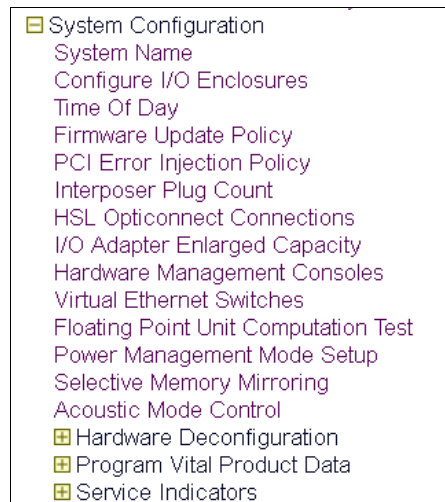


Figure 6-78 System Configuration menu

System Name option

From the System Configuration menu, you can select the system name option to display the current system name and change the system name if you choose to do so. The system name is a value that is used to identify the system or server. The system name might not be blank and might not be longer than 31 characters. To change the system name, enter a new value and click **Save settings**.

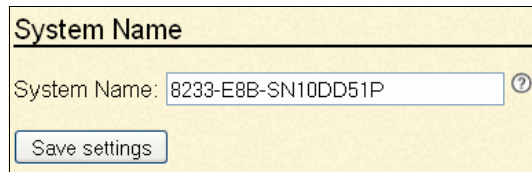


Figure 6-79 System Name

This example shows changes to the system name. The valid characters are [a-Z], [0-9], hyphen (-), underscore (_), and period (.).

The system is shipped with the default system name initialized to a 31 character value as follows (Server-*ttt*-*mmm*-SN0000000). In this default system name:

- ▶ *ttt* = Machine type
- ▶ *mmm* = Model number
- ▶ 0000000 = Serial number

Configure I/O Enclosures function

This function normally is used by your hardware service provider. After the server firmware reaches the *standby* state, you can configure I/O enclosure attributes as follows:

- ▶ Display the status, location code, rack address, unit address, power control network identifier, and the machine type and model of each enclosure in the system.
- ▶ Change the identification indicator state on each enclosure to *on* (identify) or *off*.
- ▶ Update the power control network identifier, enclosure serial number, and the machine type and model of each enclosure.
- ▶ Change the identification indicator state of the SPCN firmware in an enclosure to *enable* or *disable*.
- ▶ Remove rack and unit addresses for all inactive enclosures in the system.

When you select this option, the window that is shown in Figure 6-80 opens.

Configure I/O Enclosures										
Enclosure Configuration										
	Status	Rack address	Unit address	Power Control Network Identifier	Power Control Network Firmware Update Status	Power Control Network Firmware Version	Start Time	Type - Model	Serial number	Location code
<input type="radio"/>	Active	0x3C00	0x1	0xF0	Not Applicable			78A0-001	DNWKCXW	U78A0.001.DNWKCXW
<input type="button" value="Identify enclosure"/> ? <input type="button" value="Turn off indicator"/> ? <input type="button" value="Change settings"/> ? <input type="button" value="Collect SPCN IO trace"/> ?										
Enclosure Options										
<input type="button" value="Clear inactive enclosures"/> ? <input type="button" value="Start SPCN firmware update"/> ? <input type="button" value="Stop SPCN firmware update"/> ? <input type="button" value="SPCN loop status"/> ?										

Figure 6-80 Configure I/O Enclosures window

In this window, you have the following options:

- ▶ **Identify enclosure:** Turns on the indicator on the selected enclosure. LED flashes to identify the enclosure.
- ▶ **Turn off indicator:** Turns off the indicator on the selected enclosure.
- ▶ **Change settings:** Changes the settings for the selected enclosure. The next page displays options for changing the configuration ID, machine type-model, and serial number.
 - **Power Control Network Identifier:** Enter a hexadecimal number for the power control network identifier.

System server firmware: The system server firmware must be in *standby* state or the expansion unit must be turned off when this operation is performed.

- **Type - Model:** Enter the enclosure machine type and model in the form *TTTT-MMM*:

TTTT: The four characters of the enclosure machine type.

MMM: The three characters of the enclosure model.

The enclosure machine type cannot be 0000. All alphanumeric characters are valid.

- **Serial number:** Enter seven characters for the enclosure serial number. All alphanumeric characters except o, i, and q are valid. All lowercase letters are converted to uppercase letters.

- ▶ **Collect SPCN I/O Trace:** Displays SPCN I/O trace for selected enclosure.

Note: The remaining options described here do not display in Figure 6-79. The options are not displayed because that window capture is a partial window capture of the Configure I/O Enclosure panel.

- ▶ **Clear inactive enclosures:** Clears the rack and unit addresses of all inactive enclosures.
- ▶ **Start SPCN firmware update:** Starts pending SPCN firmware downloads if allowed by the SPCN firmware update policy. SPCN firmware downloads cannot all be attempted at the same time. Some downloads can remain in a pending state before starting while others complete. Starting SPCN downloads is done asynchronously.
- ▶ **Stop SPCN firmware update:** Stops SPCN firmware downloads that are currently in progress. SPCN firmware downloads that are stopped move to a pending state. These SPCN firmware downloads can be restarted from the beginning either automatically by the system or by using Start SPCN Firmware Update, if allowed by the SPCN firmware update policy.

Stopping the SPCN downloads is done asynchronously and can be monitored showing the power control network firmware update status.

- ▶ **SPCN firmware update policy:** If *Disabled*, no SPCN firmware downloads are allowed to start. Changing the SPCN firmware update policy to disabled does not affect SPCN firmware downloads currently in progress.

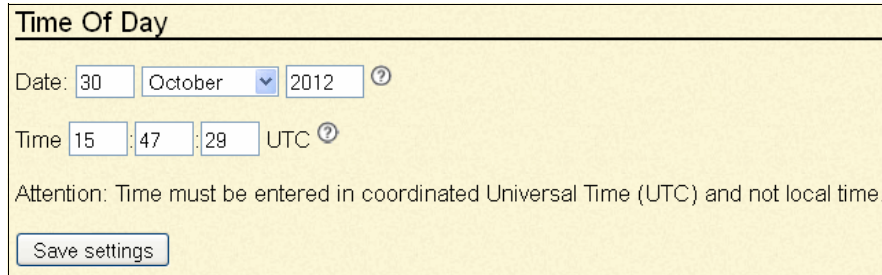
If *Enabled*, SPCN firmware downloads are allowed only over the high-speed link (HSL) interface. Changing the SPCN firmware update policy to enabled does not affect SPCN firmware downloads over the serial interface that are currently in progress.

Changing the SPCN firmware update policy setting to *Enabled* from *Disabled* does not automatically cause SPCN firmware downloads over the HSL interface to begin immediately.

If *Expanded*, SPCN firmware downloads are allowed over both the HSL and serial interfaces. Changing to the SPCN firmware update policy to *expanded* does not cause SPCN firmware downloads to begin immediately.

Time of Day function

You can display and change the current date and time of the system. This function is available if your system is turned on or off. When you select this option, the window that is shown in Figure 6-81 opens.



The screenshot shows a window titled "Time Of Day" with a yellow background. It contains the following elements:

- Date:** Three input fields containing "30", "October" (a dropdown menu), and "2012". A question mark icon is to the right.
- Time:** Three input fields containing "15", ":47", and ":29", followed by "UTC" and a question mark icon.
- Attention:** A text line stating "Attention: Time must be entered in coordinated Universal Time (UTC) and not local time."
- Save settings:** A button at the bottom left.

Figure 6-81 Time of Day function

From this window, you have the following options:

- ▶ **Date:** Enter the current date. Any change to the current date or time is applied to the service processor only, and is independent of any partition.
- ▶ **Time:** Enter the current time in Coordinated Universal Time (UTC) format. UTC is the current term for what was commonly referred to as Greenwich mean time (GMT). Zero (0) hours UTC is midnight in Greenwich, England, which lies on the zero longitudinal (or prime) meridian.

Universal time is based on a 24 hour clock. Local time is expressed as a positive or negative offset from UTC, depending on whether the local time zone is east or west of the prime meridian.

To convert local time to UTC, use 24 hour notation, then algebraically add the time zone offset to the local time. For instance, a user in the Central Daylight-savings Time zone (CDT) adds 5 hours (UTC offset -5 hours) to the local time to obtain the time in UTC. Example: 07:00 PM CDT equals 00:00 UTC.

Enter the date and time and select **Save settings**.

Firmware Update Policy

This policy defines whether firmware updates are allowed from an operating system when the system is managed by an HMC. The default setting of this policy is to not allow firmware updates through the operating system. This policy

takes effect only when a system is HMC managed. When a system is not HMC-managed, firmware updates can be made only through the operating system, so this policy setting is ignored.

When this policy is set to allow firmware updates from the operating system, firmware updates from an HMC are not allowed, unless the system is turned off.

When a system is turned off, firmware updates can be performed from an HMC, regardless of the setting of this policy. However, take care when you update firmware from both an HMC and the operating system.

When you select this option, the window that is shown in Figure 6-82 opens.

Firmware Update Policy

Update Policy: Hardware management console (HMC) ?

Note: This policy is only applicable in certain system configurations. Some system configurations may cause the firmware to override this policy.

When the system is powered off: Firmware update is possible only from the HMC.

When the operating system is running: Firmware update is allowed only from the HMC.

Save settings

Figure 6-82 Firmware Update Policy before POWER7

Some Power Systems servers require an HMC for firmware updates, as shown in Figure 6-83.

Firmware Update Policy

Update Policy: Hardware management console (HMC)

Note: This policy is only applicable in certain system configurations. Some system configurations may cause the firmware to override this policy.

When the system is powered off: Firmware update is possible only from the HMC.

When the operating system is running: Firmware update is allowed only from the HMC.

Figure 6-83 Firmware Update Policy in certain models of POWER7 Systems

PCI Error Injection Policy option

This option controls the PCI error injection policy. If enabled, utilities on the host operating system can inject PCI errors.

I/O Adapter Enlarged Capacity option

This option controls the size of PCI memory space that is allocated to each PCI slot. When enabled, selected PCI slots, including those in external I/O subsystems, receive the larger direct memory access (DMA) and memory mapped address space. Some PCI adapters might require this additional DMA or memory space, per the adapter specification. This option increases system main storage allocation to these selected PCI slots.

Enabling this option might result in some PCI host bridges and slots not being configured because the installed main storage is insufficient to configure all installed PCI slots.

Hardware Management Consoles

You can use this option to view the HMC that is connected or was connected to the managed system. See Figure 6-84. From this menu, you can also remove the disconnected HMC from your managed system. Select the HMC serial number and click **Remove Connection**.

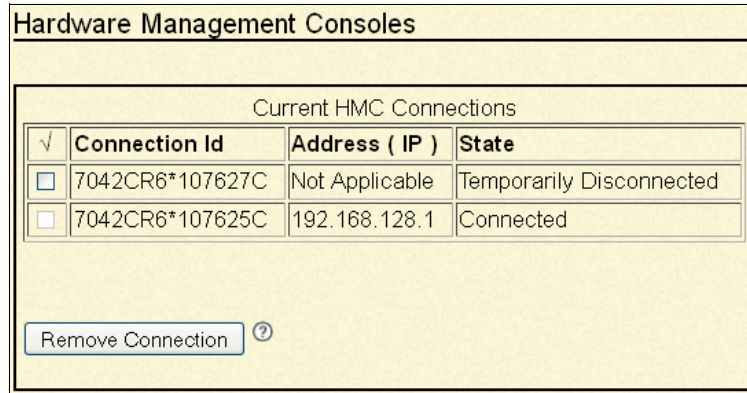


Figure 6-84 Hardware Management Consoles window

Click **Remove Connection**.

Virtual Ethernet Switches option

To use this option, enter a number between 0 to 16 for virtual Ethernet switches. This value controls the number of virtual Ethernet switches that are allocated by system server firmware. Most users leave this value set to its default of 0. A value

of 0 enables the HMC to control the number of virtual Ethernet switches that are allocated by system server firmware.

For advanced configuration, this number can be set higher to cause the system server firmware to create that many virtual Ethernet switches during platform power-on. It also disables the ability of the HMC to configure the number of virtual Ethernet switches.

If this process is done, when a virtual Ethernet adapter is created by using the HMC, the adapter is connected to a particular virtual switch depending on the virtual slot number that is chosen during creation.

The adapter's virtual slot number is divided by the number of virtual Ethernet switches. The remainder of this division operation is used to determine with which switch the adapter is associated.

Each virtual Ethernet adapter is able to communicate only with other virtual Ethernet adapters on the same virtual switch. For example, if the number of virtual Ethernet switches is set to 3, virtual Ethernet adapters in virtual slot 3, 6, and 9 are assigned to the same switch. A virtual Ethernet adapter in virtual slot 4 is assigned to a different switch, and would not be able to communicate with the adapters in slots 3, 6, and 9.

Floating Point Unit Computation Test option

This option allows you to set the floating point unit test policy or to run the test immediately. You can set one of the following functions:

- ▶ **Disabled:** Test never runs except when choosing to run the test immediately.
- ▶ **Staggered:** Test is run once on every processor in the platform over a 24 hour period. Staggered is the default setting.
- ▶ **Periodic:** Test runs at a specified time, sequentially through all processors in the system.

When choosing to run the test immediately, the current policy setting is overridden but not changed. The test is run sequentially on all the processors in the system. This feature is only available when the system is turned on.

Hardware Deconfiguration option

You can set various policies to deconfigure processors and memory in certain situations (Figure 6-85). *Deconfiguration* means that the resource is taken from a state of being available to the system, to a state of being unavailable to the system.

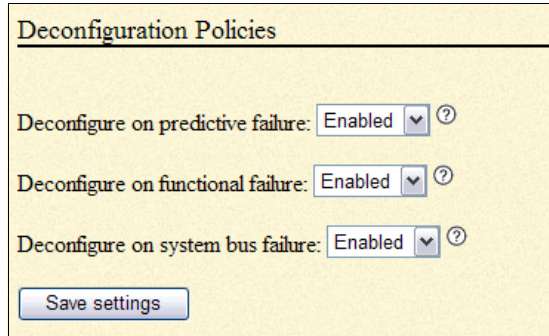


Figure 6-85 Deconfiguration Policies window

From this window, you have the following options:

- ▶ **Deconfigure on predictive failure:** Select the policy for deconfigure on predictive failures. This configuration applies to run time or persistent boot time deconfiguration of processing unit resources or functions with predictive failures, such as correctable errors over the threshold.

If enabled, the particular resource or function that is affected by the failure is deconfigured.

- ▶ **Deconfigure on functional failure:** Select the policy for deconfigure on functional failures. This configuration applies to run time or persistent boot time deconfiguration of processing unit resources or functions with functional failures, such as check stop errors or uncorrectable errors.

If enabled, the particular resource or function that is affected by the failure is deconfigured.

- ▶ **Deconfigure on system bus failure:** Select the policy for deconfigure on system bus failures. Applies to run time or persistent boot time deconfiguration of processing unit resources or functions with system bus failures, such as check stop errors or uncorrectable errors.

This policy is not applicable for systems with one processing unit node. If enabled, the particular resource or function that is affected by the failure is deconfigured.

This configuration applies to resource types such as processor, L2 cache, L3 cache, and memory.

Processor deconfiguration

In the event of a single processor failure, it might be possible to continue operating, with degraded performance, on fewer processors. You can use the panel that is shown in Figure 6-86 on page 367 to start the process of removing processors that might have failed or are beginning to generate errors. You can also see processors that might have become deconfigured because of some error condition that the system was able to detect and isolate.

All processor failures that stop the system, even if intermittent, are reported to the authorized service provider as a diagnostic dial-out for a service repair action. To prevent the recurrence of intermittent problems and improve the availability of the system until a scheduled maintenance window can be found, processors with a failure history are marked *deconfigured* to prevent them from being configured on subsequent boots. Processors marked as deconfigured remain offline and are omitted from the system configuration.

A processor is marked *deconfigured* under the following circumstances:

- ▶ If a processor fails built-in self-test or power-on self-test testing during boot (as determined by the service processor).
- ▶ If a processor causes a machine check or check stop during run time, and the failure can be isolated specifically to that processor (as determined by the processor runtime diagnostics in the service processor firmware).
- ▶ If a processor reaches a threshold of recovered failures that results in a predictive call to the service provider (as determined by the processor runtime diagnostics in the service processor firmware).

The deconfiguration policy also provides the user with the option to manually deconfigure a processor or re-enable a previous manually deconfigured processor.

To begin the process, use a panel similar to the one shown in Figure 6-86. Select the processing unit with which you want to work (one or more processing units can be shown) and click **Continue**.

Processor Deconfiguration

Total system processors: 4

Total system configured processors: 2

Total system deconfigured processors: 2

	Processing unit	Total processors	Configured	Deconfigured
<input type="radio"/>	0	4	2	2

Figure 6-86 Processor Deconfiguration window

Select the setting to configure or deconfigure for the processors and select **Save settings**.

Processor Deconfiguration

Processing unit: 0

Processor ID	Location code	State	Error type	Change settings
0	U789D.001.DQDVWZK-P2-C1	Configured	None (0)	Configured <input type="button" value="v"/> ?
1	U789D.001.DQDVWZK-P2-C1	Configured	None (0)	Configured <input type="button" value="v"/> ?
2	U789D.001.DQDVWZK-P2-C2	Manually deconfigured	None (0)	Deconfigured <input type="button" value="v"/> ?
3	U789D.001.DQDVWZK-P2-C2	System deconfigured	Predictive (E9)	Deconfigured <input type="button" value="v"/> ?

Figure 6-87 Processor Deconfiguration window

Memory deconfiguration

Most System POWER6 systems have several gigabytes (GB) of memory. Each memory bank contains two dual inline memory modules (DIMMs). If the firmware detects a failure, or predictive failure of a DIMM, it deconfigures the DIMM with the failure, and the other one. All memory failures that stop the system, even if intermittent, are reported to the authorized service provider as a diagnostic dial-out for a service repair action.

To prevent the recurrence of intermittent problems and improve the availability of the system until a scheduled maintenance window can be found, memory banks with a failure history are marked *deconfigured*. This status prevents them from being configured on subsequent boots. Memory banks marked as deconfigured remain offline and are omitted from the system configuration.

A memory bank is marked *deconfigured* under the following circumstances:

- ▶ If a memory bank fails built-in self-test or power-on self-test testing during boot (as determined by the service processor).
- ▶ If a memory bank causes a machine check or check stop during run time, and the failure can be isolated specifically to that memory bank (as determined by the processor runtime diagnostics in the service processor firmware).
- ▶ If a memory bank reaches a threshold of recovered failures that results in a predictive call to the service provider (as determined by the processor runtime diagnostics in the service processor firmware).

The deconfiguration policy also provides the user with the option to manually deconfigure a memory bank or re-enable a previous manually deconfigured memory bank.

If you select **Memory Deconfiguration** from the **Hardware Configuration** menu, you see a panel similar to the one shown in Figure 6-88 on page 369. This window allows you to view the total memory that is installed on your system. From this panel, you can select the processing unit (one or more processing units can be shown). The reason that you see *processing unit* is because the

memory is installed on the processor board. Click **Continue** to advance to the next panel. See Figure 6-88, which shows the Memory Deconfiguration window.

Memory Deconfiguration

Total system memory: 12288 MB

Total system configured memory: 11776 MB

Total system deconfigured memory: 512 MB

	Processing unit	Total memory	Configured	Deconfigured
<input type="radio"/>	0	12288 MB	11776 MB	512 MB

Figure 6-88 Memory Deconfiguration window

A new panel similar to the one shown in Figure 6-89 displays. You can then see any memory banks that might have become deconfigured because of some error condition that the system was able to detect and isolate. You can select either configured or deconfigured for each memory bank and select **Save settings**.

Memory Deconfiguration					
Processing unit: 0					
Memory dimm	Location code	Size	State	Error type	Change settings
0	U789D.001.DQDVWZK-P2-C1-C6	512 MB	Configured	None (0)	Configured <input type="button" value="v"/> ?
1	U789D.001.DQDVWZK-P2-C1-C3	512 MB	Manually deconfigured	None (0)	Deconfigured <input type="button" value="v"/> ?
2	U789D.001.DQDVWZK-P2-C1-C9	512 MB	Configured	None (0)	Configured <input type="button" value="v"/> ?
3	U789D.001.DQDVWZK-P2-C1-C12	512 MB	Configured	None (0)	Configured <input type="button" value="v"/> ?

Figure 6-89 Memory deconfiguration memory bank selection

Program Vital Product Data

The ASMI allows you to program the system (VPD such as system brand, system identifiers, and system enclosure type (Figure 6-90). To access any of the VPD-related panels, your authority level must be *administrator* or *authorized service provider*.

Starting the system: You cannot boot the system until valid values are entered for the system brand, system identifiers, and system enclosure type.

<ul style="list-style-type: none"> Program Vital Product Data System Brand System Keywords System Enclosures
--

Figure 6-90 Program Vital Product Data panel

System Brand

Enter a two-character brand type. The first character must be one of the following characters:

- D IBM Storage
- I IBM System i
- N OEM IBM System i only
- O OEM IBM System p only
- P IBM System p

The second character is reserved. A value of zero means that there is no specific information that is associated with it. This entry is write-once only, except in the case where it is all blanks, or when changing from a System p system to an IBM Storage system. Any other changes are disallowed. A valid value is required for the machine to boot. Additionally, for IBM Storage, each of the systems that constitutes the storage facility must have the first character set to D for storage to be accessible online.

System Keywords display

You can set the system-unique ID, serial number, machine type, and machine model (Figure 6-91). If you do not know the system-unique ID, contact your next level of support.

System Keywords
Machine type-model: 8233-E8B
System serial number: 10DD51P
System unique ID: 0004AC171F2F
World Wide Port Name Index: C05076045B76
RB Keyword0: 1

Figure 6-91 System Keyword display example

Machine type-model

Enter a machine type and model in the form *TTTT-MMM*, where *TTTT* is the 4-character machine type and *MMM* is the 3-character model. A valid value is required for the machine to boot. Additionally, for storage to be accessible online, this value must match exactly both systems that constitute the storage facility. This entry is write-once only.

System serial number

Enter a system serial number in the form *XXYYYYYY*, where *XX* is the code for the plant of manufacture and *YYYYYY* is the unit sequence number. Valid characters are 0 - 9 and A - Z. A valid value is required for the machine to boot. This entry is write-once only.

System unique ID

Enter a system-unique serial number as 12 hexadecimal digits. The value should be unique to a specific system anywhere in the world. A valid value is required for the machine to boot.

Worldwide port name

Enter a 16-digit hexadecimal number for the worldwide node name. This value is an IEEE-assigned 64-bit identifier for the storage facility. A valid value is required for the machine to boot. This entry is write-once only.

System enclosure

When setting the system enclosure type, ensure that the enclosure serial number field matches the original value, which can be found on a label affixed to the unit. Updating the enclosure serial field keeps the configuration and error information synchronized. This information is used by the system when you create the location codes. This task must be done by using the ASMI, not with the control panel. However, if you do not have access to the ASMI, the system still operates without updating this information. See Figure 6-92.

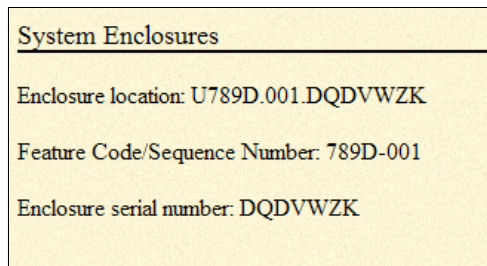


Figure 6-92 System Enclosures display example

Feature code and sequence number

Enter a feature code and sequence number in the form *FFFF-SSS*, where *FFFF* is the 4-character feature and *SSS* is the 3-character sequence number. The Feature Code/Sequence Number is used to uniquely identify the type of the enclosure that is attached to the system. A valid value is required for the machine to boot. When this value is changed, the service processor reboots so that the location codes can be updated accordingly.

Enclosure serial number

Enter an enclosure serial number in the form *XXYYYYYY*, where *XX* is the code for the plant of manufacture and *YYYYYY* is the unit sequence number. Valid characters are 0 - 9 and A - Z. This serial number must be different from the serial number on the machine. A valid value is required for the machine to boot. When this value is changed, the service processor reboots so that the location codes can be updated accordingly.

Service Indicators

From this menu, you can turn off the system attention indicator, enable enclosure indicators, change indicators by location code, and perform an LED test on the control panel.

The service indicators alert you that the system requires attention or service. It also provides a method for identifying a field-replaceable unit (FRU) or a specific enclosure within the system. A hierarchical relationship exists between FRU indicators and enclosure indicators. If any FRU indicator is in an identify state, then the corresponding enclosure indicator changes to an identify state automatically. You cannot turn off the enclosure indicator until all FRU indicators within that enclosure are in an off state.




Figure 6-93 Service Indicators menu

System attention indicator

Click this button to turn off the system attention indicator. If the indicator is off, you cannot use this option to turn the system attention indicator on again. See Figure 6-94.

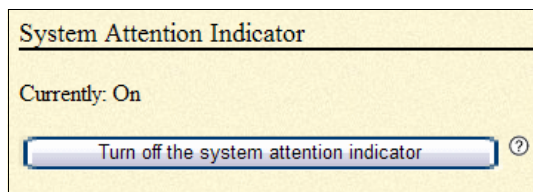


Figure 6-94 System Attention Indicator when indicator is **on**

Figure 6-95 shows the System Attention Indicator when the indicator is off.




Figure 6-95 System Attention Indicator when indicator is **off**

Enclosure indicators

You can turn on or off the identify indicators in each enclosure. An enclosure is a group of indicators. For example, a processing unit enclosure represents all of the indicators within the processing unit and an I/O enclosure represents all of the indicators within that I/O enclosure. Enclosures are listed by their location code. See Figure 6-96. Select the check box and select **Continue**.


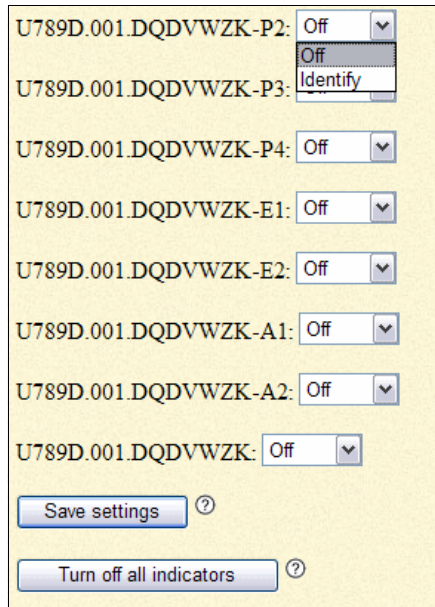


Figure 6-96 Enclosure Indicators window

Select to turn off or identify as appropriate and select **Save settings**. Alternatively, select **Turn off all indicators** to reset the LEDs. See Figure 6-97.



The screenshot shows a window titled "Enclosure Identify" with a yellow background. It contains a list of indicators, each with a dropdown menu. The indicators are:

- U789D.001.DQDVWZK-P2: Off
- U789D.001.DQDVWZK-P3: Identify
- U789D.001.DQDVWZK-P4: Off
- U789D.001.DQDVWZK-E1: Off
- U789D.001.DQDVWZK-E2: Off
- U789D.001.DQDVWZK-A1: Off
- U789D.001.DQDVWZK-A2: Off
- U789D.001.DQDVWZK: Off

At the bottom of the window, there are two buttons: "Save settings" and "Turn off all indicators". Both buttons have a question mark icon to their right.

Figure 6-97 Enclosure Identify window

Indicators by location code

You can specify the location code of any indicator to view or modify its current state. If you provide the wrong location code, the advanced system manager attempts to go to the next higher level of the location code. The next level is the base-level location code for that FRU. For example, a user types the location code for the FRU on the second I/O slot of the third enclosure in the system. If the location code for the second I/O slot is incorrect (the FRU does not exist at this location), an attempt to set the indicator for the third enclosure is initiated. This process continues until an FRU is located or no other level is available.

Lamp test

You can perform an LED test on the control panel to determine if one of the LEDs is not functioning properly. Select **Lamp Test**. Click **Continue** to do the lamp test. The test changes all indicators to the identify state for a short time (approximately 4 minutes).

6.3.7 Network Services menu

Use this menu option to configure the number and type of network interfaces according to the needs of your system (Figure 6-98). You can configure network interfaces on the system. The number and type of interfaces vary according to the specific needs of your system.



Figure 6-98 Network Services menu

Network configuration

This operation can be performed when the system is turned on and off. Because network configuration changes occur immediately, existing network sessions, such as HMC connections, are stopped. If a firmware update is in progress, do not do this operation. The new settings must be used to re-establish any network connections.

More errors can also be logged if the system is turned on. See Figure 6-99.

Network Configuration

Network interface eth0

Configure this interface? ?

MAC address: E4:1F:13:6F:8B:B6

IPv4: ?

Type of IP address: ?

Host name: ?

IP address: ?

Subnet mask: ?

Default gateway: ?

Network interface eth1

Configure this interface? ?

MAC address: E4:1F:13:6F:8B:B7

IPv4: ?

Type of IP address: ?

Host name: ?

IP address: ?

Subnet mask: ?

Default gateway: ?

Domain name: ?

IP address of first DNS server: ?

IP address of second DNS server: ?

IP address of third DNS server: ?

Figure 6-99 HMC Ethernet port configuration

From this window, you have the following options:

- ▶ **Configure this interface:** Configures this interface. If not selected, then the corresponding fields are ignored.
- ▶ **Type of IP address:** Select the IP address type for this interface. If *dynamic* is selected, then network configuration data is obtained from the DHCP server. Typically, your HMC is your DHCP server that is connected to FSP Ethernet port1.

- ▶ **Host name:** Enter a new value for the host name.
The following characters are valid: hyphen (-) and period (.), uppercase and lowercase alphabets (A to Z and a to z), and numeric (0 - 9).
The first character must be alphabetic or numeric and the last character must not be a hyphen or a period. However, if the host name contains a period, then the preceding characters must have an alphabetic character. This input is required for the static type of IP address.
- ▶ **Domain name:** Enter a new value for the domain name. All alphanumeric characters and the symbols hyphen (-), underscore (_), and period (.) are valid.
- ▶ **IP address:** Enter a new value for the IP address. This input is required for the static IP address type.
- ▶ **Subnet mask:** Enter a new value for the subnet mask. This input is required for the static IP address type.
- ▶ **Default gateway:** Enter a new value for the default gateway.
- ▶ **IP address of first DNS server:** Enter a new value for the first DNS server.
- ▶ **IP address of second DNS server:** Enter a new value for the second DNS server.
- ▶ **IP address of third DNS server:** Enter a new value for the third DNS server.
- ▶ **Reset Network Configuration:** Resets the Network Configuration settings to their default factory settings.
- ▶ **Network Configuration:** Select service processor to be configured. The default is the current service processor

Selecting **Save Settings** causes the network configuration changes to be made and the service processor to be rebooted. As the service processor reboots, your ASMI session drops and you have to reconnect your session to continue. When you reconnect, you are then using the new settings.

Network Access window

When you configure network access, you specify which IP addresses can access the service processor. You can specify a list of allowed IP addresses and a list of denied IP addresses. See Figure 6-100.

Allowed and denied lists: The allowed list takes priority over the denied list, and an empty denied list is ignored. *ALL* is not allowed in the denied list if the allowed list is empty

Network Access	
IP address: 192.168.128.1	
Allowed IP addresses ?	Denied IP addresses ?
1. <input type="text"/>	1. <input type="text"/>
2. <input type="text"/>	2. <input type="text"/>
3. <input type="text"/>	3. <input type="text"/>
4. <input type="text"/>	4. <input type="text"/>
5. <input type="text"/>	5. <input type="text"/>
6. <input type="text"/>	6. <input type="text"/>
7. <input type="text"/>	7. <input type="text"/>
8. <input type="text"/>	8. <input type="text"/>
9. <input type="text"/>	9. <input type="text"/>
10. <input type="text"/>	10. <input type="text"/>
11. <input type="text"/>	11. <input type="text"/>
12. <input type="text"/>	12. <input type="text"/>
13. <input type="text"/>	13. <input type="text"/>
14. <input type="text"/>	14. <input type="text"/>
15. <input type="text"/>	15. <input type="text"/>
16. <input type="text"/>	16. <input type="text"/>
<input type="button" value="Save settings"/>	

Figure 6-100 Network Access window

In this window, you have the following options:

- ▶ **Allowed IP addresses:** Enter up to 16 complete or partial IP addresses. A complete IP address contains all four octets.

A partial IP address has only 1, 2, or 3 octets, and must end in a period. If a login is received from an IP address, which matches a complete or partial IP address in the allowed list, access to the service processor is granted.

To allow access to the service processor from any IP address, enter ALL in the allowed list. An empty allowed list is ignored and access is granted from any IP address.

- ▶ **Denied IP addresses:** Enter up to 16 complete or partial IP addresses to be denied. Access to the service processor is not allowed if a login is received from an IP address that is listed in this list.

To deny access from any IP address, enter ALL in the list. If an incorrect IP address is entered in the allowed list and the denied list contains ALL, access to the service processor can be permanently denied. In this case, reset the network parameters by using the network reset parameters switch on the service processor card. An empty denied list is ignored and the allowed list takes priority over the denied list. For these reasons, ALL is not allowed in the denied list if the allowed list is empty.

6.3.8 Performance setup

You might enhance the managed system performance by manually or automatically changing the logical memory block size. The system kernel uses the memory block size to read and write files. By default, the logical memory block size is set to Automatic. This setting allows the system to set the logical block memory size that is based on the physical memory available. You can also manually change the logical memory block size. See Figure 6-101.




Figure 6-101 Performance setup

To select a reasonable logical block size for your system, consider both the performance that is wanted and the physical memory size. Use the following guidelines when you select logical block sizes:

- ▶ On systems with a small amount of memory installed (2 GB or less), a large logical memory block size results in the firmware consuming an excessive amount of memory. Firmware must consume at least one logical memory block. Generally, select the logical memory block size to be no greater than 1/8th the size of the physical memory of the system.
- ▶ On systems with a large amount of memory that is installed, small logical memory block sizes result in many logical memory blocks. Because each logical memory block must be managed during boot, many logical memory blocks can cause boot performance problems. Generally, limit the number of logical memory blocks to 8 K or less.

Attention: The logical memory block size can be changed at run time, but the change does not take effect until the system is restarted

Select **Logical Memory Block Size**. Select the logical memory block size and click **Save settings**.

System Memory Page setup

Improve your system performance by setting up the system with larger memory pages. Performance improvements vary depending on the applications running on your system. Only change this setting if advised by service and support.

To change the system memory page setup, select **System Memory Page Setup**. In the right pane, select the settings that you want, and then click **Save settings**.

6.3.9 On demand utilities

Activate inactive processors or inactive system memory without restarting your server or interrupting your business. CoD allows you to permanently activate inactive processors or inactive system memory without requiring you to restart your server or interrupt your business. You can also view information about your CoD resources. Important: Use this information if a hardware failure causes the system to lose its CoD or function on demand purchased capabilities, and if there never was an HMC managing the system. If an HMC is managing the system, use the HMC to do the following tasks instead of the ASMI.

Important: To decide whether you need CoD, refer to 4.5, “Capacity on Demand” on page 160.

CoD order information

After you determine that you want to permanently activate some or all of your inactive processors or memory, you must order one or more processor or memory activation features. You then enter the resulting processor or memory-activation key that is provided by your hardware provider to activate your inactive processors or memory.

To order processor or memory activation features select **On Demand Utilities** → **Select CoD Order Information**. The server firmware displays the information that is necessary to order a Capacity on Demand activation feature. Record the information that is displayed, and click **Continue**. See Figure 6-102.

CoD Order Information
System type: 8233
System serial number: 10-DD51P
Card type: 52B6
Card serial number: 00-817W000
Card ID: 1209070428616C62

Figure 6-102 CoD Order Information example

CoD activation

To activate this feature, click **Demand Utilities** → **CoD Activation**. Enter the activation key into the field and click **Continue** to do the specified operation. See Figure 6-103.

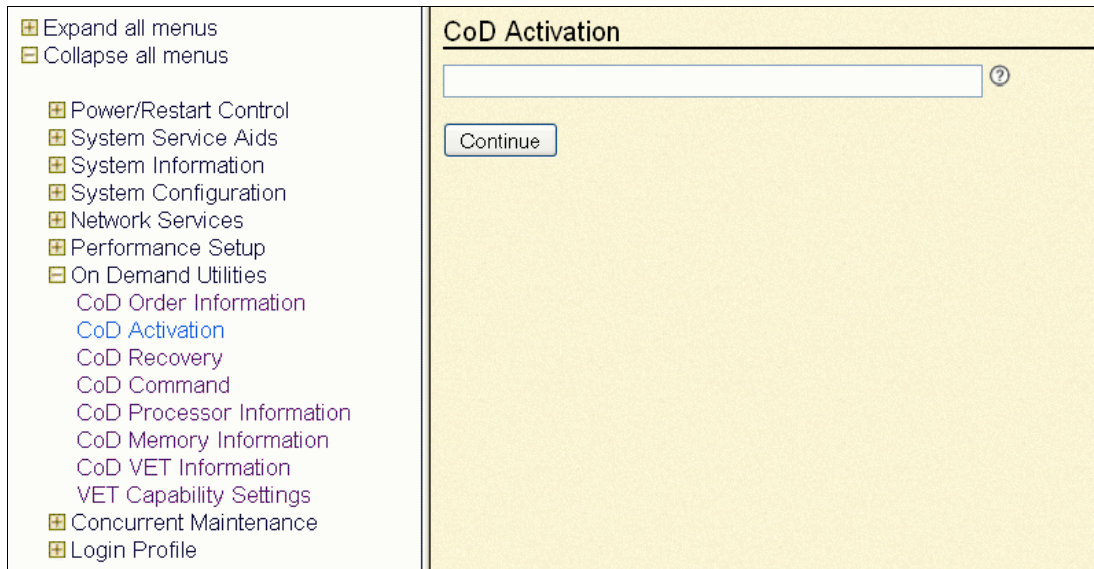


Figure 6-103 CoD Activation window

CoD recovery

This process is to resume the booting process of the server firmware after the CoD activation keys are entered. Resuming the server firmware causes the CoD key to become recognized and the hardware to become activated. This option allows the server to complete the startup process that is delayed up to one hour to place the server into the On Demand Recovery state that was needed to enter the CoD activation keys.

Select **On Demand Utilities** → **CoD Recovery**. Enter the activation key into the field and click **Continue** to do the specified operation (Figure 6-103).

CoD command

Select **On Demand Utilities** → **CoD Recovery**. Enter the command into the field and click **Continue**.

Viewing information about CoD resources

When CoD is activated on your system, you can view information about the CoD processors, the memory that is allocated as CoD memory, and Virtualization Engine technology resources.

Select **On demand Utilities**. Then, select one of the following options for the type of information that you want to view:

- ▶ **CoD Processor Information** to view information about the CoD processors.
- ▶ **CoD Memory Information** to view information about available CoD memory.
- ▶ **CoD Vet Information** to view information about available Virtualization Engine technologies.
- ▶ **CoD Capability Settings** to view information about the CoD capabilities that are enabled.

6.3.10 Login Profile

How to change passwords, view login audits, change the default language, and update the installed languages are now described.

Change password

You can change the general user, administrator, and HMC access passwords. If you are a general user, you can change only your own password. If you are an administrator, you can change your password and the passwords for general user accounts. If you are an authorized service provider, you can change your password, the passwords for general and administrator user accounts, and the HMC access password.

Passwords can be any combination of up to 64 alphanumeric characters. The default password for the general user ID is general, and the default password for the administrator ID is admin. After your initial login to the ASMI and after the reset toggle jumpers are moved, the general user and administrator passwords must be changed. The HMC access password is usually set from the HMC during initial login. If you change this password by using the ASMI, the change takes effect immediately.

Security measure: As a security measure, you are required to enter the current user's password into the current password for the current user field. This password is not the password for the user ID you want to change.

To change the password, select **Login Profile** → **Change password**. In the window that opens, enter the appropriate information and click **Continue**.

Retrieve login audits

You can view the login history for the ASMI to see the last 20 successful logins and the last 20 logins that failed. To view login audit, select **Login Profile** → **Retrieve Login Audits**.

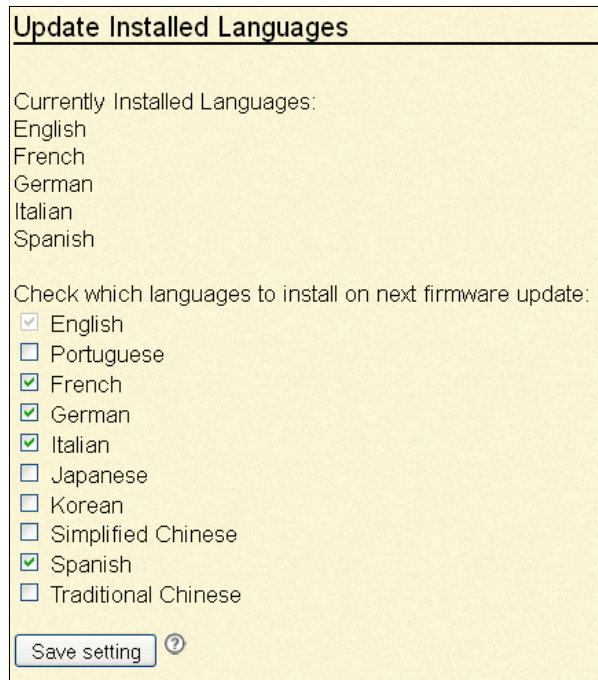
Change default language

You can select the language that is displayed on the ASMI welcome window before login and during your ASMI session if you do not choose an alternative language at the time of login. You must provide all requested input in English-language characters regardless of the language that is selected to view the interface.

To change the default language, select **Login Profile** → **Change Default Language**. Select the language and click **Save Settings**.

Update Installed Languages menu

A maximum of five languages can be supported on the service processor at any specified time. By default, English is always installed. Languages installation changes take effect when the firmware is updated. See Figure 6-104.



Update Installed Languages

Currently Installed Languages:
English
French
German
Italian
Spanish

Check which languages to install on next firmware update:

- English
- Portuguese
- French
- German
- Italian
- Japanese
- Korean
- Simplified Chinese
- Spanish
- Traditional Chinese

Save setting ?

Figure 6-104 Update Installed Languages menu

To select which language to install at the next firmware update, select **Login Profile** → **update installed language** → select five languages → **Save Settings**.

User Access Policy menu

This menu enables the admin user to grant or deny the access to service and development personnel by enabling or disabling dev, celogin, celogin1, and celogin2. Enabling access policy for celogin1 and celogin2 requires new passwords to be set even if they were set before.

To enable user access, select **Login Profile** → **user access policy**. Then, select the user ID and policy setting and click **Continue**. Enter the admin password and new password for the user.

You have the following password options:

- ▶ **Current password for user ID:** As a security measure, the current password must be supplied.
- ▶ **New password for user:** Enter the new password for the user whose password you want to change.
- ▶ **New password again:** Enter the new password for the user again for verification.



A

Introduction to IBM Systems Director

IBM Systems Director is an integrated, easy-to-use suite of tools that provide customers with flexible Systems Management capabilities to help realize maximum systems availability and lower IT costs.

This appendix provides an overview of the IBM Systems Director for managing IBM Power Systems servers. It helps you to understand the specifics of IBM Systems Director on the Power platform and to decide if IBM Systems Director is the best way to manage your environment.

Overview of IBM Systems Director

IBM Systems Director is an integrated suite of tools that provides you with a system management solution for heterogeneous environments, including the IBM Power Systems environment. IBM Systems Director works with the Hardware Management Console (HMC) to provide a comprehensive system management solution. With IBM Systems Director, IT administrators can view and track the hardware configuration of remote systems in detail and monitor the usage and performance of critical components, such as processors, disks, and memory.

IBM Systems Director is provided at no additional charge for use on IBM systems.

Extensions to IBM Systems Director are available for customers who want more capabilities from a consistent, single point of management. IBM Systems Director also complements and integrates with other popular Systems Management products by using its upward integration modules.

For more information about IBM Systems Director, see the following resources:

- ▶ IBM Systems Director:
<http://www.ibm.com/systems/software/director>
- ▶ IBM Systems Software Information Center:
<http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp?>
- ▶ *IBM Director on System p5*, REDP-4219
- ▶ *Implementing IBM Systems Director 6.1*, SG24-7694

There is a long list of features that are associated with IBM Systems Director. The following list provides the key features:

- ▶ Unifies the management of IBM systems, delivering a consistent appearance for common management tasks
- ▶ Integrates IBM best-of-breed virtualization capabilities to provide new and radically improved ways to simplify the management of physical and virtual platform resources
- ▶ Provides multi-system support for Power Systems, IBM System x®, IBM BladeCenter, IBM System z, and IBM Storage Systems
- ▶ Provides an extensible and modular foundation to advance the core Systems Management capabilities with more plug-ins
- ▶ Enables seamless integration of IBM systems with the total infrastructure

- ▶ Facilitates reduced training cost by using a consistent and unified platform management foundation and interface
- ▶ Manages non-IBM x86-based systems through a dedicated agent

IBM Systems Director is the next generation platform management solution of IBM Director that can improve the total cost of ownership by decreasing management costs and improving the usage of existing IT resources within a data center by eliminating the need to maintain multiple tools.

IBM Systems Director is designed to manage complex environments that contain many servers. Figure A-1 shows a sample environment that can be managed by using IBM Systems Director.

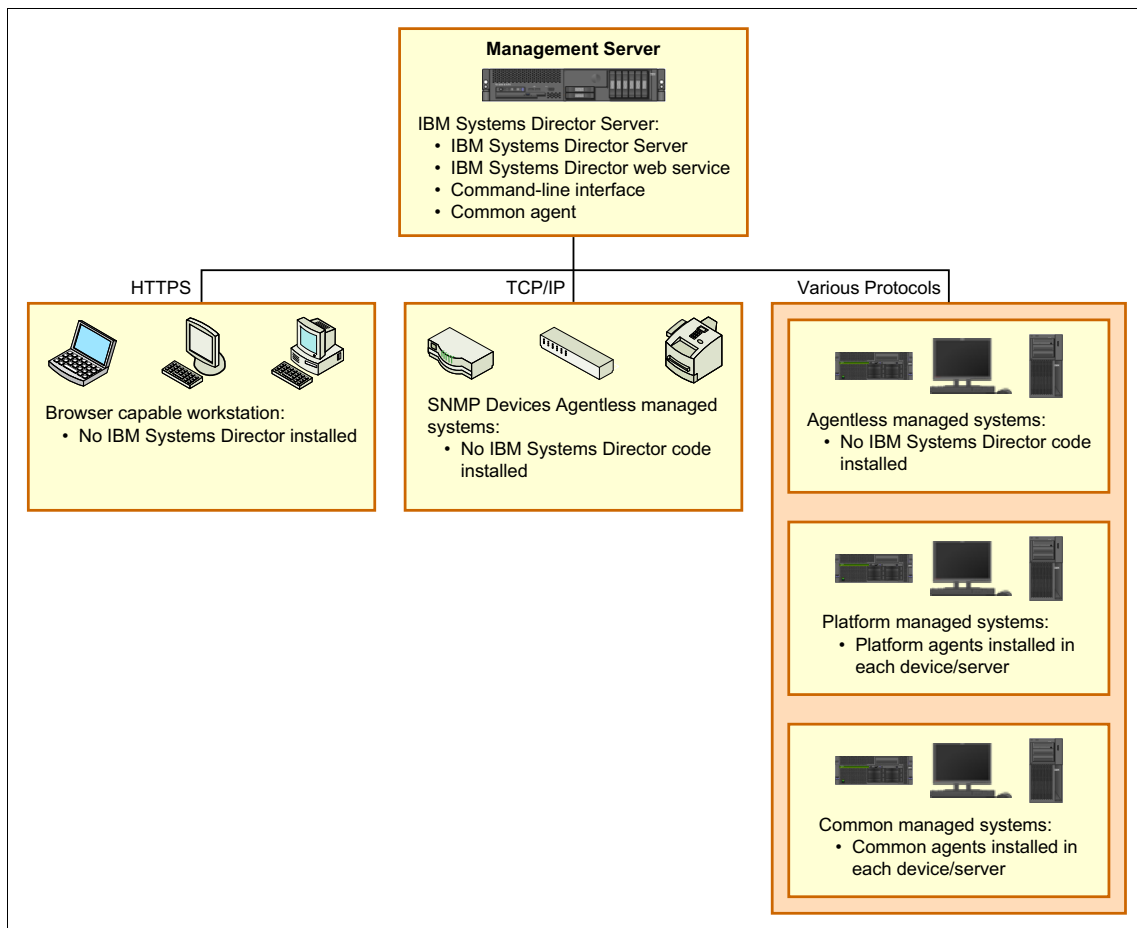


Figure A-1 IBM Systems Director environment

IBM Systems Director components

The hardware in a Systems Director environment can be divided into the following categories:

- ▶ *Management servers*: One or more servers on which IBM Systems Director Server is installed
- ▶ *Managed systems*: Servers, workstations, desktop computers, and notebook computers that are managed by Systems Director
- ▶ *SNMP devices*: Network devices, printers, or computers that have SNMP agents that are installed or embedded

IBM Systems Director software has three main components:

- ▶ Platform Agent
- ▶ Common Agent
- ▶ IBM System Director Server

Each managed endpoint in a Systems Director environment might have one or more of these components that are installed, each of which is described in the following sections.

Systems Director can manage some endpoints on which none of the previously mentioned components are installed. Such a managed system is now referred to as an agentless-managed system. This is equivalent to the Level-0 managed object terminology of a Director 5 environment.

These endpoints must at a minimum support either Secure Shell (SSH), distributed component object model (DCOM), or Simple Network Management Protocol (SNMP) in order for the Systems Director server to discover them. The function available to agentless-managed systems is limited to the following tasks, and varies based on operating system and hardware:

- ▶ Discover systems
- ▶ Collect limited operating-system inventory data
- ▶ Remotely deploy and install Platform Agent
- ▶ Remotely deploy and install Common Agent
- ▶ Perform limited remote access
- ▶ Perform limited restart capabilities

Platform Agent

Platform Agent is installed on managed systems where the smallest agent footprint is critical and management requirements are fairly simple. This agent communicates directly with both the operating system and the hardware (that is,

the service processor) to surface problems via Director Native Events, CIM indications, and SNMP traps to the management server. Platform Agent also is responsible for communicating with other Systems Management environments, which is referred to as upward integration. Platform Agent is equivalent to the Level-1 Agent or IBM Director Core Services component of a Director 5 environment.

Platform Agent provides a base set of management functionality that is used to communicate with and administer a managed endpoint. Systems that have Platform Agent (but not Common Agent) installed on them are referred to as Platform Agent managed systems.

The function available for Platform Agent managed systems is limited to the following tasks, and varies based on operating system and hardware:

- ▶ Discover systems
- ▶ Collect limited platform inventory data
- ▶ Monitor health and status
- ▶ Manage alerts
- ▶ Remotely deploy and install Common Agent
- ▶ Perform limited remote access
- ▶ Perform limited restart capabilities

Common Agent

Common Agent is the full-function management agent that is designed to provide comprehensive Systems Management capabilities. Once Common Agent is installed on an endpoint, more agent-side plug-ins can be installed to add advanced management functionality to the endpoint. For example, once Common Agent is installed on a VMware VirtualCenter Server, the IBM Systems Director Virtualization Manager plug-in can be pushed to that system to support advanced Virtualization Manager functionality that is particular to VirtualCenter. Common Agent is equivalent to the Level-2 Agent or IBM Director Agent component of a Director 5 environment.

Common Agent is installed on a managed endpoint to provide enhanced functionality for IBM Systems Director to communicate with and administer the system. Common Agent communicates with the management server through a single port (9510). This process is an improvement over the number of ports that are required for server-agent communication in Director 5, although more ports are required for certain types of functions. For example, remote command-line access to a Linux-managed system uses port 22, which is standard for the SSH protocol that is used for this operation.

Systems (IBM and non-IBM servers, desktop computers, workstations, and mobile computers, as well as virtual systems) that have Common Agent installed on them are referred to as *Common Agent managed systems*. The function available for Common Agent managed systems varies based on operating system and hardware, and includes the following tasks:

- ▶ Discover systems
- ▶ Collect comprehensive platform and operating system inventory data
- ▶ Monitor health and status
- ▶ Manage alerts
- ▶ Remotely deploy and install Common Agent
- ▶ Perform remote access, including transferring files
- ▶ Perform power management function
- ▶ Has more event support
- ▶ Monitor processes and resources
- ▶ Set critical thresholds that send notifications when triggered
- ▶ Manage operating system resources and processes

IBM Systems Director Server

IBM Systems Director Server is installed on the system that is to become the management server. Ideally, this system is a single system in the environment, but this is not always possible. In the case where multiple management servers are required, you must decide whether to install an *Agent Manager* on each Systems Director Server or to share a single Agent Manager between multiple management servers. The Agent Manager is new to IBM Systems Director 6.1 and is responsible for credentials and authentication between the IBM Systems Director Server and the Common Agent.

IBM Systems Director Server is the main component of IBM Systems Director and is completely rewritten for the Version 6.1 release. Systems Director Server contains the management data, the server engine, and the application logic. It provides basic functions such as discovery of the managed endpoints, persistent storage of inventory data, SQL database support, presence checking, security and authentication, web service, and administrative tasks. In the basic installation, Systems Director Server stores management information in an embedded Apache Derby database. You can access information that is stored in this integrated, centralized, relational database even when the managed endpoints are not available. For large-scale Systems Director solutions, you can use a stand-alone database application, such as IBM DB2® Universal Database™, Oracle, or Microsoft SQL Server.



Managing POWER processor-based blades

With the introduction of HMC V7R760, the HMC can now manage IBM POWER6 and POWER7 processor-based blades.

This appendix describes steps that are needed to perform to Power Systems blades through BladeCenter *Advanced Management Module (AMM)* in order to be managed by an HMC.

Managing a POWER processor-based blade with the HMC

Managing a POWER processor-based blade with the HMC appliance is accomplished the same way as managing any other system. The POWER7 and POWER6 blades are discovered and managed directly through the flexible service processor (FSP) IP address.

Follow these configuration steps:

1. The POWER processor-based blade firmware must be
 - for POWER7 - **01AA730_094** or later.
 - for POWER6 - **01EA350_132** or later

Firmware can be obtained from the IBM Fix Central website:

<http://www.ibm.com/support/fixcentral/>

Instruction for updating blade firmware can be found at:

- for POWER7 -

https://publib.boulder.ibm.com/infocenter/bladectr/documentation/index.jsp?topic=/com.ibm.bladecenter.ps700.doc/ps700_t_update_firmware.html

- for POWER6 -

https://publib.boulder.ibm.com/infocenter/bladectr/documentation/index.jsp?topic=/com.ibm.bladecenter.js23.doc/dw1im_t_update_firmware.html

2. The IBM BladeCenter AMM firmware must be at the latest level.

- a. To download the latest AMM firmware, see this website:

<http://www.ibm.com/support/fixcentral/systemx/selectFixes?product=ibm/systemx/8852&&platform=NONE&function=all>

Then, select **Management Module**, and then select and download the firmware.

- b. Extract the **ibm_fw_amm_<nnnnn>_anyos_noarch.zip** file, and follow installation instructions in the **readme.txt** file.

3. There is no requirement for Virtual I/O Server level. However, the latest general available (GA) level is recommended.

Virtual I/O Server upgrades and fixes can be obtained from this website:

<http://www14.software.ibm.com/webapp/set2/sas/f/vios/home.html>

4. Connect the HMC *open network* to the BladeCenter Ethernet switch in Bay 1, as shown in Figure B-1.




Figure B-1 HMC networks diagram

- By default, the FSP is not accessible in the network. AMM must be used to enable the FSP in the network. From the BladeCenter management web graphical user interface (GUI), go to the **Blade Tasks** → **Configuration** → **Management Network** tab, as shown in Figure B-2.

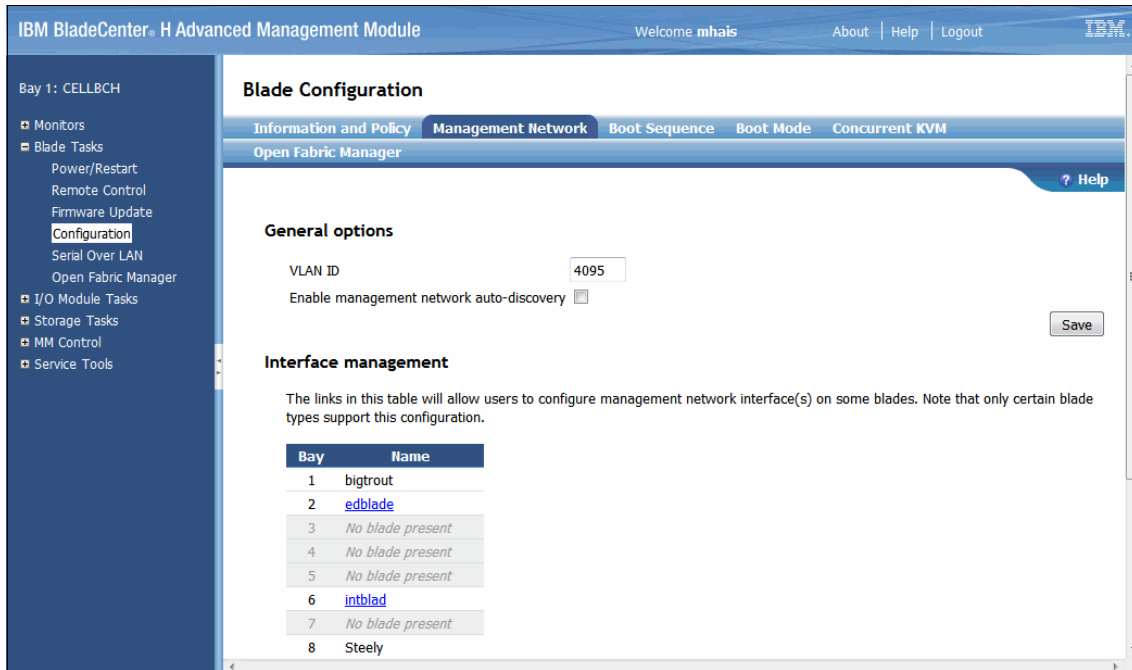


Figure B-2 Blade configuration management network

- Click the required Power Systems blade **name** in the **Interface management** list. In our example, *edblade* in Bay# 2, as shown in Figure B-2.

7. In the **General Settings** section for the **eth0** interface, change the **Enable/Disable NIC** attribute to *Enabled* and press **Save**.

In the IPv4 section for eth0 interface, change the **DHCP** attribute to *Disabled*. Use *Static IP configuration* and enter the open network **IP Address**, **Subnet mask**, and **Gateway address** for the eth0 (FSP) interface. Press **Save**.

See Figure B-3.

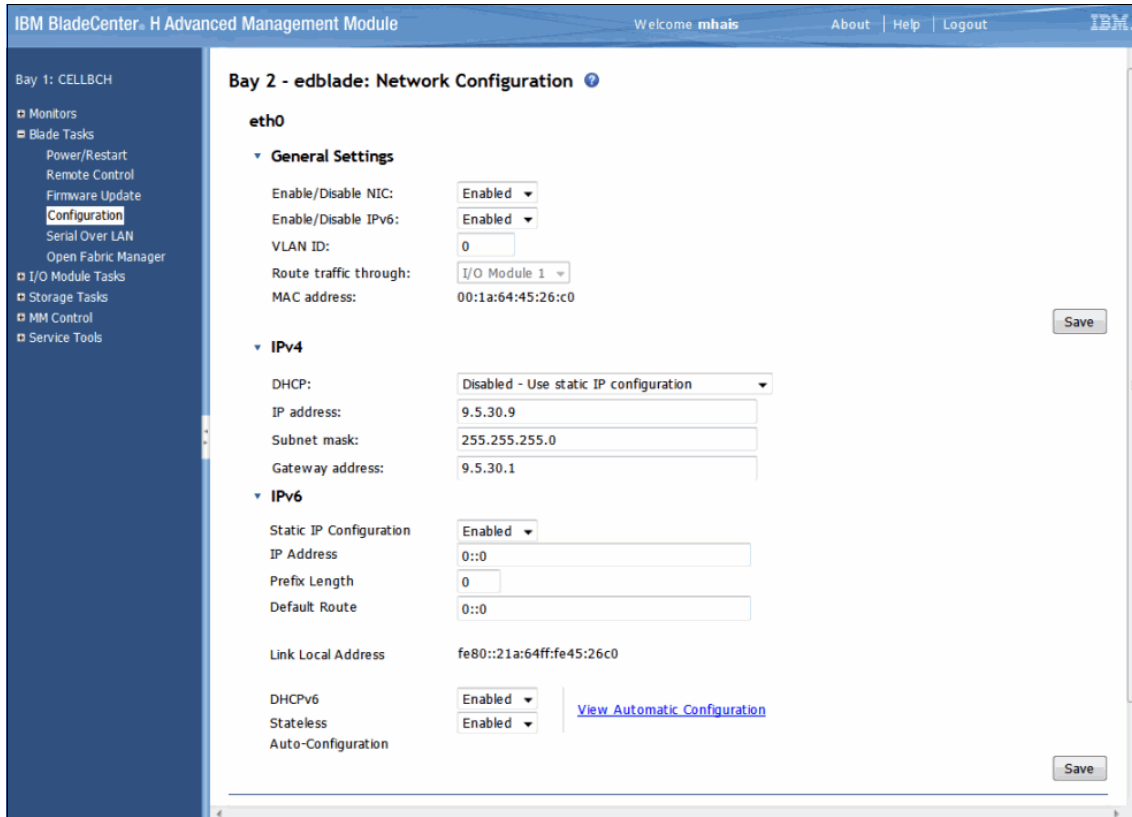


Figure B-3 Setting FSP IP address for Power Systems blade

Implementing IPv6? If you are implementing IPv6 in your environment, complete *step 7* with *IPv6 section*.

8. If HMC vterm is used for partition ID 1, then Serial Over LAN (SOL) must be disabled in AMM.
 - a. Go to **Blade Tasks** → **Configuration** → **Serial Over Lan**

- b. Select the require blade (*ebblade* in our example). In **Available actions** select *Disable Serial Over Lan* and press **Perform action**. The **green SOL status** changes to **gray**.
9. Using the web browser login to ASMI with the IP address provided in step 7., on page 397, accept the certificates.

In the first login, you are asked to provide the password for ASMI. After that login to ASMI, you get a window like what is shown in Figure B-4.

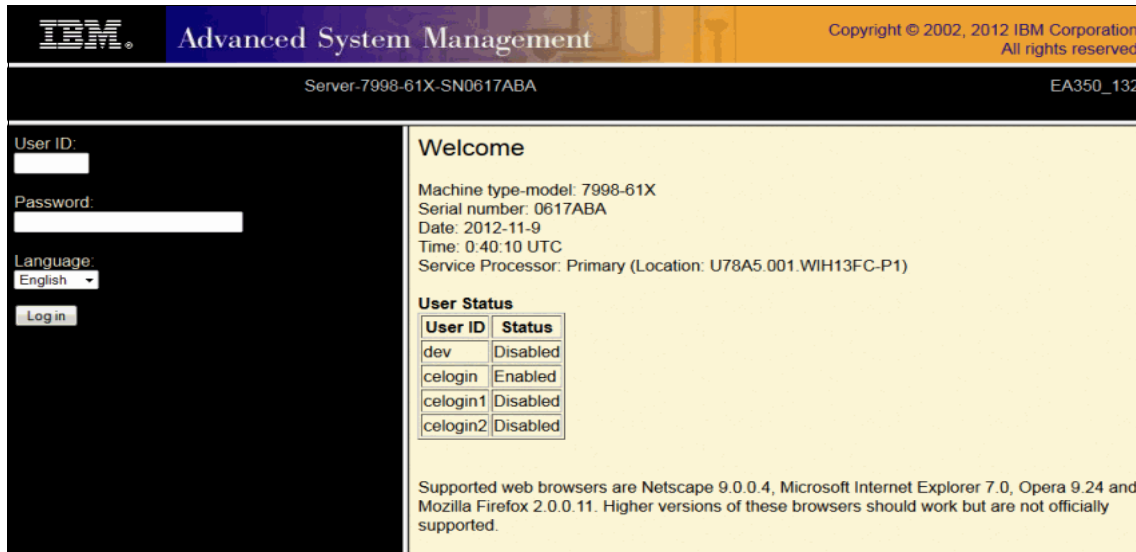


Figure B-4 Power Systems blade ASMI login

10. Log in to HMC and go to **Systems Management** → **Server**. In the **Task: Servers** pane, select **Connections** → **Add Management System**. In **Add Managed System**, enter the IP address that is provided in step 7 on page 397, and password that is provided in step 9., on page 398.

- Continue as though it was a new Power Systems server added to the HMC.
Figure B-5 shows a Power Systems blade that is managed by the HMC.

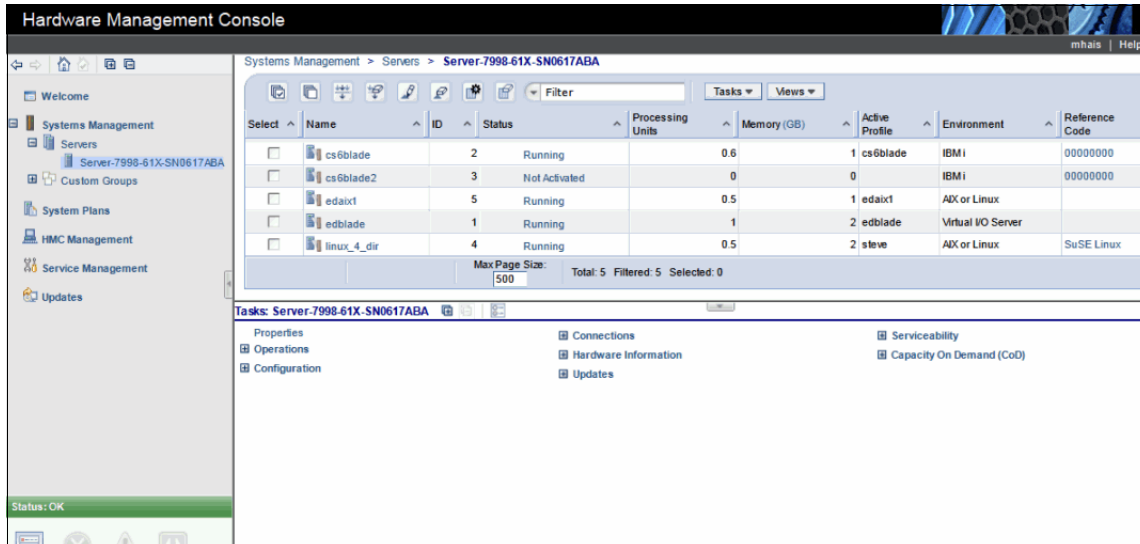



Figure B-5 Hardware Management Console



IBM product engineering debug data collection

The appendix describes the process of Hardware Management Console (HMC) IBM product engineering (PE) debug data collection. This data collection might be required to support a problem determination in an IBM Power Systems environment.

Preparing to collect the pedbg

In preparation to collect the data, we have to ensure two things: there are special users (*hscpe*) to collect the data, and the transfer method to offload debug data.

The user ID *hscpe* with task role *hmcpe* is required to exist in the HMC to run the collection.

To verify the existence of the *hscpe* user by opening the HMC terminal, there are two ways to open the terminal. The first way is to log in to the HMC management graphical user interface (GUI) by using an HMC restricted terminal in HMC management GUI. The second way is to use enabled remote command execution and do a remote access from your Secure Shell (SSH) enabled workstation.

1. Log in to the terminal by using an *hscroot* user or user with *hmcsuperadmin* roles.
2. Do a command

```
lshmcusr --filter "names=hscpe"
```

Press **Enter**. If the *hscpe* does not exist, it shows. See Figure C-1.

```
hscroot@hmc7:~> lshmcusr --filter "names=hscpe"  
No results were found.  
hscroot@hmc7:~> █
```

Figure C-1 List HMC user command shows no *hscpe* user

If the result has a *hscpe* user, skip to running the collection.

```
hscroot@hmc7:~> lshmcusr --filter "names=hscpe"  
name=hscpe,taskrole=hmcpe,description=IBM Service,pwage=99999,resource=ALL,  
authentication_type=local,remote_webui_access=0,remote_ssh_access=1,min_pwage=0,  
session_timeout=0,verify_timeout=15,idle_timeout=0,inactivity_expiration=0,resources=<ResourceID = ALL:><UserDefinedName = AllSystemResources>,password_encryption=md5,disabled=0  
hscroot@hmc7:~> █
```

Figure C-2 List HMC user command show *hscpe* user

3. To create a *hscpe* user, do the following command:

```
mkhmcusr -u hscpe -a hmcpe -d IBM Service
```
4. Prepare a removable media that you want to use to offload the logs. See “Format Media option” on page 281.

Run pedbg collection command

Sign on to the HMC with an hscpe user profile.

5. Run the PE debug collection in quiet mode by doing the following command:

```
pedbg -c -q 3
```

or

```
pedbg -c -q 3 9
```

Prompt to copy files to media.

Tips: The valid options for collecting HMC data in quiet mode are as follows:

1 = Network information only

2 = Network information + base logs

3 = Network information + base logs + extended logs

4 = Network information + base logs + extended logs + archives

5 = Collect only those files in the /home/hscpe/ibmsupt directory

9 = Run prompt to copy files to media

Or, run *pedbg* normally, which creates a file on the HMC and automatically prompts the user to move the file to a removable media or to leave it on the HMC. Do the following functions:

- Run command:

```
pedbg -c
```

Press **Enter** to proceed.

- Respond **Yes** to all questions up to the prompt to collect archives. When prompted to collect archives (“Would you like to collect archived log data?”), respond **No** unless the archives are specifically requested by the support representative.

Offload pedbg collection

When prompted with “Would you like to move zip file to a DVD or other device?”, select one of the following questions. Select **Yes** to move the data to DVD-RAM, USB media device, or to a remote secure copy server (scp Server) with **scp** command.

Verify the DVD-RAM or USB media device is correctly formatted for *service data* and *inserted* before replying **Yes**. It might take the HMC several seconds to recognize the USB devices. Then, enter the name of the device name when prompted. The device name is the mount point for the target device, for example DVD would be `/media/cdrom` and USB devices would be `/media/sdb1`.

For the network option, you are prompted for user name, IP address for the SCP enabled server, and the working directory. You might also be prompted to accept the RSA key fingerprint of the SCP server. Choose **Yes** to accept the key, then wait for the offload to start.

Select **No** if any other method is used to copy the data from the HMC. The HMC keeps the collection on the hard disk drive.

Figure C-3 shows the process of offloading logs to a USB media device.

```
These additional logs would be collected on request of PE support.
No

Created /dump/HMClogs.hmc71108H44.zip

Would you like to move zip file to a DVD or other device?
The answer no will write to the HMC hardfile
Please type yes or no.
yes

The file is 114683212 bytes. Ensure you have enough room on the device.
Use formatted Read/Write media only
Enter a device name from the list below. Example /media/cdrom

/media/cdrom=CD/DVD
/media/sdb1=USB flash memory device
network=move to another system via the scp command

/media/sdb1

Copy to /media/sdb1 in progress
```

Figure C-3 Sample of the process of offloading logs to a USB media device

Figure C-4 shows the process of offloading logs to a network scp server.

```
/media/cdrom=CD/DVD
/media/sdb1=USB flash memory device
network=move to another system via the scp command

network

    Enter the user name on the destination host.
ftpuser
    Enter the destination host. example myhost.mycompany.com
172.16.254.42
    Enter the destination filesystem on the host. example /tmp/hmcdir
    The directory must exist and have correct permissions.

/Users/ftpuser
    scp /dump/HMClogs.hmc71108H45.zip ftpuser@172.16.254.42:/Users/ftpuser
The authenticity of host '172.16.254.42 (172.16.254.42)' can't be established.
RSA key fingerprint is 86:a4:f4:c9:47:b5:33:9f:39:37:29:eb:99:1c:81:6f.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '172.16.254.42' (RSA) to the list of known hosts.
Password:
HMClogs.hmc71108H45.zip          100% 111MB 705.7KB/s   02:41
The send was successful
hscpe@hmc7:~> █
```

Figure C-4 Sample of process offloading logs to a network scp server



D

Live Partition Mobility support log collection

This appendix describes information to gather for problems that involve IBM i Live Partition Mobility (LPM) errors. This process ensures that the log files do not wrap or get overwritten. Having this log beforehand when calling IBM Support helps clients get a faster response in problem determination.

Background information

You must provide background information of the current LPM configuration or any changes that were made before the problem occurs.

- ▶ Did this work in the past?
If so, what changed if known? In other words, was the HMC, server firmware, Virtual I/O Server (VIOS), Virtual I/O Server adapter microcode, IBM i, IBM AIX, or SAN updated since the last failure?
- ▶ The name of the server and partition that failed.
- ▶ The approximate date and time (both the HMC and partition times) of the failure.
- ▶ Virtual I/O Server LPAR level, partition name, and ID, both source and target.
- ▶ HMC level and fixes that run the LPM.
- ▶ How the LPM was initiated, graphical user interface (GUI) or command-line interface (CLI)?

HMC log collection

For configuration with more than one HMC, you must repeat the steps on both the source HMC and target HMC:

- ▶ pedbg HMC log. Details in Appendix C, “IBM product engineering debug data collection” on page 401.
- ▶ Rename the logs with problem management record (PMR) number `PMRno.ZZZ.000.HMClogs*.zip` if the LPM moves across the HMC. See Example 6-1.

Example 6-1 HMC pedbg logs for LPM for more than one HMC

```
PMRno.ZZZ.000.HMClogs*.source.zip  
PMRno.ZZZ.000.HMClogs*.target.zip
```

Operating system log collection

There are certain level requirements for LPM activity:

▶ IBM i

Verify that at minimum IBM i 7.1 technology refresh is installed. Issue the **WRKPTFGRP** command.

Do a **DSPPTF OUTPUT(*PRINT)** and **WRKPTFGRP OUTPUT(*PRINT)**. Get the spooled file to your local workstation

▶ AIX

- a. AIX LPAR level, partition name, and ID.
- b. Collect the **snap** for AIX LPAR and make a copy to your workstation. See Example 6-2.

Example 6-2 Provide AIX client SNAP report

```
# snap -r
# snap -ac
# move /tmp/ibmsupt/snap.pax.Z to
/tmp/ibmsupt/PMRno.ZZZ.000.client.src_msp.target_msp.snap.pax.Z
```

Take a snap dump: You should have one client LPAR snap. For more clients failing, you need an operating system log for every client LPAR.

Virtual I/O Server Mover Server Partition snaps

Collect snaps data from each Virtual I/O Server Mover Server Partition (MSP). This process includes one or two (if redundant, Virtual I/O Server is configured) source partitions and one or two (if redundant, Virtual I/O Server is configured) target partitions.

1. snap

To collect the snap data on the Virtual I/O Server partition, take the following steps:

- a. Log on to the Virtual I/O Server.
- b. Run the **snap** command and press **Enter**. Wait until the command completes the process.
- c. The snap file is in `/home/padmin` and is named `snap.pax.Z`.

- d. Rename the file to include the PMR number and indicate if it is the storage or target. See Example 6-3.

Example 6-3 If you have redundant VIO (four VIO)

```
mv /home/padminsnap.pax.Z PMRno.ZZZ.000.vio1.source.snap.pax.Z
```

If you have redundant HMC, your files will be:
PMRno.ZZZ.000.vio1.source.snap.pax.Z
PMRno.ZZZ.000.vio2.source.snap.pax.Z
PMRno.ZZZ.000.vio1.target.snap.pax.Z
PMRno.ZZZ.000.vio2.target.snap.pax.Z

2. ctsnap

This log is needed if there is an RMC-related issue.

- a. Enter the commands that are shown in Example 6-4.

Example 6-4 Collecting ctsnap

```
$ oem_setup_env  
# ctsnap -x runrpttr
```

This command produces a report log file in /tmp/ctsupt/ctsnap*.tar.gz.

- b. Rename the file to include the PMR number and indicate if it is the storage or target. See Example 6-5.

Example 6-5 Rename the file to include PMR number

```
# mv /tmp/ctsupt/ctsnapx.tar.gz  
/tmp/PRMNo.ZZZ.000.source.ctsnap.pax.Z
```

If you have redundant HMC, your files will be:
PMRno.ZZZ.000.vio1.source.ctsnap.pax.Z
PMRno.ZZZ.000.vio2.source.ctsnap.pax.Z
PMRno.ZZZ.000.vio1.target.ctsnap.pax.Z
PMRno.ZZZ.000.vio2.target.ctsnap.pax.Z

Hypervisor HMC resource dump

The hypervisor HMC resource dump is to be collected only if requested by an IBM service representative. If you have dual HMCs to manage the system, you need only one resource dump. Do not run it again from the redundant HMC. Instructions for collecting hypervisor HMC dump is on “Manage Dumps option”

on page 282. Or, you can do it by using the HMC restricted terminal. See Example 6-6.

Example 6-6 Collect hypervisor resource dump through a terminal

To find your managed system do:

```
# lssyscfg -r sys -F name
```

Then initiate the system dump

```
# startdump -t resource -m {managed system} -r "system"
```

File will be created in /dump directory.

```
# ls -l /dump
```

will show the file. Do not offload the dump unless there is a file named .IN_PROGRESS extension removed by the HMC, this file indicate that the system dumps still in progress offloading from managed system FSP to HMC.

After the process in Example 6-6 is completed, you can offload the dump by using the manage dump GUI in HMC.

Sending logs to IBM

There are several ways to send the logs to IBM. There is also an option to send the log individually by using HMC outbound connectivity. However, we describe the transfer by creating a single file:

- ▶ FTP to IBM
- ▶ IBM Enhanced Customer Data Repository (ECuRep) website

A single file is preferred. If the file is too large, a couple of files can be transferred. Move all above pax.Z compressed and log files in to a single directory, then archive this directory. See Example D-1.

Example D-1 Example to archive log files in AIX

```
$ mkdir -p /tmp/pmrnumber/pmdata
```

```
move or ftp or scp data to sample directory above
```

```
$ cd /tmp/pmrnumber
```

```
$ pax -xpax -vw pmdata | gzip -c > data_collected.pax.gz
```

FTP to IBM

After a single file is created, rename the file to include the PMR number as follows:

For example, if your PMR is 12345,999,000

- ▶ 12345 is the PMR number
- ▶ 999 is the branch number
- ▶ 000 is the country code

Do the following command:

```
$ mv data_collected.pax.gz 12345.999.000_data_collected.pax.gz
```

Then, transfer the files to IBM through FTP:

```
$ ftp testcase.software.ibm.com
```

```
login: anonymous
```

```
passwd: (your email address in format your_email_id@your_email_domain)
```

```
ftp> cd /toibm/aix (for IBM AIX) or ftp> cd /toibm/os400 (for IBM i)
```

```
ftp> bin
```

```
ftp> put <filename>
```

```
ftp> quit
```

IBM ECuRep

IBM Enhanced Customer Repository (IBM ECuRep) is another way to transfer logs to IBM. IBM ECuRep is a secure and fully supported data repository with problem determination tools and functions. It updates PMRs and maintains full data lifecycle management. IBM ECuRep needs an IBM ID and qualified PMR/Incident Number to send the data.

For more information, see the IBM ECuRep website:

<http://www.ibm.com/de/support/ecurep/index.html>

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- ▶ *A Practical Guide for Resource Monitoring and Control (RMC)*, SG24-6615
- ▶ *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940
- ▶ *Implementing IBM Systems Director 6.1*, SG24-7694
- ▶ *Hardware Management Console V7 Handbook*, SG24-7491
- ▶ *IBM PowerVM Live Partition Mobility*, SG24-7460
- ▶ *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590
- ▶ *Converting Hardware Management Console (HMC) 7042-CR6 or 7042-CR7 Models to RAID1*, REDP-4909

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other publications

This publication is also relevant as a further information source:

- ▶ *Operations Guide for the Hardware Management Console and Managed Systems*, SA76-0085

Online resources

These websites are also relevant as further information sources:

- ▶ IBM Power Systems Hardware Information Center
<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp>
- ▶ HMC POWER Code Matrix
<http://www-304.ibm.com/webapp/set2/sas/f/power5cm/home.html>
- ▶ IBM United States Hardware Announcement 112-175, dated October 3, 2012
<https://www.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=897/ENUS112-175&infotype=AN&subtype=CA>
- ▶ VIOS sizing rule of thumb link to the Nigel Griffiths AIXpert Blog
https://www.ibm.com/developerworks/mydeveloperworks/blogs/aixpert/entry/rule_of_thumb_sizing_the_virtual_i_o_server78?lang_en
- ▶ VIOS Performance Advisor
<https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/Power%20Systems/page/VIOS%20Advisor>
- ▶ IBM AIX information center for virtual device management
<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7hcg/mkvdev.htm>
- ▶ Utility Capacity on Demand
<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7ha2/utilitycapacityondemandkick.htm>
- ▶ On/Off Capacity on Demand
 - <http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7ha2/onoffcodbillchange.htm>
 - <http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7ha2/onoffcodchangerequest.htm>
- ▶ Power Systems Capacity on Demand
<http://www.ibm.com/systems/power/hardware/cod/index.html>
- ▶ IBM System Capacity on Demand Information
<http://www-912.ibm.com/pod/pod/>
- ▶ Accessing the ASMI without HMC
http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7hby/connect_asmi.htm

- ▶ Capacity on Demand Activation
<http://www-03.ibm.com/systems/power/hardware/cod/activations.html>
- ▶ Capacity on Demand Recovery
<http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/i pha2/tcodend.htm>
- ▶ HMC PDF information resources
<http://www14.software.ibm.com/webapp/set2/sas/f/hmcl/resources.html>
- ▶ IBM Electronic Support Agent website
<http://www.ibm.com/support/electronic>
- ▶ IBM ID registration
<http://www.ibm.com/registration>
- ▶ Fix Level Recommendation Tool website
<https://www-304.ibm.com/support/customer/care/flrt/>
- ▶ IBM Fix Central website
<http://www-933.ibm.com/support/fixcentral/options>
- ▶ Compatibility HMC Software and managed system website
<http://www.ibm.com/support/fixcentral/firmware/supportedCombinations>
- ▶ IBM Systems Director
<http://www-03.ibm.com/systems/software/director>
- ▶ IBM Systems Software Information Center
<http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp?>
- ▶ POWER6 BladeCenter update instruction
https://publib.boulder.ibm.com/infocenter/bladectr/documentation/index.jsp?topic=/com.ibm.bladecenter.js23.doc/dw1im_t_update_firmware.html
- ▶ IBM AMM BladeCenter AMM firmware download
<http://www.ibm.com/support/fixcentral/systemx/selectFixes?product=ibm/systemx/8852&platform=NONE&function=all>
- ▶ VIOS upgrades and fixes
<http://www14.software.ibm.com/webapp/set2/sas/f/vios/home.html>
- ▶ IBM Enhanced Customer Data Repository
<http://www-05.ibm.com/de/support/ecurep/index.html>

- ▶ Must Gather LPM IBM i
http://www-912.ibm.com/s_dir/slkbases.NSF/DocNumber/633439208
- ▶ Must Gather LPM AIX
ftp://ftp.software.ibm.com/systems/virtualization/vio/ztools/lpm-data-collection/lpm_diagnostic_data_requirements.doc
- ▶ PE Debug Collection knowledge database
http://www-912.ibm.com/s_dir/slkbases.NSF/DocNumber/451766819

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Abbreviations and acronyms

AMM	Advanced Management Module	DHCP	Dynamic Host Configuration Protocol
AMS	Advanced Management System	DIMM	dual inline memory module
ARP	Address Resolution Protocol	DLPAR	Dynamic Logical Partition
ASCII	American Standard Code for Information Exchange	DMA	Direct Memory Access
ASM	Advanced System Management	DNS	Domain Name Server
ASMI	Advanced System Management Interface	DST	Dedicated Service Tools
BIOS	Basic Input Output Setup	DVD	Digital Versatile Disk
BPA	Bulk Power Assembly	DVD-RAM	Digital Versatile Disk-Random Access Memory
BPC	Bulk Power Controller	ESA	Electronic Service Agent
BPH	Bulk Power Hub	FLRT	Fix Level Recommendation Tool
BSR	Barrier Synchronization Register	FP	Fix Pack
CBU	Capacity Back Up	FRU	field-replaceable unit
CCD	Critical Console Data	FSP	flexible service processor
CD	compact disc	FTP	File Transport Protocol
CD-ROM	compact-disc read-only memory	GA	general availability
CDT	Central Daylight Time	GB	gigabyte
CE	Customer Engineer	GMT	Greenwich mean time
CEC	central electronics complex	GPFS	General Parallel File System
CIM	Common Information Model	GUI	graphical user interface
CLI	command-line interface	GUID	Globally Unique Identifier
CoD	Capacity on Demand	HACMP	High Availability Cluster Multi Processing
CPU	central processing unit	HBA	Host Bay Adapter
CSM	Cluster Systems Management	HCA	Host Channel Adapter
DCOM	Distributed Component Object Model	HDD	hard disk drive
		HEA	Host Ethernet Adapter
		HMC	Hardware Management Console
		HPS	High Performance Switch
		HSL	high-speed link

HTTPS	Hypertext Transfer Protocol Secure	MES	Miscellaneous Equipment Specification
IBM	International Business Machines Corporation	MPT	Modem Parameter Table
ICMP	Internet Control Message Protocol	MSP	Managed Service Provider
ID	identification	MTM	Machine Type Model
IDE	Integrated Drive Electronics	NDP	Neighbor Discovery Protocol
IEEE	Institute Electrical Electronics Engineers	NFS	Network File System
IM	Information Management	NIC	network interface card
I/O	input/output	NIM	Network Installation Management
IOA	input/output adapter	NPIV	N-Port ID Virtualization
IOP	input/output processor	NTP	Network Time Protocol
IP	Internet Protocol	NVRAM	Non-Volatile RAM
IPL	Initial Program Load	OEM	Original Equipment Manufacturer
IS	IBM support	OS	operating system
IT	Information Technology	OSI	Open Systems Interconnection
ITSO	International Technical Support Organization	PC	personal computer
IVE	Integrated Virtual Ethernet	PCI	Peripheral Component Interconnect
KDC	Key Distribution Center	PDF	Portable Document Format
KVM	Keyboard, Video, Mouse	PE	Product Engineer
LAN	Local Area Network	PHYP	Power Hypervisor
LDAP	Lightweight Directory Access Protocol	PM	Performance Management
LED	light-emitting diode	PMR	Problem Management Record
LHEA	Logical Host Ethernet Adapter	PPP	Point to Point Protocol
LIC	License Internal Code	PSP	Preventive Service Planning
LP	Logical Partition	RAID	Redundant Array of Independent Disks
LPAR	Logical Partition	RAS	reliability, availability, and serviceability
LPM	Live Partition Mobility	RDMA	Remote Direct Memory Access
LV	logical volume	RIO	remote input/output
LVT	LPAR Validation Tool	RMC	Resource Monitoring and Control
MAC	Media Access Control	RPO	Record Purpose Only
MB	megabyte		

RPQ	request for price quotation	URL	Uniform Resource Locator
RSA	Remote Supervisor Adapter	US	United States
SAN	storage area network	USB	Universal Serial Bus
SAS	serial-attached SCSI	UTC	Universal Time Clock
SCP	Secure Copy Protocol	VIO	Virtual I/O
SCSI	Small Computer System Interface	VIOS	Virtual I/O Server
SDMC	Systems Director Management Console	VLAN	virtual local area network
SF	Support Facility	VM	virtual machine
SFP	Service Focal Point	VPD	vital product data
SFTP	Secure FTP	VPN	virtual private network
SMS	System Management Server	WWN	worldwide name
SMTP	Simple Mail Transport Protocol		
SNAP	Sub Network Access Protocol		
SNI	Switch Network Interface		
SNMP	Simple Network Management Protocol		
SOL	Serial Over LAN		
SP	Service Pack		
SPCN	System Power Control Network		
SPT	System Planning Tool		
SQL	Structured Query Language		
SRC	System Reference Code		
SSH	Secure Shell		
SSL	Secure Socket Layer		
SSP	System Support Program		
SSR	System Service Representative		
SSS	Software Service Support		
TCP	Transmission Control Protocol		
TCP/IP	TCP/Internet Protocol		
UDP	User Datagram Protocol		
UEFI	Unified Extensive Firmware Interface		

Index

Symbols

/opt/hsc/data/sysplan 14

Numerics

7310-CR2 60

7310-CR3 61

7310-CR4 60

A

Activate 205

Additional HMC Users 66

Administrator mailing address 81

ambiguity 29

ASMI

 default IP addresses for service processors 343

Automatic allocation 70

B

Back up 234

Bulk Power Assembly (BPA) 220

 status 222

C

cage number 220

Central Electronic Complex 29

Change Network Settings 235

Change Password 195

Change User Interface Settings 237

Change User Password 241

CLI

 common usage examples 225–226

 lshwres

 examples 227–229

 lssyscfg 227–229

 lssysconn 225–226

CoD 203

 benefits 160

 enhancements 10

Configure Connectivity to Your Service Provider 82

Configure Customizable Data Replication 253

Configure DNS 77

Configure Domain Suffix 78

Configure HMC Firewall 74

Configure HMC Gateway 76

Configure HMC Network Settings 71

connection monitoring 154

Connections

 disconnect Another HMC 200

 reset or remove 200

Create Logical Partition 195

Create System Plan 21

Customizable Data Replication 248

D

Data Replication 248

Date and Time 65

 change 239

dedicated processor idle cycles 9

Delete 211

Deploy 30, 32, 37

Deploy System Plan 32

Deployment 31–32, 34, 37–38

Deployment examples 31

DHCP server 70, 89

DIAG_DEFAULT 156

DIAG_STORED 156

directory path for a sysplan 14

Discovered Console Information 253

Display the processing unit identifier 357

DLPAR

 Memory 216

 Physical Adapters 216–217

 Processor 216

DNS server IP address 77

Domain 76

Domain suffix 78

Dump Retry 208

E

Electronic Service Agent 87

Enabling hardware inventory collection from active partitions 22

Export System Plan 20

F

- Fast power off 189
- Firewall 73
- Firewall Settings
 - RMC 24
- Firmware update policy 357
- Format Media 234
- Frames
 - change password 195, 221
 - initialize 220
 - rebuild information 221
 - reset connections 222

G

- GUI
 - differences 6
- Guided Setup wizard 240
 - launching 65

H

- hardware inventory
 - enabling detailed retrieval 22
- hardware inventory collection
 - enabling 22
- Hardware Management Console 13
- hardware validation during system plan deployment 29
- HMC 29
 - firewall settings 73
 - hardware validation - important note 29
 - viewing hardware details with a system plan 14
- HMC Management 231
- HMC Version 231
- Host Ethernet Adapter 12, 224
 - configuration 201
- Host Name 76
- hscroot 65

I

- I/O Units 220
- IBM Director 387–388
 - capabilities 390
 - components 390
- Identify LED 191
- Import System Plan 17
- Initialize
 - resetting all configuration data 262, 264

- initialize a frame 220–221
- initiating deployment using the HMC graphical interface 31
- Internet Secure Sockets Layer (SSL) 82
- Internet Virtual Private Network (VPN) 82
- inventory scout 22
- invscout 22
- IOA 14
- IOP 14
- IP ranges 70

L

- LAN adapter 67
- Language 246
- language 64
- LHEA 12
- Locale 246
- locale 64
- logical HEA 12

M

- Manage Certificates 243
- Manage Custom Groups 213, 222
- Manage Data Replication 252
- Manage Dumps 282
- Manage Profiles 212
- Manage Remote Connections 279
- Manage Tasks 241
- Manage User Profiles 241
- Manage User Profiles and Access 114
- Manage Users 242
- Managed Resource Roles 116
- Managed Systems
 - connecting to HMC 89
- Master-to-subordinate replication 250, 252
- mksysplan 25

N

- network settings 66
- node status 112
- Nonroutable IP address ranges 70
- Normal power off 189

O

- Open 5250 Console 233
- Open Restricted Shell Terminal 233
- OPEN_FIRMWARE 156

P

- Partition auto start 91, 187
- Partition Availability Priority 10
- partition availability priority 195
- Partition communication 23
- Partition Data
 - Backup 265
 - manage 197
- Partition standby 91, 187
- partition validation 30
- partition validation during deployment 30
- Pass-Through System 85
- password 65
 - changing 115
- PCI error injection policies 357
- Peer-to-Peer Replication 250
- Pending Authentication 90
- Post Guided Setup tasks 88
- POWER Hypervisor 129
- Private Network 68
- private network 70
- processor enclosures (CEC) 29
- profile data
 - Backup priority 263
 - Delete 266
 - Full restore 263
 - Managed system priority 263
 - restoring 262

R

- Rebuild 194
- Redbooks website 413
 - Contact us xxv
- Redundant HMC configuration 43
- Redundant HMC considerations 45
- Remote Command Execution 243
- Remote Operation 246
- remote support requests 278
- Remote Virtual Terminal 245
- Remove System Plan 28
- Removing system plan on the HM 28
- Resource Management and Control (RMC)
 - firewall settings 24
- Resource Monitoring and Control 22
- Resource roles 241
- ResourceLink 274
- Restart 233
 - Dump 208

- HMC 233
 - Operating System 208
 - Operating System Immediate 208
- Restore HMC Data 235, 314
- RMC 22
- root password 66

S

- Save Current Configuration 214
- Save Upgrade Data 235
- Schedule Operations 210, 234
- Selected deployment wizard action 35
- Server processor 89
- Service Focal Point (SFP) 154
- service processor 89
- Service Processor Status 199
- Service Provider 82
- Session Preservation 176
- Shared Ethernet Adapter 123
- Shut Down 208, 233
 - Delayed 210
 - Immediate 210
 - Operating System 210
 - Operating System Immediate 210
- SMS 156
- SMTP 79
- ssh 245
- step 38
- steps 36–37
- sysplan 14
- System Management Services 156
- system plan 14, 28
 - creating 21
 - deployment 28
 - exporting 18
 - functions 16
 - hardware validation 28
 - importing 17
 - partition validation 28
 - removing 28
 - validation 33
 - viewing 25
- System Plan Viewer 35
- System Planning Tool 30
- System Planning Tool (SPT) 14
- System Plans 195
 - HMC menu 16
- System profile 91, 187

manage 197

T

Test LED 191

Test Network Connectivity 235

Tip of the Day 236

two LAN adapters 67

U

Utilization Data 194

V

Validation 28

validation of hardware 29

View HMC Events 231

View Network Topology 235

View RIO Topology 222

View System Attention LED 190

View System Plan 22, 25

Virtual I/O Server 123

W

Welcome Text 247

workload management groups 196



IBM Power Systems HMC Implementation and Usage Guide

(0.5" spine)
0.475" <-> 0.875"
250 <-> 459 pages



IBM Power Systems HMC Implementation and Usage Guide



**Practical guide to
using the IBM HMC in
virtualized Power
Systems servers**

**Documents the RAID
1 feature on IBM
HMC CR6 and CR7**

**Updated to include
HMC V7R760 and
IBM POWER7**

The IBM Hardware Management Console (HMC) provides systems administrators a tool for planning, deploying, and managing IBM Power Systems servers. This IBM Redbooks publication is designed for system administrators to use as a desk-side reference when managing partition-capable IBM Power Systems servers by using the HMC.

The major functions that the HMC provides are Power Systems server hardware management and virtualization (partition) management.

The following topics are described:

- ▶ Plan to implement the HMC
- ▶ Configure the HMC
- ▶ Operate the HMC
- ▶ Manage software levels on the HMC
- ▶ Use service functions on the HMC
- ▶ Update firmware of managed systems
- ▶ Use System Planning Tool deployments

In addition, there is an explanation on how to use the new HMC graphical user interface and the new HMC commands that are available with HMC Version 7, Release 7, modification 60.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks