

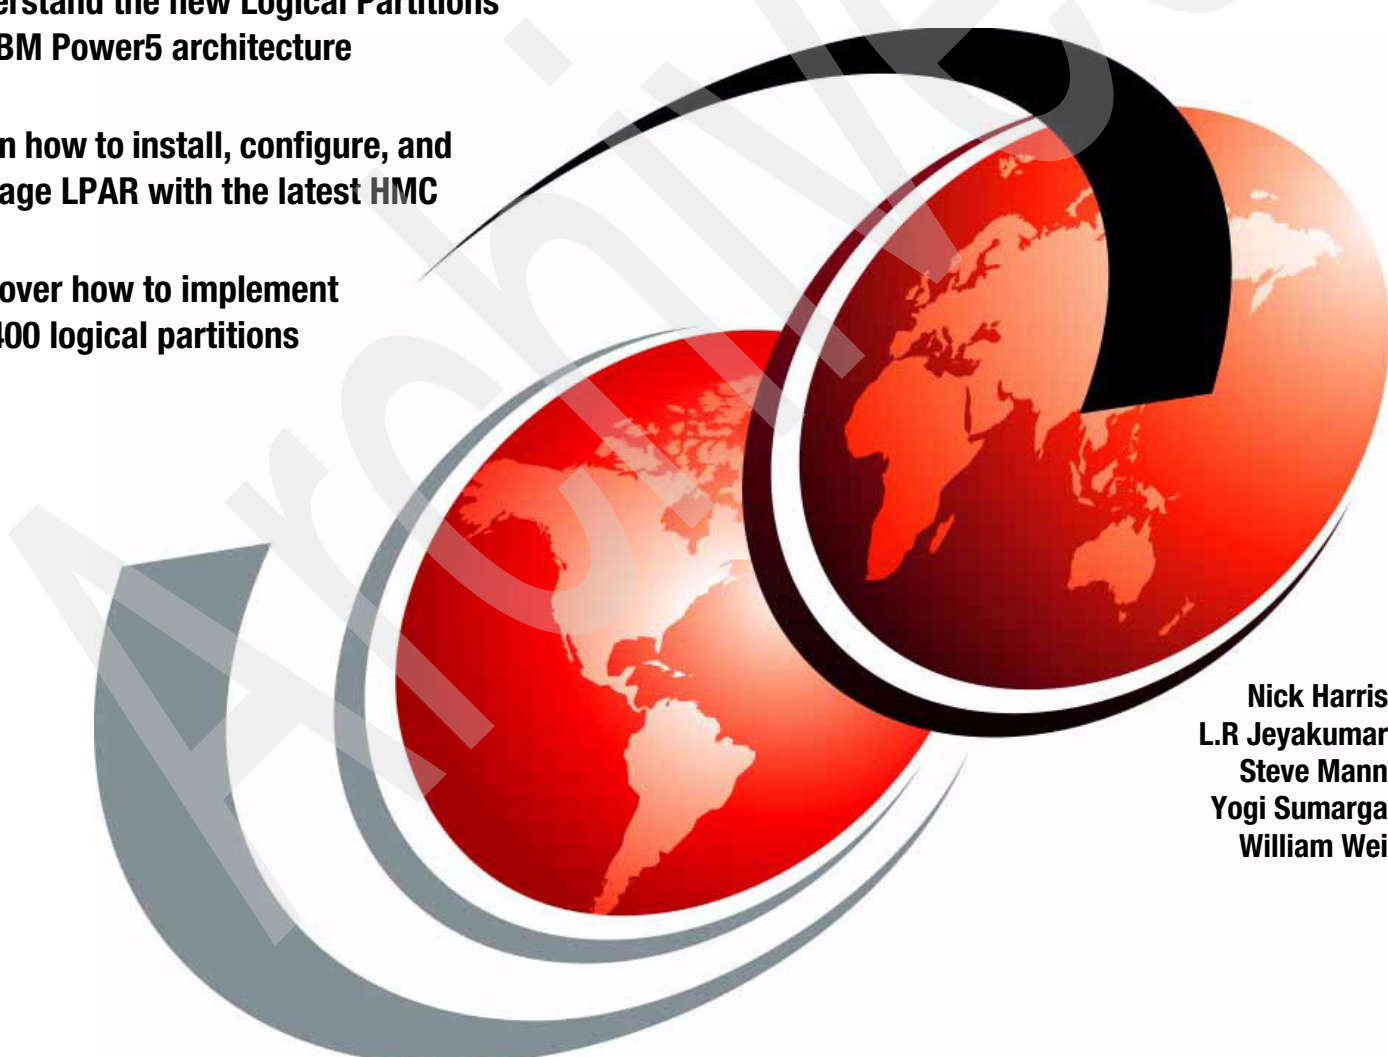
Logical Partitions on System i5

A Guide to Planning and Configuring LPAR with HMC on System i

Understand the new Logical Partitions for IBM Power5 architecture

Learn how to install, configure, and manage LPAR with the latest HMC

Discover how to implement OS/400 logical partitions



**Nick Harris
L.R. Jeyakumar
Steve Mann
Yogi Sumarga
William Wei**

Redbooks



International Technical Support Organization

**Logical Partitions on System i5
A Guide to Planning and Configuring LPAR
with HMC on System i**

May 2006

Archived

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

Archived

Second Edition (May 2006)

This edition applies to i5/OS Version 5, Release 3, and the System i5™ system products.

© Copyright International Business Machines Corporation 2005, 2006. 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

Notices	ix
Trademarks	x
Preface	xi
The team that wrote this redbook	xi
Become a published author	xiii
Comments welcome	xiii
Summary of changes	xv
May 2006, Second Edition	xv
Chapter 1. Introduction to LPAR on IBM System i5	1
1.1 HMC, Hypervisor, and partitions	2
1.2 Software requirements	5
1.2.1 What's new for licensing	5
1.2.2 Licensing requirements	5
1.2.3 Simple scenario with dedicated and shared capped partitions	6
1.2.4 Normal scenario with shared capped and uncapped partitions	6
1.2.5 Complex scenario with shared uncapped partitions	7
1.2.6 A more complex scenario with shared uncapped partitions and capacity upgrade on demand	7
1.3 Processor use in System i5 LPARs	8
1.3.1 System i5 processor overview	8
1.3.2 Logical partition resources	11
1.3.3 Dedicated processor	12
1.3.4 Shared processor	13
1.3.5 Shared uncapped processor	14
1.3.6 Shared capped processor	15
1.3.7 Virtual processor	15
1.3.8 Configuring dedicated processors for the logical partition	20
1.3.9 Configuring shared capped processors for logical partition	21
1.3.10 Configuring shared uncapped processors for logical partition	23
1.3.11 Considerations on using shared processors	24
1.3.12 Considerations in configuring virtual processors	25
1.4 Memory use in System i5 LPARs	25
1.4.1 Memory granularity	25
1.4.2 Memory and LPAR	26
1.4.3 Hardware page table	26
1.4.4 Memory assignment at initial LPAR start	26
1.4.5 Memory allocation for the i5 OS logical partition	27
Chapter 2. i5/OS consoles under System i5	31
2.1 Console history	32
2.2 Console for stand-alone systems	32
2.3 Console for partitioned (LPAR) systems	34
2.4 Console IOA placement rules for stand-alone systems	35
2.4.1 Model 520/550 CEC	35
2.4.2 Model 570 CEC	35
2.5 Console IOA placement rules for stand-alone systems, including IXS considerations	36

2.5.1	Model 520/550 CEC	36
2.5.2	Model 570 CEC	36
2.6	Connecting to a 5250 console remotely	37
2.6.1	5250 console remotely configured	37
Chapter 3.	HMC overview and planning	45
3.1	HMC concepts and initial setup	46
3.1.1	HMC and Service Processor levels, and information sources	46
3.1.2	Types of HMC	47
3.1.3	Initial setup of the HMC	48
3.2	Installing the HMC	50
3.3	HMC networking options	53
3.3.1	Private direct networking	53
3.3.2	Private indirect networking	54
3.3.3	Private and open networking	55
3.3.4	Open networking	55
3.3.5	Working with HMC in open networks	55
3.4	Initial tour of the desktop	56
3.4.1	Terminal menu	56
3.4.2	Net menu	57
3.4.3	Lock menu	58
3.4.4	Shutting down or logging off the HMC	58
3.5	Server and partition	60
3.5.1	Managed system properties	61
3.5.2	Other system-wide options	71
3.5.3	Available partition options	74
3.5.4	HMC configuration	97
3.6	Service applications	99
3.6.1	Problem determination	99
3.6.2	Inventory Scout services	100
3.6.3	Service Focal Point	100
3.6.4	Remote support	104
Chapter 4.	HMC Guided Setup	105
4.1	HMC Guided Setup	106
4.1.1	HMC Guided Setup wizard checklist	106
4.1.2	Using the Guided Setup wizard	107
Chapter 5.	Partition creation using the HMC	139
5.1	System and partition profiles	140
5.1.1	System profiles	140
5.1.2	Partition profiles	140
5.1.3	Uses of partition profiles	140
5.2	Creating an i5/OS logical partition through the HMC	141
5.2.1	Partition creation overview	141
5.2.2	Starting the create partition wizard	142
5.2.3	Specifying the partition name, partition ID, and partition type	144
5.2.4	Workload management group	146
5.2.5	Partition profile name	147
5.2.6	Partition memory	148
5.2.7	Partition processors	149
5.2.8	Interactive (5250 OLTP) capacity	155
5.2.9	Allocate physical I/O resources	155
5.2.10	Virtual IO adapters	157

5.2.11	Tagging partition resources	161
5.2.12	Opticonnect	169
5.2.13	Specifying power control partitions	170
5.2.14	Miscellaneous profile settings	170
5.2.15	Review profile summary	171
5.2.16	New partition profile has been created	172
5.3	Creating another profile for an existing partition	173
5.4	Changing the default profile for a partition	175
Chapter 6.	Working with the HMC	177
6.1	Accessing LPAR functions	178
6.2	Viewing partition properties	179
6.2.1	Partition properties	179
6.2.2	Partition profile properties	184
6.3	Starting and stopping partitions	191
6.3.1	Changing IPL side and mode	191
6.3.2	Starting a powered off partition	192
6.3.3	Restarting a running partition	193
6.3.4	Stopping (powering down) a running partition	194
6.4	Performing dynamic LPAR (DLPAR) functions	195
6.4.1	Physical adapters	195
6.4.2	Processors	196
6.4.3	Memory	197
6.4.4	Virtual IO adapters	197
6.5	Performing service functions on partitions	198
6.5.1	Displaying reference code information	199
6.5.2	Posting DST to the console	199
6.5.3	Performing main storage dumps	200
6.5.4	Working with remote service	200
6.5.5	Deleting a partition	201
6.5.6	Working with IOP functions	203
6.5.7	Concurrent Maintenance Power Domain	204
6.6	Remote management	205
6.6.1	Web SM	208
6.7	iSeries control panel functions versus HMC options	210
6.8	Licensed Internal Code Updates	211
6.8.1	HMC Code Update	212
6.9	Troubleshooting	213
6.9.1	Management Server State	213
6.9.2	Correcting the managed system operating state	214
6.9.3	List of progress codes	216
6.9.4	Using progress codes	217
6.9.5	Service processor progress codes	218
6.9.6	Service processor reset	225
6.10	HMC boot up fails with "fsck"	232
6.11	Determining the HMC serial number	233
Chapter 7.	HMC security and user management	235
7.1	System Manager Security	236
7.2	Overview and status	237
7.2.1	Certificate Authority (CA)	237
7.2.2	Server Security	237
7.2.3	Object Manager Security	238

7.3 HMC User Management	238
7.3.1 HMC Users	238
7.3.2 Add, modify, copy, or remove user profile.	239
7.3.3 Customizing user task roles and managed resource roles.	242
Chapter 8. HMC duplication and redundancy.	245
8.1 Redundant HMC configurations	245
8.2 Multiple managed system configuration	247
8.3 Cloning HMC configurations	248
8.4 Redundant HMC configuration considerations	249
Chapter 9. Migration of existing LPAR profiles to HMC.	251
9.1 Migration planning	252
9.1.1 Use DST or SST to print a copy of your current system configuration	252
9.1.2 Get a printout of your resource allocations	253
9.1.3 Review the resources in the Primary partition.	253
9.1.4 Bus numbering	253
9.1.5 Use the LVT tool to create a graphical worksheet	253
9.1.6 Review the list of supported and unsupported hardware	254
9.1.7 Create a customized planning checklist	254
9.1.8 Schedule a Solutions Assurance Review	254
9.1.9 Order the System i5	255
9.2 Preparing your system for migration to System i5 and HMC	255
9.2.1 Load V5R3	255
9.2.2 Update the Linux kernel	255
9.2.3 iSeries Navigator.	255
9.3 Migrating Linux partitions	256
9.4 Migration scenarios	256
9.4.1 Information gathering and pre-sales planning.	257
9.4.2 Post-sales customer tasks for both scenarios.	272
9.4.3 Post sales tasks: IBM	276
9.4.4 Customer migration tasks	277
9.5 Backing up HMC profiles.	279
9.5.1 Backing up Critical Console Data	280
9.5.2 Scheduling and reviewing scheduled HMC backups	282
Chapter 10. Using the Advanced Systems Management Interface	283
10.1 ASMI introduction	284
10.2 ASMI concepts and initial setup	284
10.2.1 Types of ASMI	284
10.2.2 Initial setup of terminal interface to ASMI	284
10.3 Initial tour of the interface	285
10.3.1 Power/restart control.	287
10.3.2 System service aids	292
10.3.3 System information	299
10.3.4 Network services.	323
10.3.5 Performance setup	327
10.3.6 On demand utilities	329
10.3.7 Concurrent maintenance.	332
10.3.8 Login Profile	336
Chapter 11. OpenSSH.	341
11.1 Utilities available in Open SSH	342
11.2 Installing the OpenSSH licence program in i5/OS	342

11.3	Using the HMC from i5/OS with OpenSSH	343
11.4	Running DLPAR scripts from i5/OS	350
11.5	Scheduling the DLPAR function from i5/OS	355
11.5.1	Scheduling the i/o movement	356
11.6	Scheduling the DLPAR function from Windows	368
Chapter 12. Using Work Management to influence System i5 resources		371
12.1	Main storage and processor affinity concept.	372
12.1.1	Processor affinity concept.	372
12.1.2	Memory affinity concept	372
12.2	Work Management enhancements for memory and processor affinity	373
12.2.1	QTHDRSCAFN (thread affinity)	374
12.2.2	QTHDRSCADJ (automatically adjust thread resources).	374
12.2.3	ADDRTGE command — new parameters	375
Chapter 13. Virtual Partition Manager		377
13.1	Introduction to Virtual Partition Manager for eServer System i5 systems	378
13.2	Planning for Virtual Partition Manager	379
13.2.1	Planning considerations	379
13.2.2	Planning for Linux logical partitions.	380
13.2.3	Perform capacity planning	380
13.2.4	Design and validate your partition configuration	381
13.2.5	Understand use of virtual resources for Linux partitions.	381
13.2.6	Ordering your eServer i5 system	381
13.3	Getting started with Virtual Partition Manager.	381
13.3.1	Minimum configuration requirements	382
13.3.2	Complete initial setup of your eServer i5	383
13.4	Preparing your system for Virtual Partition Manager.	384
13.4.1	Removing resources from i5/OS.	384
13.4.2	Understanding virtual I/O slots and requirements for IPLs	396
13.5	Creating Linux partitions using Virtual Partition Manager	397
13.5.1	Defining new Linux partitions	397
13.6	Establishing network connectivity for Linux partitions	404
13.6.1	Defining i5/OS components that support Proxy ARP	405
13.7	Setting up i5/OS virtual I/O resources for Linux partitions	410
13.7.1	Network Server Descriptor	410
13.7.2	Network Server Storage Space.	412
13.7.3	Starting the Linux installation	417
13.8	Virtual media management	419
13.8.1	Linux native backup with virtual tape	419
Chapter 14. Firmware maintenance.		421
14.1	HMC firmware maintenance	422
14.1.1	How to determine the HMC installed software level	423
14.1.2	Obtaining HMC update software.	424
14.1.3	Changing the i5/OS default FTP attributes	426
14.1.4	HMC backup of critical console data.	427
14.1.5	Updating the HMC software	430
14.1.6	Installing an individual HMC fix	438
14.2	Licensed internal code updates	440
14.2.1	Firmware overview	440
14.2.2	Setting the i5 system firmware update policy	442
14.2.3	Displaying the current firmware levels	446
14.2.4	Updating firmware through the HMC (out-of-band).	448

14.2.5 Updating firmware through an i5/OS service partition (in-band)	457
Chapter 15. HMC Access Password Reset Using Advanced System Management Interface (ASMI)	463
15.1 Accessing the ASMI using the HMC	464
15.2 Accessing the ASMI using a Web browser	464
Appendix A. HMC command list	473
HMC CLI introduction	474
HMC CLI commands listed by task	474
HMC CLI commands by name	476
HMC CLI commands by category	477
Working with the managed system	477
Working with logical partitions	478
Dynamic LPAR (DLPAR) operations	480
Working with partition profiles	483
Working with system profiles	485
Listing hardware resources	486
HMC CLI command usage	487
Change hardware resources (chhwres)	487
Change system configuration (chsyscfg)	488
Change system state (chsysstate)	489
List hardware resources (lshwres)	490
List system configuration (lssyscfg)	491
Create (make) system configuration (mksyscfg)	492
Remove system configuration (rmsyscfg)	493
List dumps (lsdump)	493
HMC CLI command attributes	494
Glossary	507
Index	509

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 give 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites 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.

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 illustrates 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

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

AIX®
AIX 5L™
AS/400®
Electronic Service Agent™
@server®
eServer®
eServer™
IBM®
iSeries™

i5/OS®
Lotus®
OpenPower™
OS/400®
Power PC®
PowerPC®
POWER™
POWER4™
pSeries®

Redbooks™
Redbooks (logo) ™
System i5™
System p5™
Tivoli®
Virtualization Engine™
Wake on LAN®
WebSphere®
xSeries®

The following terms are trademarks of other companies:

Java, Javadoc, JDBC, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

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

Preface

This IBM Redbook gives a broad understanding of the new System i5™ architecture as it applies to logically partitioned System i5 systems. This functionality is delivered through a new configuration and management interface called the Hardware Management Console (HMC).

Reading this redbook will help you design your server partition scheme from scratch. We also discuss the requirements to create a solution to migrate from existing iSeries™ servers with and without logical partitions.

This redbook will help you install, tailor, and configure logical partitions (LPARs) on System i5 systems. You will understand how to set up the server and HMC via the Guided Setup Wizard. We also cover the creation of multiple i5/OS® partitions.

The team that wrote this redbook

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

Nick Harris is a Consulting IT Specialist for the iSeries and has spent the last seven years in the International Technical Support Organization, Rochester Center. He specializes in LPAR, iSeries hardware and software, external disk, Integrated xSeries® Server for iSeries, and Linux. He writes and teaches IBM classes worldwide on areas of i5, iSeries, and AS/400® system design and server consolidation. He spent 13 years in the United Kingdom (UK) AS/400 Business, and has experience in S/36, S/38, AS/400, and iSeries servers.

You can contact him by sending e-mail to: <mailto:niharris@us.ibm.com>

L.R Jeyakumar has been with IBM since May 2000, working as a IT Specialist in IBM Global Services India (pvt) Ltd in Bangalore. He had six years of experience in the IT technical support field before he started with IBM. He works on IBM System i5, iSeries, and AS/400, as well as IBM System p5™, pSeries®, and Openpower systems.

At present he is also the Regional Specialist for iSeries platforms. His job responsibility includes planning, implementation, and support for all the iSeries platforms. He has been actively involved in RISC to RISC migrations, OS upgrades, LPAR implementations, Linux installations, and upgrades on iSeries, IBM System i5, and OpenPower™ systems.

You can contact him by sending an e-mail to: <mailto:njeyakum@in.ibm.com>

Steve Mann is an iSeries Advisory I/T Specialist working for IBM Switzerland. He has 15 years experience in working with the AS/400 and iSeries servers. He is both an IBM Certified Solutions Expert for iSeries and a Certified Tivoli® Storage Manager consultant. His areas of expertise include logical partitioning planning and implementation, systems management, BRMS, IBM Tivoli Storage Manager, and system recovery.

He has extensive experience with LPAR, IBM Tivoli Storage Manager for OS/400® PASE and BRMS, providing planning and implementation services for customers. He coauthored the following IBM Redbooks™: *LPAR Configuration and Management: Working with IBM iSeries Logical Partitions*, SG24-6251; and *Integrating Backup Recovery and Media Services and IBM Tivoli Storage Manager on the IBM iSeries Server*, SG24-7031-00.

Yogi Sumarga is an Account Product Services Professional working for IBM Global Services in Indonesia. He specializes in LPAR design and configuration, IT Maintenance, OS support and Hardware System Service for IBM System i5, iSeries, and AS/400.

He has planned and implemented LPAR and OS/400 V5R3 installation at three customer sites using System i5 model 520s. He has also planned and installed System i5 model 520, Model 570, and iSeries Model 870, all with LPAR.

You can contact him by sending an e-mail to: <mailto:syogi@id.ibm.com>

William Wei has worked for IBM China as iSeries Field Technical Sales Specialist for four years. He is the technical skill owner of LPAR configuration and implementation for i5/OS, Linux, and AIX®. He also works on WebSphere and Domino. He has planned and performed LPAR setup on iSeries for key IBM iSeries customers in China.

Since 1995, he has written 11 articles about information and computer technology in several journals. He also wrote Chapter 1 of the iSeries technology guides for users in China published in 2004, and he translated from English to Chinese Chapters 9 & 10 of “The Inside Story of the IBM iSeries” written by Frank G. Soltis and published in 2003.

You can contact him by sending an e-mail to: <mailto:weijh@cn.ibm.com>

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

IBM Rochester:

Dan Benke
Tim Covalt
Selwyn Dickey
Joe Gibbons
Jim Hermes
Mike Konkel
Edith Lueke
Mark Manges
Naresh Nayar
Jonathan Van Niewaal
Kris Ryan
Tracy Smith
Brad Voigt
Geoff Warren
Allyn Walsh
Craig Wilcox

IBM HMC Development:

Bob Foster
Andy McLaughlin
Duc Nguyen
Ming Nguyen
Bob Whitehurst

IBM ITSO, Austin Center:

Scott Vetter

IBM ITSO, San Jose Center:

Yvonne Lyon

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

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 Redbooks to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

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

ibm.com/redbooks

- ▶ Send your comments in an Internet note to:

redbook@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. JLU Building 107-2
3605 Highway 52N
Rochester, Minnesota 55901-7829

Archived

Summary of changes

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

Summary of Changes
for SG24-8000-01
for Logical Partitions on System i5
as created or updated on May 17, 2006.

May 2006, Second Edition

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

This IBM® Redbook was first produced in draft in April 2004 to match the i5 announcements and General Availability 1 (GA1).

It is important to remain current with both your HMC code level and the Service processor code level. You should review the latest information at:

http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/index.htm

New information

New chapters include:

- ▶ HMC duplication and redundancy
- ▶ Open SSH
- ▶ Virtual Partition Manager
- ▶ Firmware maintenance
- ▶ HMC Access Password Reset

Changed information

The following information has been changed:

- ▶ In this most recent update we have included changes made in HMC support, Service Processor code support, and available hardware models.
- ▶ We have removed VMware — the information remains available in *Implementing VMware ESX Server 2.1 with IBM TotalStorage FAStT*, SG24-6434.
- ▶ We have removed section 6.5.7, Concurrent Maintenance Power Domain. This function is not available in the HMC GUI, it is available within the ASMI GUI.
- ▶ Other chapters have been updated to reflect the DS4000 products and features that were current as of May 2006.

Archived



Introduction to LPAR on IBM System i5

This chapter provides an overview of the System i5 and its LPAR capabilities as explained in the following topics:

- ▶ HMC, Hypervisor, and partitions
- ▶ Software requirements
- ▶ Processor use in System i5 LPARs
- ▶ Memory use in System i5 LPARs

1.1 HMC, Hypervisor, and partitions

The IBM @server System i5 systems provide a new system architecture for logical partitioning (LPAR) and Capacity Upgrade on Demand (CUoD):

- ▶ The LPAR Hypervisor is now shipped as a firmware part of all @server System i5 models.

It is stored in the non-volatile random access memory (NVRAM) of the Service Processor. Previously, it was a part of the System Licensed Internal Code (SLIC) shipped with OS/400. Once loaded, the LPAR hypervisor runs in main memory.

As the Hypervisor is now independent from the operating systems, there is no longer a primary partition concept for LPAR. Memory is the only resource used by the Hypervisor.

- ▶ A *Hardware Management Console* (HMC) device is required to perform LPAR, Dynamic LPAR (DLPAR), and CUoD configuration and management.

This device is an IBM customized Linux appliance featuring an IBM NetVista specific workstation. It is pre-loaded with the HMC software and cannot be used for any other purpose or loaded onto another PC. A DVD-RAM drive is included for loading new versions of the HMC code, and backup and recovery. A diskette drive is also included for migration purposes.

The HMC is connected via the first Ethernet port of the @server. A system can be managed by two HMCs. A single HMC can manage 48 i5 systems with 254 partitions, in HMC terminology these are known as *Managed Systems*. A second HMC should be connected for redundancy.

In Figure 1-1 we show the new LPAR arrangement with System i5 systems, including the two types of HMC. In the server you can see the components: firmware, i5/OS partitions, a Linux partition, and an AIX partition. In this case the i5/OS partition is designated as a Service Partition. We describe this arrangement and components in more detail in a later chapter.

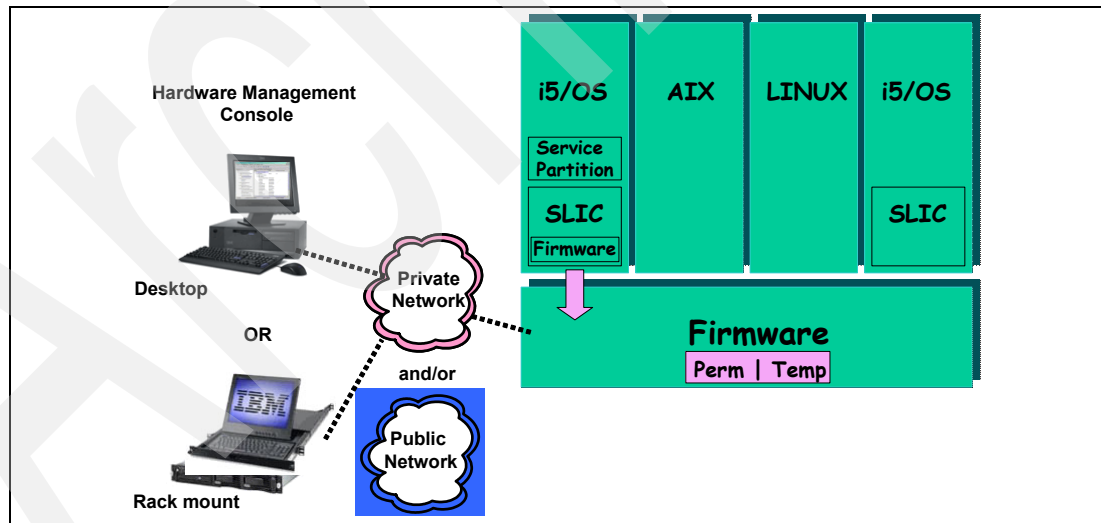


Figure 1-1 New Logical Partitioning arrangement for System i5 systems

A 5250 terminal function is provided to open an operating system console session on each managed i5/OS partition from the HMC. This is an additional choice to the current OS/400 Twinax, Operations Console, and LAN Console options, called the HMC type console.

An integrated modem is shipped with the HMC; this is to be used as a focal point for the IBM Service Agent function that can report problems to IBM Support. I/O error reporting requires an additional Ethernet link between the HMC and the @server.

The HMC is ordered as a required priced feature of any LPAR or CuOD configuration for new orders or upgrades (MES), or shipped as a mandatory part of all high-end models.

The new System i5 systems have a new scheme storing and managing partition information. In this new scheme you have Partition Profiles and System Profiles.

► Partition Profiles:

A partition profile is used to allocate resources such as processor units, memory and I/O cards to a partition. Several partition profiles may be created for the same partition, because no resource availability checking is performed when a partition profile is created. Partitions can be created with access to all resources with one single check box, or user defined resources. The use all resources check box is typically not used, since a use all resources arrangement would be for non-partitioned servers.

As any partition can be defined to have access to any resource, there are two ways to validate that partition profiles have no resource conflicts when creating an additional partition:

- Attempt to power on the second partition. An error will be given if there are conflicts, and depending on the resource conflict, the second partition may not start.
- Create a *System Profile* which is a collection of partition profiles. The creation of the system profile will validate the resources. A partition profile cannot be added to a system profile if the partition resources are already committed to another partition profile and are considered part of that system profile.

Processor and memory resources are specified as minimum, maximum, and a new type called *Desired*. If insufficient resources are available at partition activation time, the desired configuration is not guaranteed, and the partition could be started with the minimum option.

I/O resources can be configured as Required or Desired. A *Required* resource will prevent the partition from starting if it is not available, but a partition can start if a *Desired* I/O resource is not available.

Any OS/400 partition may be configured as a *Service Partition* to be used as an alternative or backup way of reporting errors from other OS/400 partitions to IBM Service when the HMC is not available.

► System Profiles

A System Profile is an ordered list of partition profiles for a particular System i5. Multiple system profiles can be stored on the HMC, allowing various partition configurations to be activated/deactivated from one point. To change from one system profile to another, all the running partitions must be shut down.

Several system profiles may be created to materialize different validated partition profiles to be activated. For example:

Restructuring resources without using DLPAR. If during over night hours, holiday, or period end processing, certain partitions were not being used (such as development), it is possible to define a system profile that takes unused resources and allocates them to another partition (such as production), as shown in Figure 1-2.

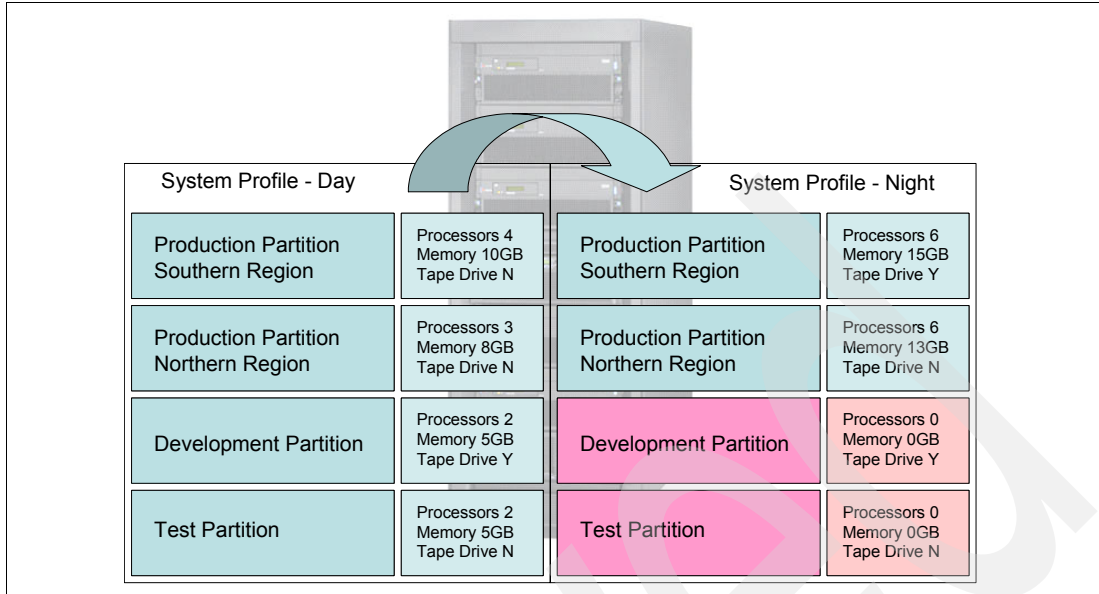


Figure 1-2 Reallocating resources with system profiles

Note: In many cases this reallocation of resources through system profiles may not be practical, as it requires all partitions to be stopped twice, once at the start of night processing and again to reset for the start of day processing.

- ▶ A disaster recovery site. Businesses have disaster recovery sites providing backup for multiple partitioned servers. It is possible to create a system profile for each server to be recovered, and then within that system profile, to define the partition profiles.

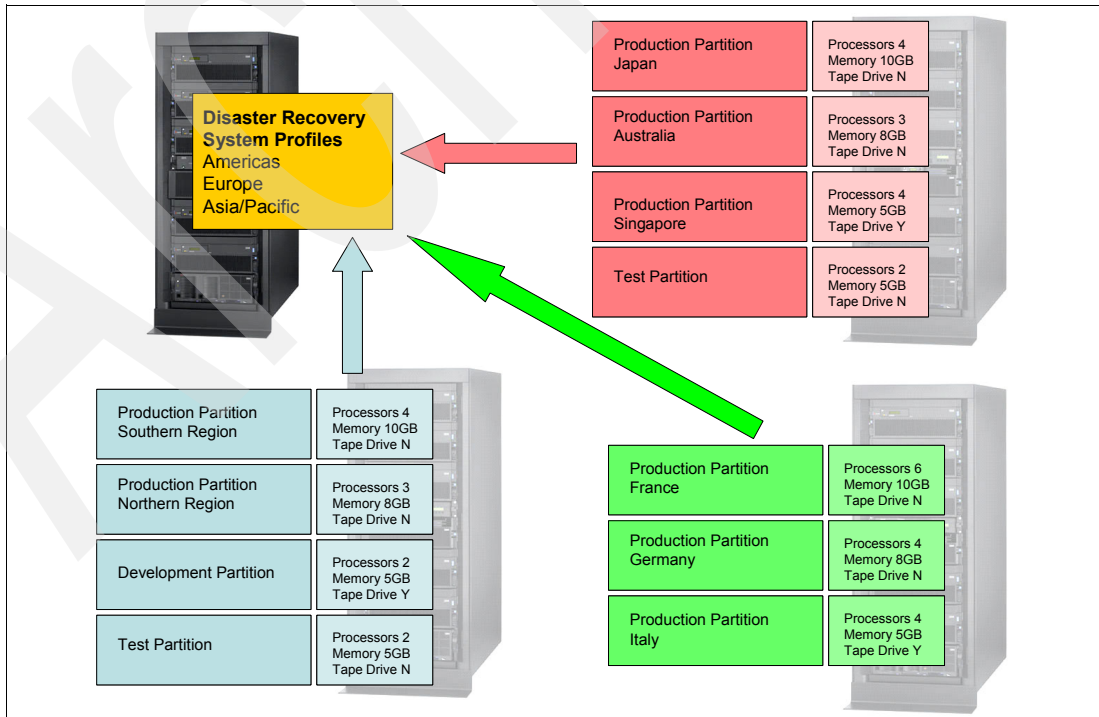


Figure 1-3 Disaster recovery site

This allows one system to adopt multiple personalities. In this type of scenario we are assuming the partitions will be reloaded at the disaster recovery site with no data on the disk.

1.2 Software requirements

The IBM @server System i5 systems require one of the following levels of operating system:

- ▶ OS/400 version 5 release 3 or later.
- ▶ AIX 5L™ version 5.2 with native IO support. AIX 5L version 5.3 will support hosted IO as well as native IO support.
- ▶ Linux version 2.6 for PowerPC®, which relates to Red Hat Enterprise Linux AS for POWER™ Version 3 and SUSE Linux Enterprise Server 9 for POWER. For more information on i5 Linux support, visit:

<http://www-1.ibm.com/servers/eserver/series/linux/>

1.2.1 What's new for licensing

These are some new licensing features:

- ▶ The number of OS licenses is now based on what the server is *configured* to use. Previously, the server price included the licenses for all startup processors.
- ▶ An OS/400 license key is now shipped to ensure compliance for the number of processors purchased versus used. It applies to licensed program 5722-SS1 feature 5051 of OS/400 V5R3.
- ▶ The LPAR Validation Tool (LVT) is adding a new function to report licensing requirements.

1.2.2 Licensing requirements

These are the basic rules for Operating System licenses on @server System i5 machines:

- ▶ For dedicated partitions, the number of licenses equals the *desired* number of processors configured.
- ▶ For capped partitions, the number of licenses equals the total of the *desired* number of processing units configured rounded up to the next whole number.
- ▶ For uncapped partitions, the number of licenses equals the *maximum* number of virtual processors configured, up to the limit of the number of *physical* processors in the shared pool.
- ▶ For On/Off Capacity on Demand (CoD), there are no additional licensing charges associated with a temporary processor activation.

1.2.3 Simple scenario with dedicated and shared capped partitions

Here is a simple example of a 4-way System i5 system with two OS/400 partitions using the shared processors pool and one AIX partition using 2 dedicated processors.

Table 1-1 System i5 system with two OS/400 partitions

Partition ID - OS	Partition Type	Processing Units - Virtual Processors	Licenses
P1 - OS/400	Shared capped	1.5	1.5 + 0.5 = 2
P2 - OS/400	Shared capped	0.5	
P3 - AIX	Dedicated	2	2

- ▶ As the total number of processing units configured for the OS/400 partitions using the shared processors pool equals 2.0, the number of OS/400 licenses required is 2.
- ▶ As the number of dedicated processors for the AIX partition is 2, the number of AIX licenses required is 2.
- ▶ Note that if the AIX partition is powered down, the unused dedicated processors are added to the shared processors pool, increasing the available processing units from 2.0 to 4.0.

However, the total number of processing units used by OS/400 partitions cannot exceed 2.0 to match the number of purchased licenses.

1.2.4 Normal scenario with shared capped and uncapped partitions

Here is a normal example of a 4-way System i5 system with two OS/400 uncapped partitions and one AIX capped partition all using the shared processors pool.

Table 1-2 System i5 system with two OS/400 uncapped partitions and one AIX capped partition

Partition ID - OS	Partition Type	Processing units - Virtual Processors	Licenses
P1 - OS/400	Shared uncapped	1.5 - VP=2	2 + 1 = 3
P2 - OS/400	Shared uncapped	0.5 - VP=1	
P3 - AIX	Shared capped	2.0	2

- ▶ As the total of maximum virtual processors configured for the uncapped OS/400 partitions is 3, the number of OS/400 licenses required is 3, although the desired processing unit total is 2.
- ▶ As the number of processing units configured for the AIX partition is 2.0, the number of AIX licenses required is 2.
- ▶ Note that here, unused processor resources configured for the AIX capped partition are available for use by the uncapped partitions, up to the maximum number of virtual processors configured.

1.2.5 Complex scenario with shared uncapped partitions

Here is a complex example of an 8-way System i5 system with two OS/400 uncapped partitions and one AIX uncapped partition, all using the shared processors pool.

Table 1-3 ISystem i5 system with two OS/400 uncapped partitions and one AIX uncapped partition

Partition ID - OS	Partition Type	Processing Units - Virtual Processors	Licenses
P1 - OS/400	Shared uncapped	4.0 - VP=7	7 + 3 = 10 but shared pool max = 8
P2 - OS/400	Shared uncapped	1.0 - VP=3	
P3 - AIX	Shared uncapped	3.0 - VP=6	6

- ▶ As the total number of virtual processors for OS/400 partitions is more than the number of processors available in the shared processors pool, only 8 licenses are required for OS/400 although the total number of virtual processors configured is 10.
- ▶ As the number of virtual processors configured for the AIX partition is 6, the number of AIX licenses required is 6.
- ▶ Note that here, it is a customer's choice to balance the cost of additional software licenses with idle processor capacity.

1.2.6 A more complex scenario with shared uncapped partitions and capacity upgrade on demand

Here is a more complex example of the same 8-way System i5 system used in the previous example, when we want to add dynamically one stand-by processor via CUoD for use by OS/400 partitions.

Table 1-4 ISystem i5 system from previous example, adding one stand-by processor via CUoD

Partition ID - OS	Partition Type	Processing Units - Virtual Processors	Licenses
P1 - OS/400	Shared uncapped	4.0 - VP=7	7 + 3 = 10 but shared pool now = 9
P2 - OS/400	Shared uncapped	1.0 - VP=3	
P3 - AIX	Shared uncapped	3.0 - VP=6	6

- ▶ First, to activate a CUoD processor, we need to enter the processor activation code delivered by IBM after the customer order is processed.
This is done on the HMC using the Activate option of the Capacity Upgrade on Demand tab on the managed system.
The activation of the stand-by processor materializes as an additional processor resource in the shared processors pool.
- ▶ We now have to enter the license key for the additional processor resource to be used by OS/400 partitions. The additional software license is generated by the IBM configurator if the order of the CUoD feature specifies OS/400 use.
This is done using the WRKLICINF command and selecting option 1, Add License Key for the product 5722-SS1 feature 5051 in both partitions P1 and P2.

1.3 Processor use in System i5 LPARs

In this section we discuss the basic concepts of dedicated processor, shared capped processor, and uncapped processor; how they work, how to configure in logical partition configuration, and several aspects regarding the use of shared capped and uncapped processors. Additionally, we also describe memory allocation for i5 OS logical partitions. These are some of the topics we cover:

- ▶ System i5 processor overview
- ▶ Logical partition resources
- ▶ Dedicated processor
- ▶ Shared processor
- ▶ Shared uncapped processor
- ▶ Shared capped processor
- ▶ Virtual processor
- ▶ Configuring dedicated processors for the logical partition
- ▶ Considerations on using shared processors
- ▶ Memory allocation for the i5 OS logical partition

1.3.1 System i5 processor overview

Basically, processors and memory are grouped into nodes. The System i5 system unit may contain more than one node. For example, System i5 system model 570 has two nodes which contain two processors per node as shown in Figure 1-4. Memory cards are plugged into each node in quad. These nodes are interconnected and can use each other memory. We recommend that you spread all memory cards evenly across all nodes for better system performance.

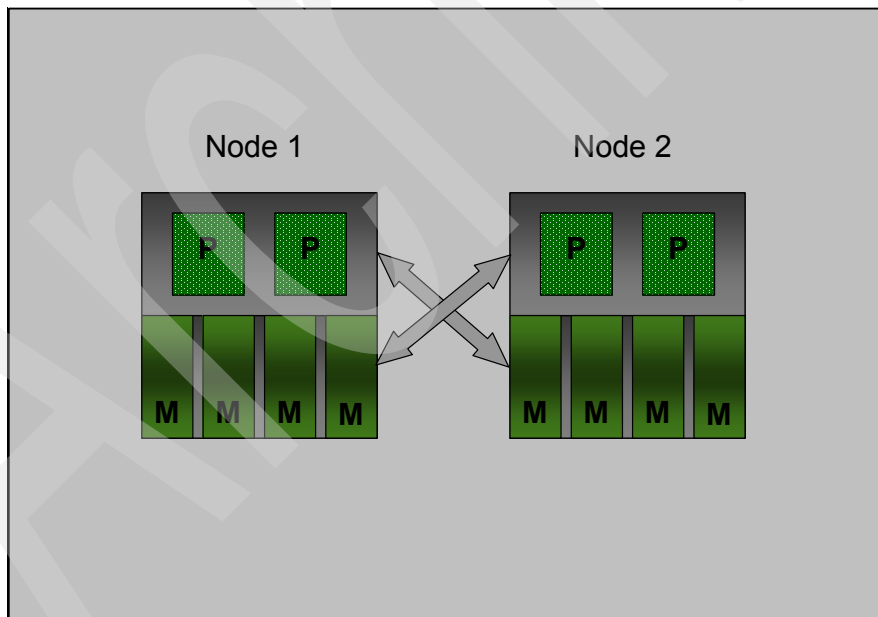


Figure 1-4 Interconnected processors and memory nodes in System i5 system

The iSeries processors cycle every 10 milliseconds. So the processors usage by all logical partitions are limited within this 10 millisecond cycle. After 10 milliseconds, the processor cycle will repeat automatically. For example, a logical shared partition with 0.6 processing units will have up to 6 milliseconds of processing time in every processor cycle (10 milliseconds).

If there are three processors in the System i5 system unit, with all three processors in the shared processors pool, if a logical partition is allocated 0.8 processing units, then the logical partition will have 8 milliseconds of processing time from 30 milliseconds of processing time. For every 10 millisecond time slot, you give 8 milliseconds out of 30 milliseconds of processing time available.

Figure 1-5 illustrates the processing time assignment from the shared processors pool to a logical partition. The same condition applies for a logical partition with 1.5 processing units. The logical partition will have 15 milliseconds of processing time available out of 30 milliseconds of processing time every 10 millisecond time slot. However, this logical partition needs 2 virtual processors in the shared processors pool.

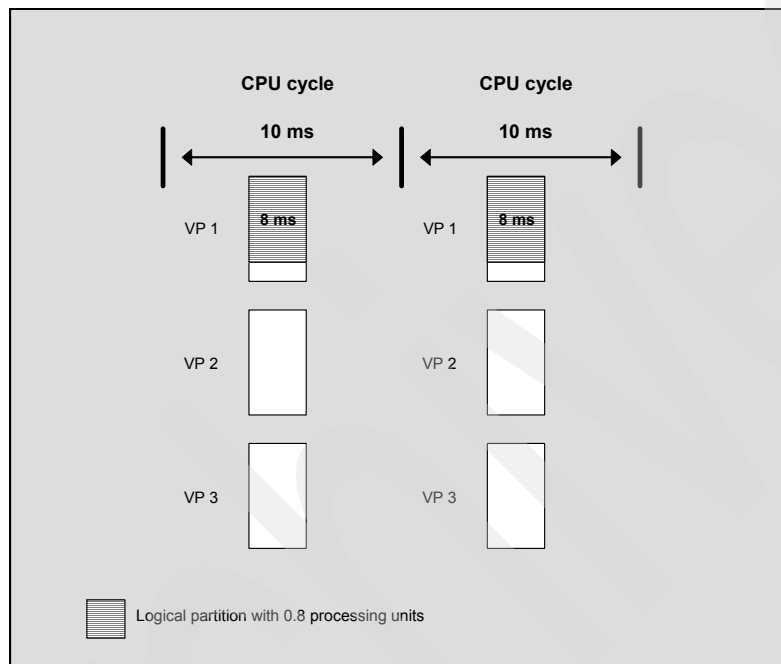


Figure 1-5 Processing time for logical partition with 0.8 processing units

Once the logical partition has used its processing time in every 10 millisecond time slot, it has to wait for next the available processing time in the next 10 millisecond physical time slot.

For a dedicated logical partition, the processing time available for that partition is equal to a number of physical processors assigned multiplied by 10 milliseconds. A logical partition with a dedicated processor has better processing time utilization than a shared logical partition. For capped logical partitions, it cannot exceed the processing time assigned to them. For uncapped logical partitions, it is possible to use idle processing time in the shared processors pool.

Before a processor can run a job, it must first get the data from DASD to main storage/memory. The data from memory is then transferred into the cache memory. The data in the cache memory is available to use by the processor. The process of moving data from DASD into the memory is slower than reading data from cache memory, so it will take a longer time to load data from DASD to cache for the first time. Frequently accessed data will remain in the cache or memory in order to maintain fast data access. Figure 1-6 below illustrates the basic flow of data loaded from DASD into memory, then into the cache memory, before being used by the processor.

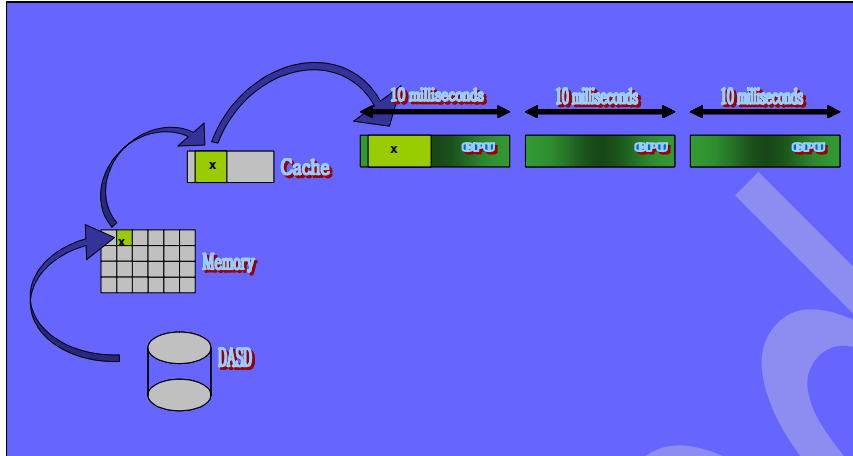


Figure 1-6 Basic flow of data from disk to processor

Next we show an example of two shared capped logical partitions sharing a single processing unit from the shared processors pool. There are two logical partitions, logical partition A and logical partition B. Logical partition A has two jobs, A1 and A2. Logical partition B has two jobs, B1 and B2. Logical partition A is assigned 0.8 processing units. Logical partition B is assigned 0.2 processing units. Figure 1-7 illustrates the processor time usage every 10 millisecond time slot. This example is only used to show concepts.

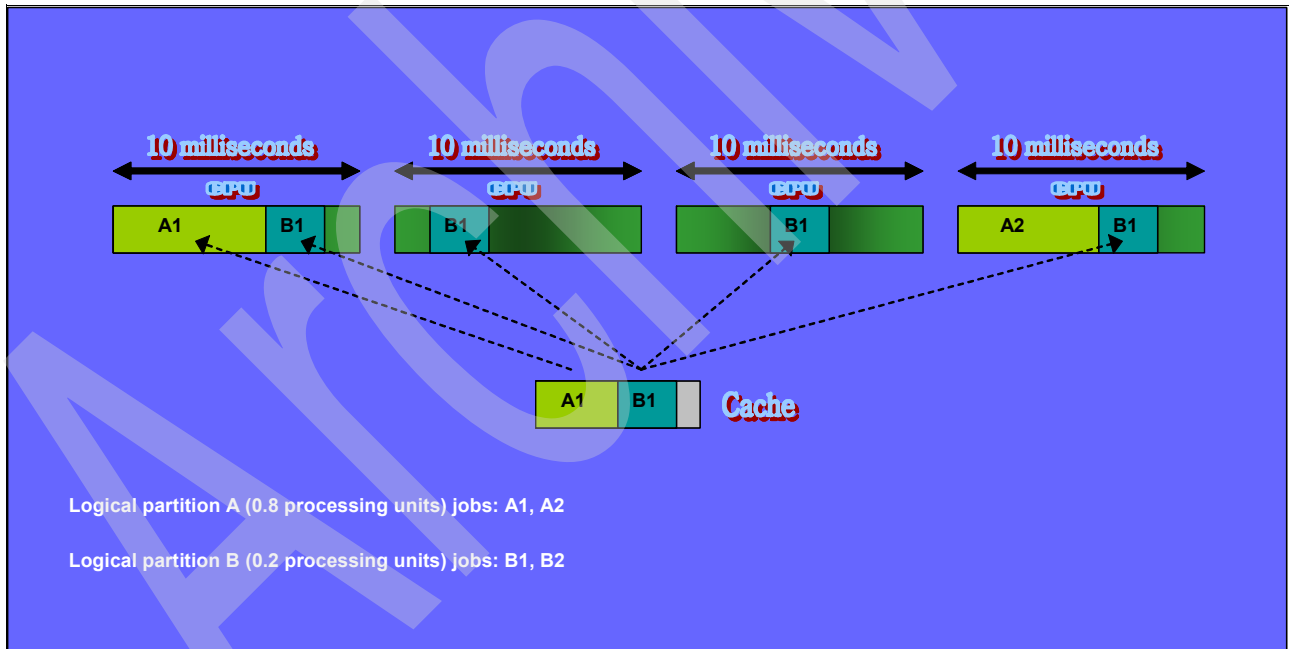


Figure 1-7 Two logical partitions sharing one processing unit

Jobs A1 and B1 are loaded into the cache memory. Since only one job can run on the given processor, job A1 loaded first into the processor, as it is a high priority job. Job A1 runs for 7 milliseconds and is complete. Job B1 is dispatched to the processor and runs for 2 milliseconds. Job B1 must stop after 2 milliseconds because logical partition B has only 0.2 processing units assigned to it. Job B1 must wait for the next processor cycle to be dispatched again into the processor to complete. In the next processor cycle, job B1 is resumed and running.

However, job B1 must stop after 2 milliseconds in the second processor cycle again because the processing units assigned are very small. In the next processor cycle, job B1 is resumed again, and has to stop after 2 milliseconds. In the next processor cycle, job A2 is submitted and finishes after 6 milliseconds. Jobs B1 is resumed again, being processed for 2 milliseconds, and is complete.

This example shows two logical partitions sharing a single processing unit from a shared processors pool. A logical partition with adequate processing units assigned to it will have its jobs complete faster than if it had smaller processing units. It is very important to determine the processing units required by a logical partition for better logical partition performance.

Simultaneous Multi Threading (SMT)

System i5 Processors implement Simultaneous Multi Threading (SMT) technology to process the instructions given. The instructions from the operating system are loaded simultaneously into the processor and executed. When the instructions are loaded simultaneously into the processor, they have to wait their turn for registers that needed by each instruction at the same time. This makes total time longer to execute each instruction, but total time to execute all the instructions loaded simultaneously is shorter. In this case, the SMT technology implemented in the System i5 processor results in greater system capacity because multiple instructions can be executed at shorter time.

During periods of low user requests, instructions are infrequently executed and reach the processor to be executed. In this period, there will be only one thread in the processor pipes at any point in time. The absence of contention allows the processor to complete each instruction faster. When demand for high processing power arises due to an increase in the frequency of user requests, SMT technology allows greater parallelism in the processing of instructions and reduces overall elapsed time for multiple requests.

The Simultaneous Multi Threading (SMT) is controlled by the QPRCMLTTSK system value of the OS/400 operating system. Each operating system in the logical partitions can define its own QPRCMLTTSK system value independently. The logical partition must be restarted in order to apply changes in QPRCMLTTSK system value.

There are three available values to set the SMT in QPRCMLTTSK system value. Value '0' will turn off SMT; the result is that there will be only one task per physical processor. Value '1' will turn on SMT and make two tasks available per physical processor. Value '2' will turn on SMT, but the operating system might switch to single thread mode.

When SMT is turned on, up to two threads are dispatched to each physical processor, thus resulting in an environment with two logical processors. If SMT is turned off, only one thread will be dispatched and the system will only operate in an environment with one logical processor. When the value of '2' is given to QPRCMLTTSK, which means that the SMT environment control is given to the system, the system will monitor the rate of instruction execution requests to the processor and switch between the two environments to achieve the best system performance.

1.3.2 Logical partition resources

Logical partitioning on the System i5 system allows us to create several independent logical partitions with their own processors, memories, and I/O resources. The logical partition profile defines these resource allocations for each logical partition. Each logical partition can run independently within the System i5 system.

The Power Hypervisor allocates the Server processors and memory during startup and during dynamic reallocation of processor and memory using Dynamic Logical Partitioning (DLPAR). Each operating system (i5/OS) allocates the processors and the memory available

for the partition to the tasks submitted for processing. The Power Hypervisor attempts to put a partition back onto the last processor it used. However, if that processor is currently busy or not available, it will go to any available one.

Resources of one partition are isolated from other partition resources of other logical partitions, but there are ways for these active logical partitions to share their processing power, memory, and I/O resources to another active logical partition. Dynamic LPAR/DLPAR allows us to add, move, or remove processors, memory, and I/O resources to, from, or between active partitions manually without having to restart or shut down the logical partition.

The LPAR only allocates physical processors to every logical partition. Micro-partitioning allows logical partitions to share the processors in the shared processors pool. The shared processors pool includes all processors on the server that are not dedicated to specific logical partitions. Each logical partition that uses the shared processors is assigned a specific amount of processing units from the shared processors pool. A logical partition that uses the shared processors pool can always use the processing units assigned to them. However, if the logical partition demands more processing power than its assigned amount, the logical partition is set by default to use the unused processing units in the shared processors pool.

The number of processing units that the logical partition can use is limited only by the virtual processor setting of the logical partition and the amount of unused processing units available in the shared processors pool.

The Power Hypervisor will remember the last active configuration on the server for the next managed system power up to standby or operating mode. The Power Hypervisor will assign the resources for dedicated partitions to provide the optimum configuration for the dedicated logical partition.

In order to create optimum resource allocation for logical partitions, the Power Hypervisor will assign processors to the fewest number of nodes possible in the system. The reason is that processors in a node share the same L2 and L3 cache. Using the processors in the same node will improve logical partition performance since each processor will have better chance to find data in the cache.

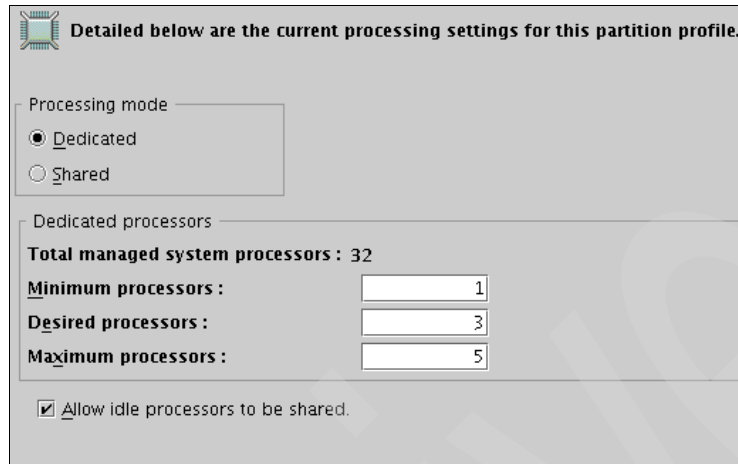
The memory portion dedicated for the Hardware Page Table (HPT) will be taken from memory in a node that contains processors. The Power Hypervisor will assign as much memory as possible in the same node as the processors. This is done to avoid too many inter-node *remote memory access* operations, which may slow down the logical partition operation. The rest of the processor resources will be assigned to the shared processors pool.

1.3.3 Dedicated processor

Dedicated processors are whole processors that are assigned to a single partition. If you choose to assigned dedicated processors to a logical partition, the minimum number of dedicated processors you must assign is one processor. If you choose to use dedicated processors, once the logical partition is activated, only that partition can use the processors. You cannot assign a partial processor to a logical partition that will use a dedicated processor.

When the logical partition with dedicated processors is activated, the desired processor resource will be assigned to that partition. If the logical partition demands more processor resources, it will get only from the minimum up to the desired processor resources value. If there is a processor failure, the system will attempt to accommodate the minimum processor sizes for all partitions. If the minimum processor resource is not available for the logical partition, the logical partition cannot be activated.

When a partition with dedicated processors is powered down, their processors will be available to the shared processors pool. This capability is enabled from the partition profile setting. Check the **Allow idle processors to be shared** option to enable this feature as shown in Figure 1-8. If this feature is not selected, the dedicated processors from inactive logical partitions will not be available in the shared processors pool. Idle processors from active partitions with dedicated processors can be used by any uncapped partition that requires additional processing units to complete its jobs.



Detailed below are the current processing settings for this partition profile.

Processing mode

Dedicated

Shared

Dedicated processors

Total managed system processors : 32

Minimum processors :

Desired processors :

Maximum processors :

Allow idle processors to be shared.

Figure 1-8 Allow idle processors to be shared option in partition profile

If logical partition configuration allows the dedicated processors to be shared in the shared processors pool, these dedicated processors are still committed to the logical partition that these processors belong to. When the inactive logical partition is activated, it will regain its dedicated processors from the shared processors pool.

You can add more processors for a logical partition using Dynamic Logical Partition functions. However, this is limited by the maximum number of processors for this logical partition. Changing the maximum processors for the current logical partition requires a logical partition restart.

Dynamically moving processors in and out of a logical partition with dedicated processors could result in processors in a partition being spread across nodes. There is no guarantee that the best configuration can be maintained. The optimum allocation of processors and memory is performed by the Power Hypervisor during the initial startup to standby or operational mode when you start or power on the managed system. Any changes to the configuration of partitions after this through Dynamic Logical Partitioning (DLPAR) could disrupt the optimum allocation.

A full system restart is required to return the system to the optimal configuration. Set the system configuration to the desired configuration before restarting.

1.3.4 Shared processor

The shared logical partition can use unassigned processors and processors that are assigned to other logical partitions but are currently unused by that partition. The amount of processing power that the logical partition can use is limited only by the virtual processor settings of the logical partition and the amount of unused processing units available in the shared processors pool.

The shared processors pool is created from the processors left over after the dedicated processors are assigned to logical partitions that use dedicated processors. This will typically be spread over multiple nodes. These processors are shared among all the partitions that use shared processors. You can allocate at least 0.1 of a shared processor or up to the total number of processors in the system.

One single physical processor is equal to 1.00 processing units, two physical processors are equal to 2.00 processing units, and so on. The number of processing units is equal to the number of physical processors available in the shared processors pool. A processor must be in the shared processors pool to become processing units.

Every logical partition that uses shared processors will receive some processing units from the shared processors pool. These processing units will appear to the operating system as a number of processors called virtual processors. The Power Hypervisor is responsible to communicate the virtual processors' information to the operating system. The number of virtual processors for each logical partition that uses the shared processor is configured in the logical partition profile.

In partitions using the shared processors pool, if a partition is ended and restarted, the Power Hypervisor will attempt to re-assign the same processor(s). This is to attempt to increase cache hits for the workload on that partition.

1.3.5 Shared uncapped processor

By default, logical partitions that use the shared processors pool are uncapped logical partitions, which means that they can use unused processing units from the shared processors pool. The distribution of unused processing units to uncapped partitions will be based on the weight assigned to each partition if more than one uncapped partitions demand additional processing units at the same time.

Uncapped weight is a number in the range of 0 through 255 that you set for each capped partition in the shared processors pool. By setting the uncapped weight, any available unused capacity is distributed to contending logical partitions in proportion to the established value of the uncapped weight. The default value for uncapped weight is 128.

For example, if there are three processors available in the shared processors pool, and if two uncapped logical partitions demand more processing units at the same time, if logical partition A has an uncapped weight of 80 and logical partition B has an uncapped weight of 160, then logical partition A will receive 1 processing unit and logical partition B will receive 2 processing units from the shared processors pool. The uncapped weight is set in the partition profile as shown in Figure 1-9. You can set this uncapped weight when the partition profile is being created.

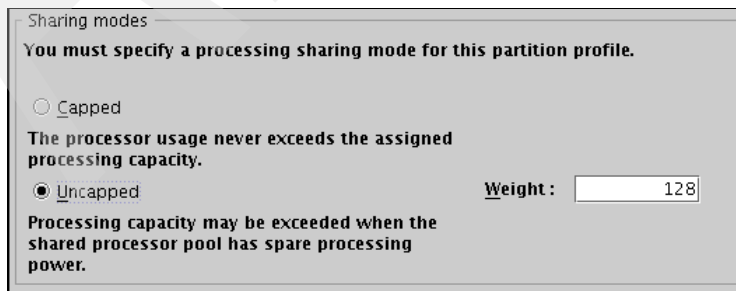


Figure 1-9 Uncapped weight setting in partition profile

It is possible for uncapped partitions to have their processing capacity exceed the current processing capacity or more than 100% CPU utilization when they use the unused processing units from the shared processors pool. This usually happens when the uncapped partitions demand more processing power to complete their tasks. It is different in the case of a partition with a capped processor. A capped partition will never exceed the assigned processing capacity.

Although an uncapped logical partition can use more processor power than its assigned processing capacity, the uncapped logical partition can never use more processing units than its assigned number of virtual processors.

Attention: When you try to add more processing units for an uncapped logical partition, the number of desired physical processing units as the result of the addition must be equal to or less than the desired number of virtual processors.

1.3.6 Shared capped processor

It is possible to configure a logical partition so that other logical partitions can use its unused processing units, but the logical partition cannot use more processing units than its assigned amount. Such a logical partition is known as a capped logical partition. The sharing modes in the logical partition profile, as shown in Figure 1-10, determine if the logical partition is using shared capped processors.

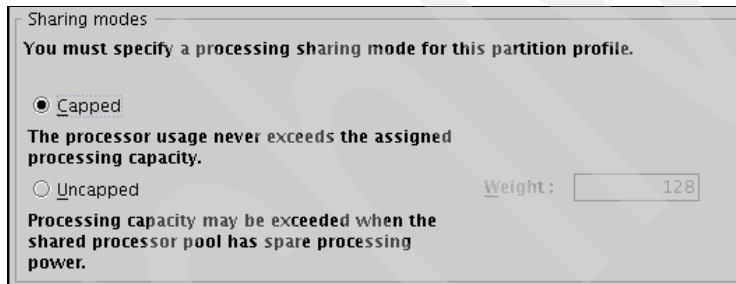


Figure 1-10 Capped processing sharing mode

Capped partitions cannot use idle processors in the shared processors pool when it demands more processing power to complete their jobs. Additional processors must be allocated manually using Dynamic Logical Partitioning (DLPAR) to a capped partition up to its maximum virtual processor setting.

Attention: When you try to add more processing units for a capped logical partition, the number of desired physical processing units as the result of the addition must be equal to or less than the desired number of virtual processors.

1.3.7 Virtual processor

A virtual processor is a representation of a physical processor to the operating system or logical partition that uses the shared processors pool. A virtual processor is a method for presenting fractional processors assigned to logical partitions as a whole number of processors, because the operating system cannot calculate a whole number of processors from fractional processors that are assigned to this logical partition. The Server firmware is responsible for representing the processing units available for logical partitions into a whole number of virtual processors.

The server firmware distributes the processing units evenly among all virtual processors assigned to a logical partition. For example, if logical partition A has 1.4 processing units and 2 virtual processors assigned, then each virtual processor will be equal to 0.7 physical processing units. These 2 virtual processors will support the logical partition workload.

The number of processing units available for each virtual processor is limited. The minimum number of processing units for each virtual processor depends on the server model on which the logical partitions are created. The maximum number of processing units for each virtual processor is always 1.00.

The logical partition operations will be faster and produce better results if the number of virtual processors assigned is very close to the number of processing units from fractional processors assigned to that logical partition. The operating system in the logical partition will be able to manage the workload more effectively.

In some cases, you may increase the number of virtual processors for the logical partition in order to give more concurrent operations for the operating system. The addition of some virtual processors may slightly improve the performance of the operating system. The operating system will be able to shift processing powers among processes very well.

However, you have to consider the impact of adding too many virtual processors for a shared processor logical partition, since this will degrade the performance of each virtual processor. If there are too many virtual processors for a logical partition whose processing unit is not increased, the number of processing units per new virtual processor is smaller than for the previous virtual processor. Then the operating system cannot shift processing powers among processes very well.

When you create a partition profile with shared processors, HMC will calculate the number of minimum, desired, and maximum virtual processors for the logical partition based on the minimum, desired, and maximum processing units assigned to this logical partition. By default HMC will calculate these default virtual processors as follows:

- ▶ The minimum virtual processor for logical partitions is the number of minimum processing units rounded to the next whole number.
- ▶ The desired virtual processor for logical partitions is the number of desired processing units rounded to the next whole number.
- ▶ The maximum virtual processor for logical partitions is the number of maximum processing units divided by minimum processing units for each virtual processor. The result is rounded to the previous whole number.

An example of virtual processor calculation done by HMC is shown in Figure 1-12.

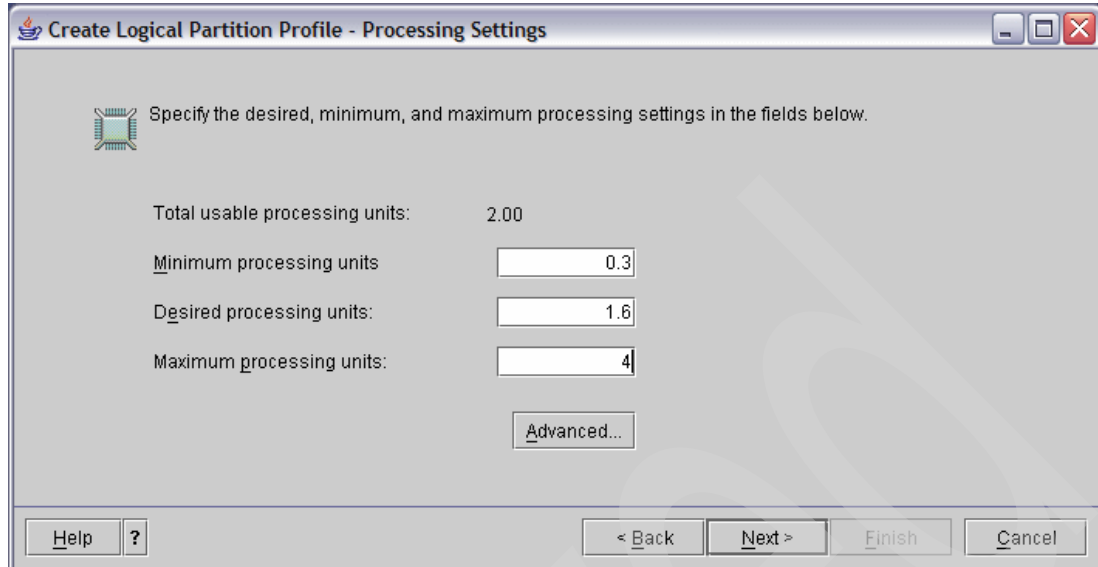


Figure 1-11 Shared processor processing units setting for logical partition

In the foregoing example, we give the minimum processing units of 0.3, desired processing units of 1.6, and maximum processing units of 4. Click the **Advanced...** button to display shared processor mode and virtual processor setting window as shown in Figure 1-12.

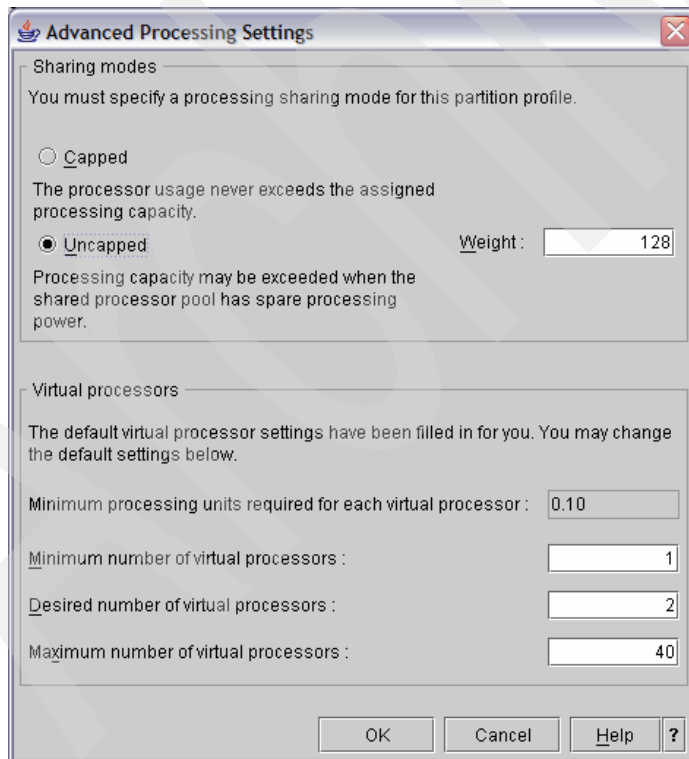


Figure 1-12 Virtual processor setting window

From Figure 1-12, you can see the number of minimum, desired, and maximum virtual processors calculated by HMC per the rules as explained previously. HMC calculates 1 minimum virtual processor, 2 desired virtual processors, and 40 maximum virtual processors. You may change these virtual processors settings to satisfy your workload requirements.

Configuring the virtual processors setting is similar to telling the logical partition how many processors it can run jobs on simultaneously. Figure 1-13 shows an illustration of the Virtual Processor configuration. The figure shows the difference of the implementation of virtual processor configuration of one shared logical partition with 2 virtual processors and with 4 virtual processors. There are 4 processors available in the shared processors pool. The logical partition is assigned 1.6 processing units out of 4 processing units from the shared processors pool.

Please note that creating many virtual processors may slow down your logical partition operation. Please calculate the proper number of virtual processors for your logical partition. Click **OK** to save changes.

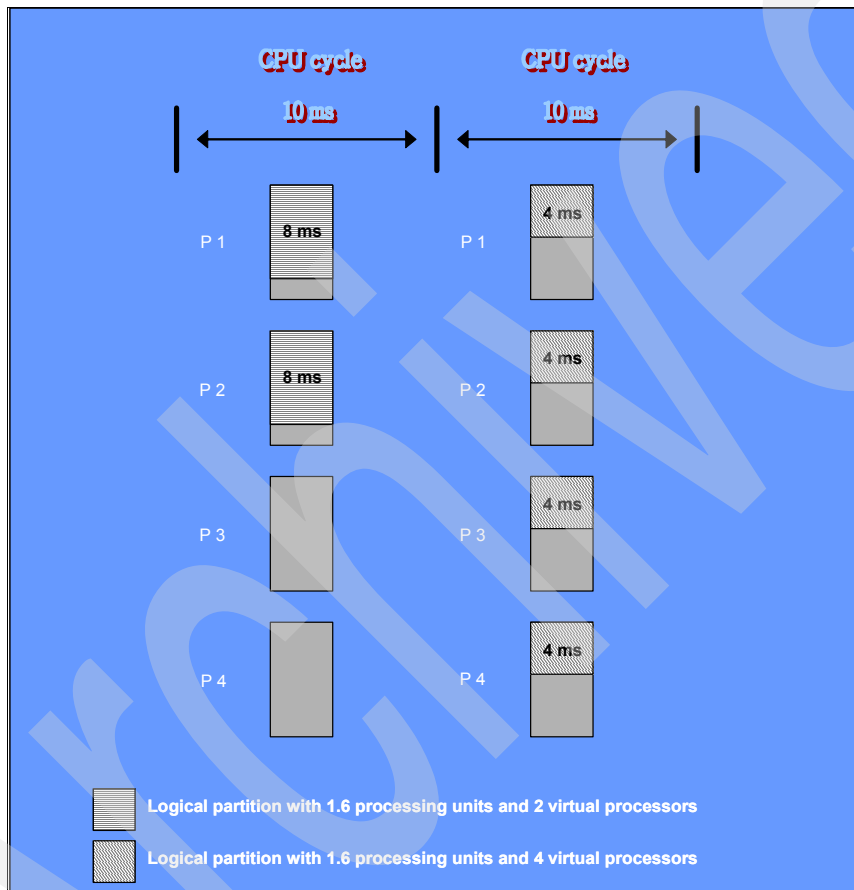


Figure 1-13 A logical partition with two and four virtual processors

When you activate the logical partition using a partition profile with a shared processor, the desired number of virtual processors will be assigned to the logical partition. Based on the configuration in Figure 1-12 on page 17, the logical partition will have 2 virtual processors available to support its workload.

You may increase the number of processing units and virtual processors for a logical partition using Dynamic Logical Partitioning (DLPAR).

For example, suppose your current logical partition processor setting is similar to the setting in Figure 1-14. Then if you want to add 2 additional virtual processors to this logical partition, you can use Dynamic Logical Partitioning to add processing units:

Selected → **Dynamic Logical Partitioning** → **Processor Resources** → **Add**

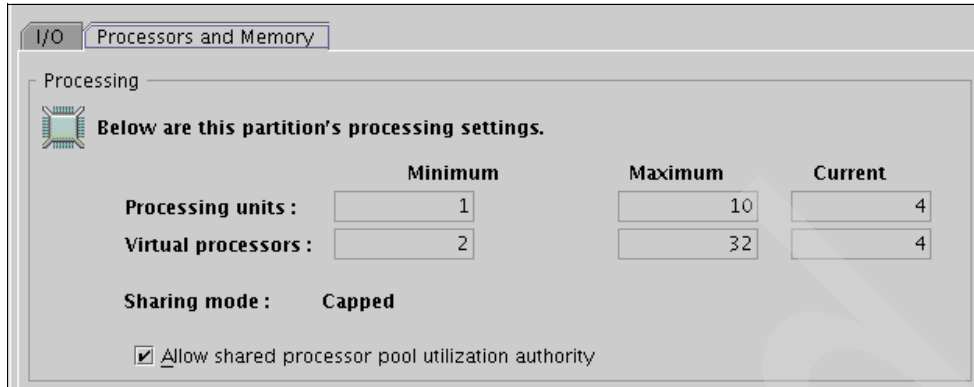


Figure 1-14 Logical partition profile processing unit configuration

Enter the number of virtual processor you want to add, and click **OK** to add new virtual processors. Figure 1-15 shows how to add 2 virtual processors to the logical partition. Click **OK** to add virtual processors.

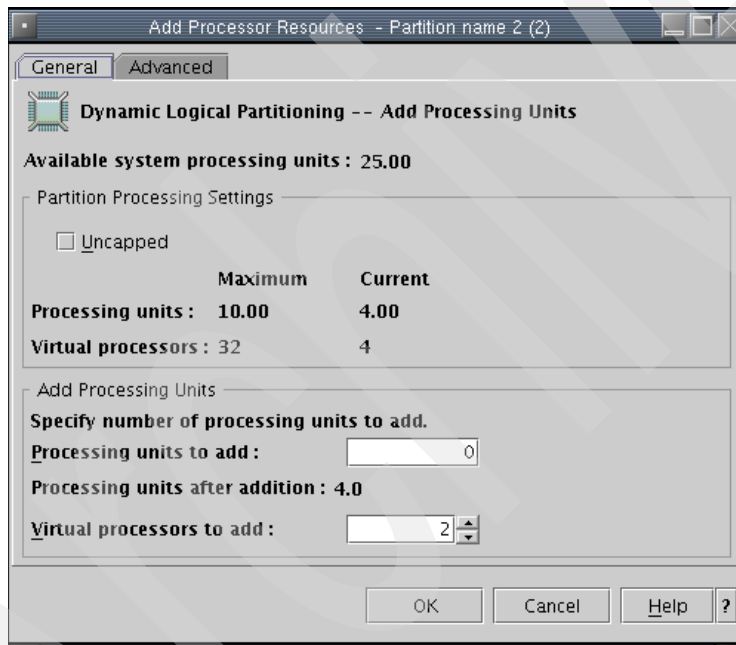


Figure 1-15 Add new virtual processors to logical partition

As a result, now the logical partition has 6 virtual processors (Figure 1-16).

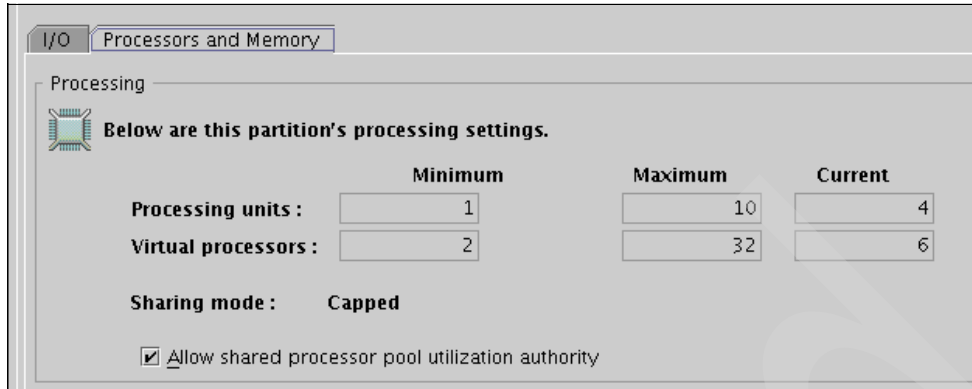


Figure 1-16 New number of current virtual processors after addition

Attention: When you try to add more processing units for a shared processor logical partition, the number of desired physical processing units as the result of the addition must be equal to or less than the desired number of virtual processors.

1.3.8 Configuring dedicated processors for the logical partition

Processor configuration is stored in the logical partition profile. When the logical partition is started, a number of physical processors are assigned to this logical partition. When started, the logical partition will be assigned the desired amount of processors. If the desired amount of processors is not available, the minimum amount of processors will be assigned to this logical partition. If no physical processors are available — for example, because of the failure of physical processors — the logical partition cannot be activated.

At the first time the logical partition is created, you have to define the processor configuration for the logical partition. To configure the logical partition with dedicated processors, select **Dedicated from Create Logical Partition Profile - Processors** as shown in Figure 1-17.

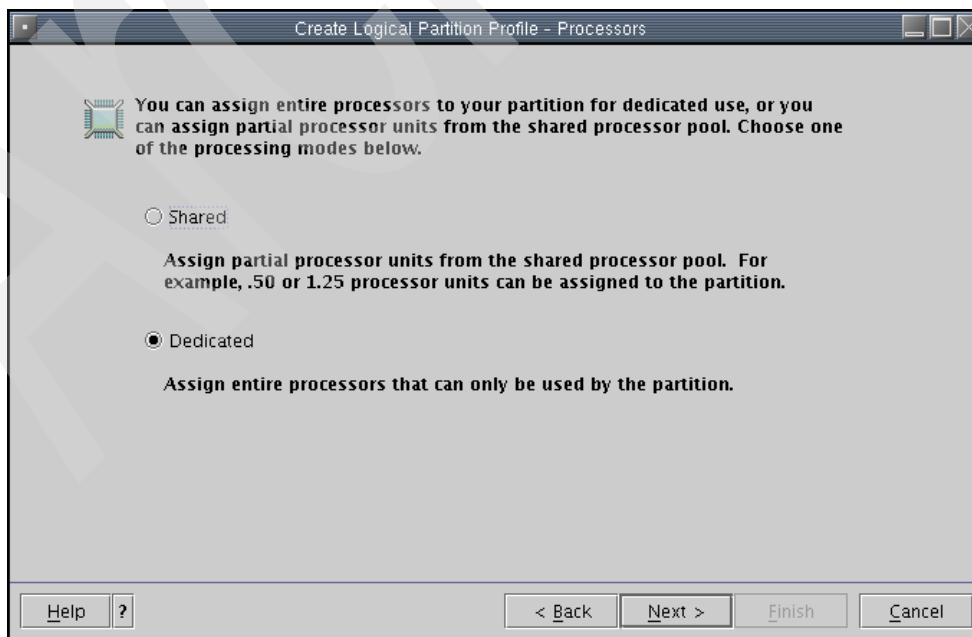


Figure 1-17 Dedicated processing mode for logical partition with dedicated processors

Click the **Next** button to continue.

Enter minimum, desired, and maximum processors for the logical partition. Figure 1-18 shows an example of minimum, desired, and maximum processors for the dedicated logical partition. Adjust these values with the logical partition workload.

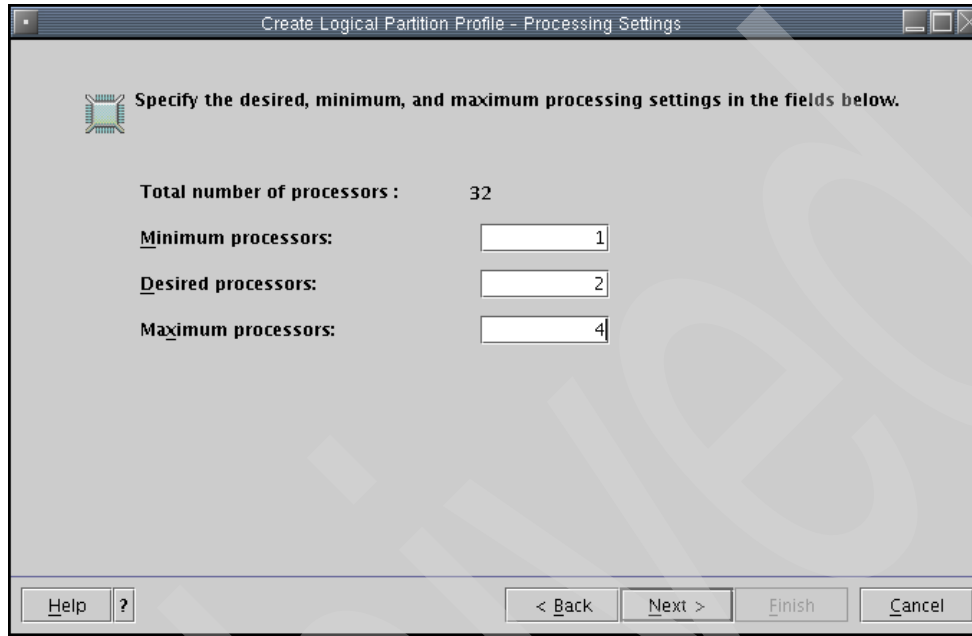


Figure 1-18 shows a dialog box titled "Create Logical Partition Profile - Processing Settings". The dialog box contains a message: "Specify the desired, minimum, and maximum processing settings in the fields below." Below the message, there are four rows of labels and input fields: "Total number of processors : 32", "Minimum processors: [input field with 1]", "Desired processors: [input field with 2]", and "Maximum processors: [input field with 4]". At the bottom, there are buttons for "Help ?", "< Back", "Next >", "Finish", and "Cancel".

Figure 1-18 Minimum, desired, and maximum processors for dedicated logical partition

1.3.9 Configuring shared capped processors for logical partition

The shared processors configuration is stored in the logical partition profile. When the shared logical partition is started, a number of processing units assigned to the logical partition will be based on the shared processors configuration in the logical partition profile. These processing units are available from the shared processors pool. When started, the logical partition will be assigned the desired amount of processing units. If the desired amount of processing units is not available from the shared processors pool, the minimum number of processing units will be assigned to the logical partitions. If no processing units are available from the shared processors pool, the logical partition cannot be activated.

At the first time the logical partition is created, you have to define the processor configuration for the shared logical partition. You have to decide if you will use shared capped processors or shared uncapped processors mode. To configure shared capped processors for the logical partition, select **Shared** from the Create Logical Partition Profile - Processors window, as shown in Figure 1-19.



Figure 1-19 Shared processing mode for logical partition with shared processors

Next, fill in the amount of minimum, desired, and maximum processing units for the logical partition. The example in Figure 1-20 shows an example of minimum, desired, and maximum processing units for the logical partition. Adjust these processing unit values with the logical partition workload.

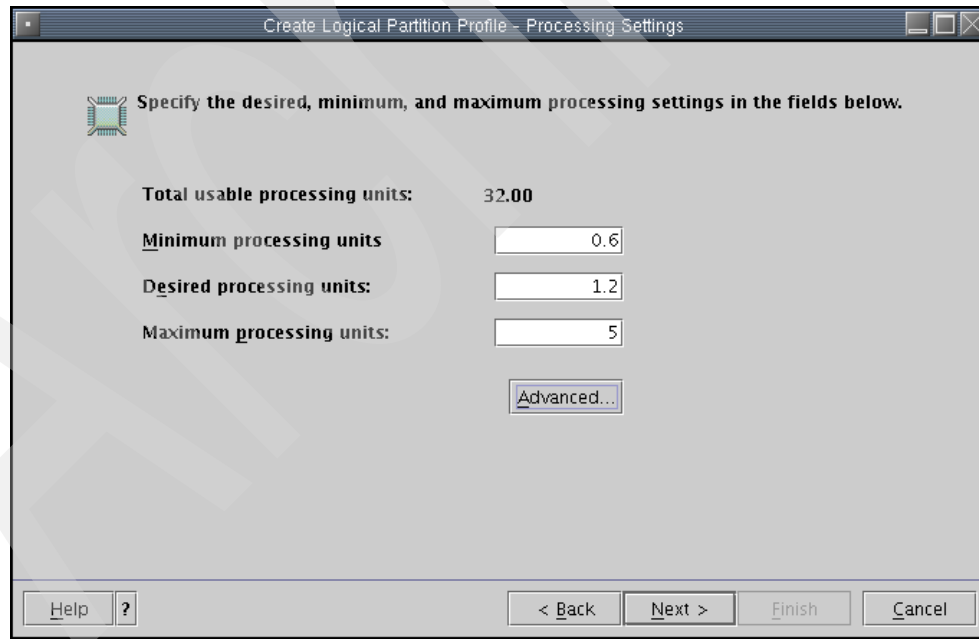


Figure 1-20 Minimum, desired, and maximum processing units for logical partition

Click the **Advanced** button to configure Capped/Uncapped sharing modes and virtual processors. Select **Capped** sharing mode to configure the logical partition with shared capped processors as shown in Figure 1-21.

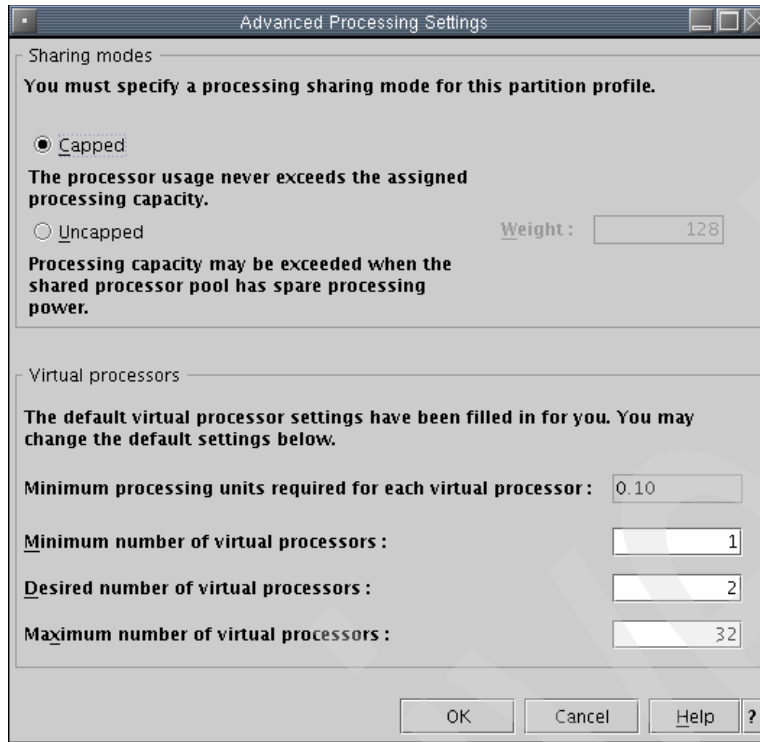


Figure 1-21 Capped sharing mode for shared capped processor configuration

HMC will automatically calculate the minimum, desired, and maximum number of Virtual Processors for logical partition. You may change the virtual processors setting now or later using Dynamic Logical Partitioning (DLPAR).

1.3.10 Configuring shared uncapped processors for logical partition

The shared processors configuration is stored in the logical partition profile. When the logical partition is started, a number of processing units assigned to the logical partition will be based on the shared processors configuration in the logical partition profile. These processing units are available from the shared processors pool.

When started, the logical partition will be assigned the desired amount of processing units. If the desired amount of processing units is not available from the shared processors pool, the minimum processing units will be assigned to logical partitions. If no processing units are available from the shared processors pool, the logical partition cannot be activated.

At the first time the logical partition is created, you have to define the processor configuration for the logical partition. To configure shared uncapped processors for the logical partition, select **Shared** from the Create Logical Partition Profile - Processors window as shown in Figure 1-19 on page 22, then select **Uncapped** from the Advanced Processing Setting as shown in Figure 1-22.

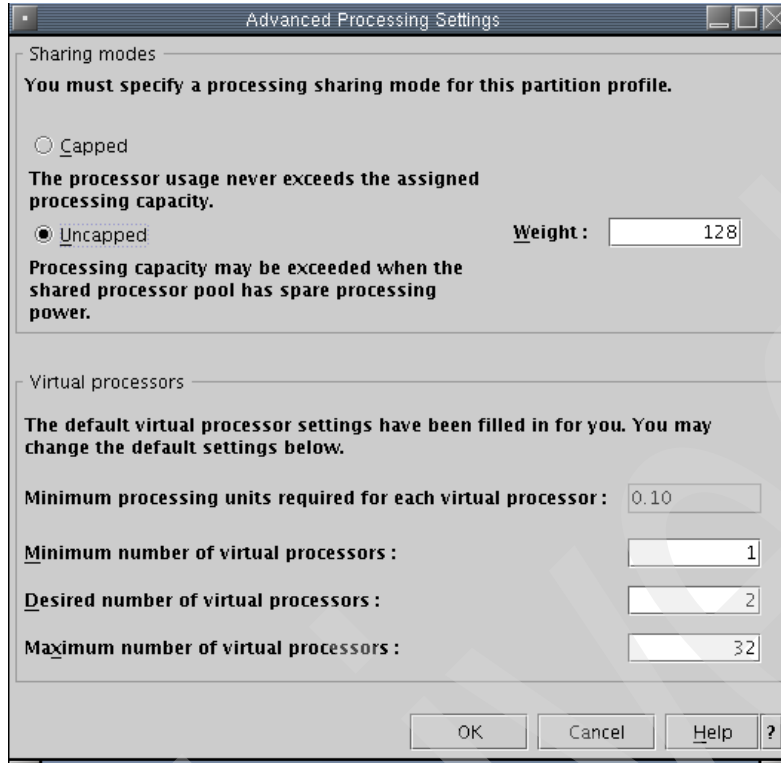


Figure 1-22 Uncapped sharing mode for shared uncapped processor configuration

Adjust the uncapped weight for this logical partition. The uncapped weight will be used to determine the portion of free processing units that will be distributed to this logical partition among all shared uncapped logical partitions when two or more shared uncapped logical partitions demand more processing units from the shared processors pool.

For example, if logical partition 1 has an uncapped weight of 90 and logical partition 2 has an uncapped weight of 180, then logical partition 1 will receive 1/3 available processing units from the shared processors pool when partition 1 and partition 2 demand more processing units at the same time and there are available processing units in the shared processors pool.

HMC will automatically calculate the minimum, desired, and maximum number of Virtual Processors for logical partition. You may change the virtual processors setting now or later using Dynamic Logical Partitioning (DLPAR).

1.3.11 Considerations on using shared processors

Logical partitioning using shared processors allows greater flexibility and maximum processor resources utilization. However, you must understand that using shared processors incurs some risks and may impact logical partitions performance. Using shared processors from the shared processors pool makes much more work for the Power Hypervisor.

Each logical partition with shared processors is only given some processing time in a few milliseconds, determined from the processing units assigned to it — usually less than processing times of one physical/dedicated processor. So, the logical partition does not always have enough time to complete its job in one CPU cycle. This problem causes the job's data to be re-dispatched to memory and cache for the next CPU cycle. It increases the chance that the data is no longer in the cache, or even in the memory any more, because much data for many jobs is loaded into memory, then into cache. It must be reloaded from memory, or at the worst case, from disk (DASD), which may take a longer time to complete.

If you could determine the processor requirement for your partition workload and predict the workload growth in the future that demands more processing power, we recommend that you create a logical partition with dedicated processors for the best logical partition performance. If you have a limited number of physical processors and want to create several logical partitions with flexibility in processor resources usage among all logical partitions, then a shared processor could be your choice.

1.3.12 Considerations in configuring virtual processors

The virtual processor is used to define a whole number of processors owned by a shared logical partition from fractional processors assigned to it. The logical partition operations will be faster and produce better results if the number of virtual processors assigned is very close to the number of processing units from fractional processors assigned to that logical partition.

The addition of some virtual processors may slightly improve the performance of the operating system because the operating system assumes that it has more processors to do more concurrent jobs. However, adding too many virtual processors to a shared processor logical partition will degrade the performance of each virtual processor. If you add too many virtual processors for a logical partition and the processing unit is not increased, the processing units for each new virtual processor will be smaller than for the previous virtual processor. The operating system cannot shift processing powers among processes very well.

An uncapped shared logical partition can use idle processing units in the shared processors pool. If the virtual processors for an uncapped shared logical partition is too low, it cannot use the maximum processing units in the shared processors pool. Setting too many virtual processors may impact logical partition performance as explained in the previous paragraph.

Determine carefully the proper number of virtual processors that will be configured for the shared logical partition for optimum logical partition operations.

1.4 Memory use in System i5 LPARs

Memory on System i5 based System i5 systems is used and managed very differently than in the past. In the following sections we discuss the new concepts of and changes to the manipulation and use of memory.

1.4.1 Memory granularity

In the past we have been able to assign or move memory from a partition at the single MB level. Memory in the System i5 machines cannot be manipulated that finely. Instead it must be assigned or moved in blocks.

These blocks are referred to as Logical Memory Blocks (LMBs) but are more technically described as Memory Regions. This is similar to the method by which the current Model 690 pSeries servers handle memory.

The LMBs are sized in multiples of 16 MB, and at initial announce, this will be the only size that can be used. Later code updates will allow for use of 16, 32, 64, 128, and 256 MB blocks. The designation of the memory region size you intend to use is a system wide parameter, so that everything on a given system will be using the same size blocks.

What this means is that whatever LMB size is specified for the system, that will be the smallest amount of memory that can be manipulated. All memory allocations will have to be in multiples of the LMB size, and memory can only be moved between partitions at the LMB size level.

So rather than a memory granularity of 1MB, the System i5 memory granularity will be between 16 MB and 256 MB. Again, this is determined by the system setting for the memory region size.

With the June 2004 level of code, you can view the LMB size from the HMC by displaying the Managed Server properties, selecting the memory tab, and you will see the current LMB setting. In a future code release, you will be able to change the LMB size through the ASM interface. There will be an option under Performance settings.

Actual memory used for paging and cache are not affected by or have anything to do with the memory region size.

1.4.2 Memory and LPAR

The assignment of memory to logical partitions in a System i5 V5R3 environment is handled differently than it has been in the past. Functionally, one of the main benefits for LPAR is that a partition can utilize all the memory assigned to it.

1.4.3 Hardware page table

The hardware page table (HPT) on systems prior to System i5 was contained within the memory assigned to the partition. This was treated as hidden or reserved memory. There is no longer the concept of reserved memory on System i5 hardware. The HPT now exists in memory that is controlled by pHyp or hypervisor memory. When a partition profile is defined, the amount of memory designated in the “max” memory parameter affects the size of that table. This is based on a specified ratio applied against the “max” parameter.

Even though there is no reserved memory, you should still take care when deciding on memory size and be reasonable about the maximum memory parameter. Because there is no reserved or “hidden” memory, a partition gets full use of its assigned memory. The Phyp memory is a part of the total physical memory on a system and it is not specifically for the HPT. There are other system factors that also affect the amount of memory in Phyp. However, we are only looking at LPAR memory requirements in this chapter.

1.4.4 Memory assignment at initial LPAR start

When you create a partition profile, you specify the min/current/max memory numbers. Click the **Recalculate** button and it shows you what is the minimum memory requirement for that partition. The minimum required memory is determined by the current/max numbers and the operating system type that will be run in that partition. For OS/400 partitions, this is either 128 MB or 256 MB depending on the max value. The more common one seems to be 256 MB. This is referred to as the Real Memory Offset (RMO).

When you start a partition for the very first time, what it does is to look for 256 MB (that minimum amount) of contiguous memory to be the RMO. This means that with a 16 MB memory region size, it will need to find 16 contiguous blocks of memory. This memory is owned by the partition and is entirely usable.

The rest of the partition's memory does not have to be contiguous but it is also assigned at the initial start up. Once the partition is associated with all of its physical memory, it then owns it. That ownership is persistent even if the partition gets powered down. No other partition may use that memory unless the profile information is overtly changed. When you restart a partition, it uses the same RMO and memory each time.

1.4.5 Memory allocation for the i5 OS logical partition

The logical partition is assigned a minimum, desired, and maximum memory as defined in the logical partition profile. The memory assigned to i5 OS logical partition is only for the logical partition. It does not include the memory reserved for Hardware Page Table (HPT). For example, the logical partition will receive full 1024 GB of memory as configured in the logical partition profile. Additional memory for HPT is allocated separately by the Power Hypervisor. In previous iSeries servers, the memory allocation for a logical partition includes memory allocated for HPT.

The virtual I/O resources also require some amount of Power Hypervisor memory, but this only requires a small amount of memory. It only changes in 16 MB increments on the current i5 systems.

To configure logical partition memory, determine the minimum, desired, and maximum memory for the logical partition. Desired memory is the amount of memory configured for the logical partition. This is the memory which the logical partition needs to have when the logical partition is powered on. If the managed system does not have the desired amount of memory but only has lesser memory, those uncommitted memory resources will be assigned to the logical partition when the logical partition is activated.

Minimum memory is the minimum amount of memory which needed by the logical partition to start. The logical partition will fail to start if the minimum memory is not available in the system. The current memory for the logical partition cannot be lower than the configured minimum memory. If you change the minimum memory setting for a powered on logical partition, it will take effect after the next logical partition restart.

Maximum memory is the total amount of memory that can be allocated to this logical partition. You can allocate additional memory to this logical partition using Dynamic Logical Partitioning (DLPAR), but it must not exceed the maximum memory value. Figure 1-23 shows an example of how to configure the memory in the logical partition profile.

Figure 1-23 shows a screenshot of the 'Create Logical Partition Profile - Memory' dialog box. The dialog box contains the following information:

- Installed memory (MB): 11534336
- Current memory available for partition usage (MB): 10485504
- Minimum memory: 2 GB, 0 MB
- Desired memory: 4 GB, 512 MB
- Maximum memory: 6 GB, 512 MB

The dialog box also includes a 'Help' button and navigation buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

Figure 1-23 Configure memory for logical partition

When the partition is activated, the configured memory is assigned to them and the Hardware Page Table (HPT) is created by Power Hypervisor.

Memory allocation must be calculated precisely, in order to avoid lack of memory, which could have an impact to the logical partition performance. If the memory allocation for all logical partitions is not calculated precisely, the last powered on logical partition will receive less memory than the configured amount of memory. If the minimum memory requirement for the last powered on logical partition is not met, the logical partition could not be started.

From the managed system properties, the memory status can be viewed. Figure 1-24 shows the memory information for the current managed system. It provides much information, for example, Installed Memory, which displays the total number of memory units that are installed on the managed system. Current memory available for partition usage is an amount of memory that is available for use by all the logical partitions configured in the managed system. This value will vary depending on the number of logical partitions configured and the memory installed on the managed system. Partition memory usage displays current memory in use by all active logical partitions.

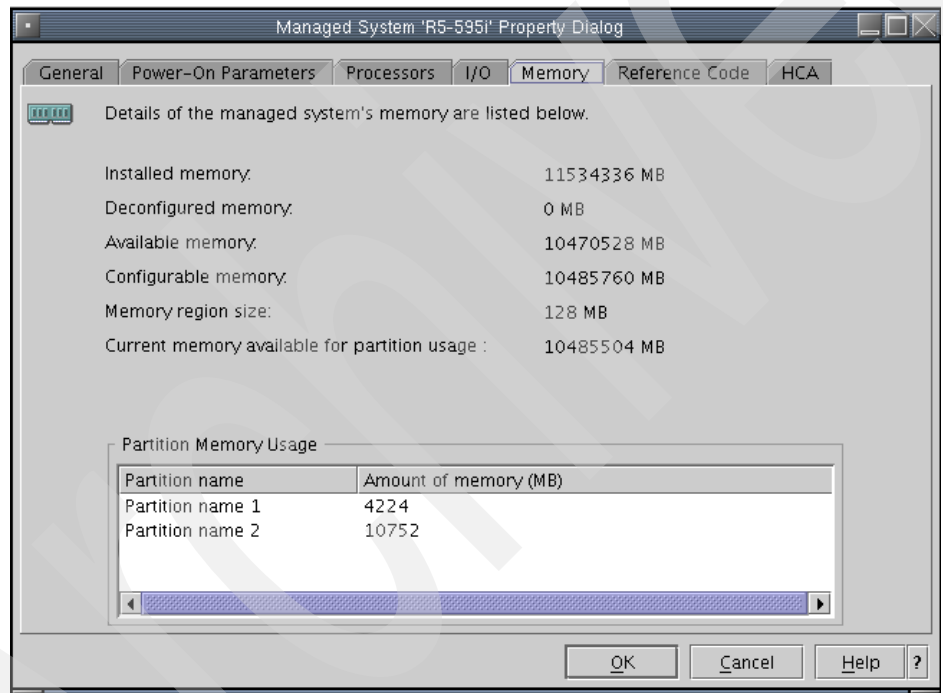


Figure 1-24 Managed system memory information

For example, partition name 2 has the memory configuration as shown in Figure 1-25. When the logical partition was powered on, it received all desired memory, which is $10240 \text{ MB} + 512 \text{ MB} = 10752 \text{ MB}$. This value can be found in the managed system memory information shown in Figure 1-25.

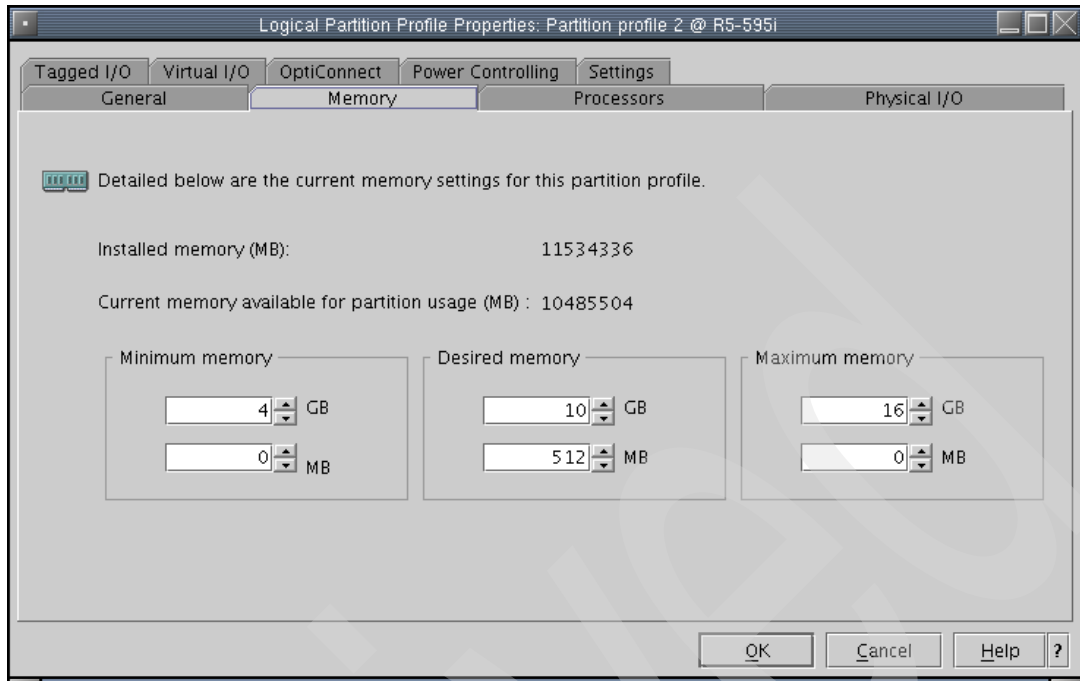


Figure 1-25 Logical partition memory configuration

The Hardware Page Table for this logical partition is calculated by dividing the total memory available for the logical partition by 64. The result is then rounded to the next whole number. For the configured memory of the logical partition in Figure 1-25, the HPT size is $10752 \text{ MB} / 64 = 168 \text{ MB}$. The Power Hypervisor will then allocate 256 MB memory for this logical partition HPT.

Another memory allocation in System i5 system is for Power Hypervisor memory. The Power Hypervisor requires some amount of memory to perform logical partitioning functions in System i5 system. There are several factors affecting the Power Hypervisor memory requirements:

1. Logical partition memory configuration
2. Maximum amount of memory assigned to the logical partition
3. Virtual I/O used by logical partition
4. Total number of logical partition in the System i5 system

“Processor use in System i5 LPARs” on page 8 discusses the Processors, which are grouped into nodes in System i5 system. Each node is interconnected to share its memory. Memory access between nodes is slower than memory access from local node. It means that memory access by processors from within the same node is faster than memory access by processors from a different node.

The Power Hypervisor will attempt to group processors and memory together for better logical partition performance, especially for dedicated partitions. The Power Hypervisor tries to avoid too many remote memory access operations, which may slow down the job’s execution, because it takes a longer time to access memory not in the same node with processors. The Power4 firmware allocates memory equally from all nodes, which leads to more remote memory access than System i5 Hypervisor does.

Memory can be moved between logical partitions dynamically. However, this may cause configurations to become less optimal, because the memory moved/removed will probably be spread across all nodes. Determining an optimum memory configuration is very important to achieve logical partition best performance.

Archived

i5/OS consoles under System i5

In this chapter we provide details of console options and rules on the Power 5 platform as it relates to iSeries for stand-alone (non-partitioned) systems and for LPAR systems. For stand-alone systems, the console IOA placement requirements are very specific and must be followed. For LPAR systems, the placement rules do not apply, but there are some concerns.

The following topics are covered in this chapter:

- ▶ Console for stand-alone systems
- ▶ Console for partitioned (LPAR) systems
- ▶ Console IOA placement rules for stand-alone systems
- ▶ Console IOA placement rules for stand-alone systems, including IXS considerations
- ▶ Connecting to a 5250 console remotely

2.1 Console history

There are now a few types of consoles that can be used on the iSeries. Originally, in the System 3X world, the Twinax console was the one and only kind of console that could be used. The interface was a 5250 green screen. When the AS/400 was announced, Twinax was still the only console. Async Console was added, which used a 5250 session in Client Access via a serial port on a PC. This is no longer supported.

Operations Console followed and consists of console GUI software on a PC, again via a serial port. This is referred to as Direct Operations Console. With OS/400 V5R1, the LAN Console was introduced. This is the same as the Operations Console, except that it is LAN attached. It can be a console for more than one system or partition at a time by running multiple console sessions on the same PC. Each system or partition must have a LAN card dedicated to the console function.

With i5 systems, each system or partition still needs to have its own console, and there are now four types of consoles supported: 1=Twinax, 2=Operations Console Direct cable, 3=Operations Console LAN, and 4=Hardware Management Console (HMC).

The newest, of course, is the HMC. On some System i5 systems, attaching an HMC is optional, as they can still use the other three kinds of console. Some of the System i5 systems require the use of the HMC even though the other console types can still be used. All System i5 partitioned systems must have an HMC because all LPAR functions on System i5 systems are moved to the HMC and are no longer in DST/SST. This is also true for Capacity Upgrade on Demand (CUoD) functions. In any of the above cases, this does not mean that the HMC has to be the selected console type, but all systems attaching an HMC can use the HMC as the system console if so desired. The HMC does not use any physical system adapters, as it attaches to special ports in the Service Processor (SP).

2.2 Console for stand-alone systems

Twinax, Direct Operations Console, and LAN Operations Console have no functional or operational difference from OS/400 V5R2 to V5R3. However, some newer generation PCs may not have a serial port to use with Direct Operations Console. In that case, a USB to Serial cable can be used. Also, the Direct Operations Console cable on a stand-alone system will connect to the ECS card.

Operations Console LAN configurations have always had the LAN adapter serial number stored as part of the configuration. If the serial number of the configured LAN card cannot be found during IPL time, then an additional 30 minute wait occurs to see if the adapter reports in. If not, then another adapter can be used. This only applies if the load source is migrated or the load source has a prior LAN Operations Console configuration.

The 30 minute wait can be overcome by clearing the LAN configuration on the load source just before you migrate it. Clear and store the LAN configuration but do not deactivate it. Then power off the system and migrate the load source.

Function 65 and 21 (disable WAN and DST), when used together within 30 seconds to change console modes, will cause the current 5250 session to be closed and a new 5250 session to be opened (new 5250 session ID). Function 65/21 will only change console modes if the resources are tagged before you try to switch them.

For example, if you have an IOP tagged for console and the IOP has a WAN and LAN card in the IOP, you can change modes between direct cable or LAN connected Operations Console.

If you also tag an IOP for ALTERNATE console, it must have a Twinax card. Then you can switch between Direct, LAN, or Twinax. The ALTERNATE console cannot install the operating system or LPPs and may have limited DST functions.

See Figure 2-1, Figure 2-2, and Figure 2-3 for an illustration of these situations.




Figure 2-1 HMC as the Console - no system adapters are required - virtual I/O is used

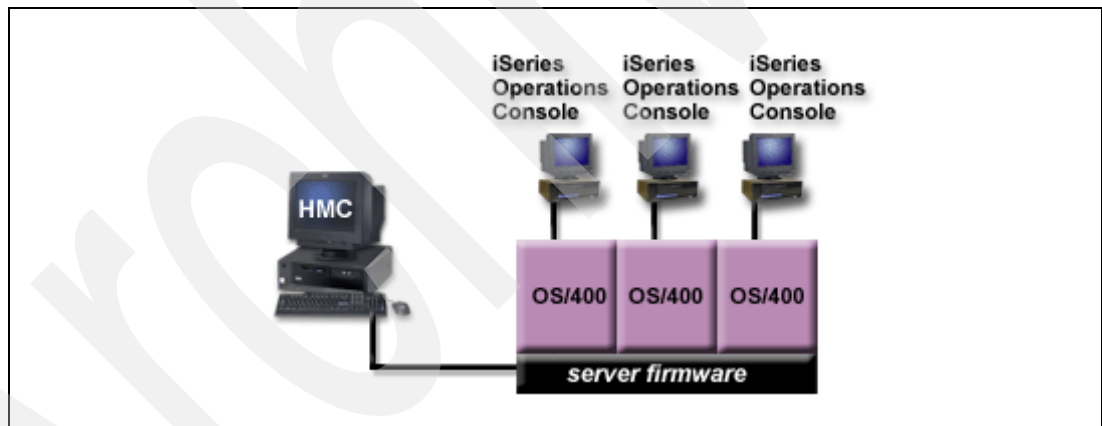


Figure 2-2 Operations Console LAN or Direct cable connections - system adapters are used - V.24 (ASYNC) or LAN

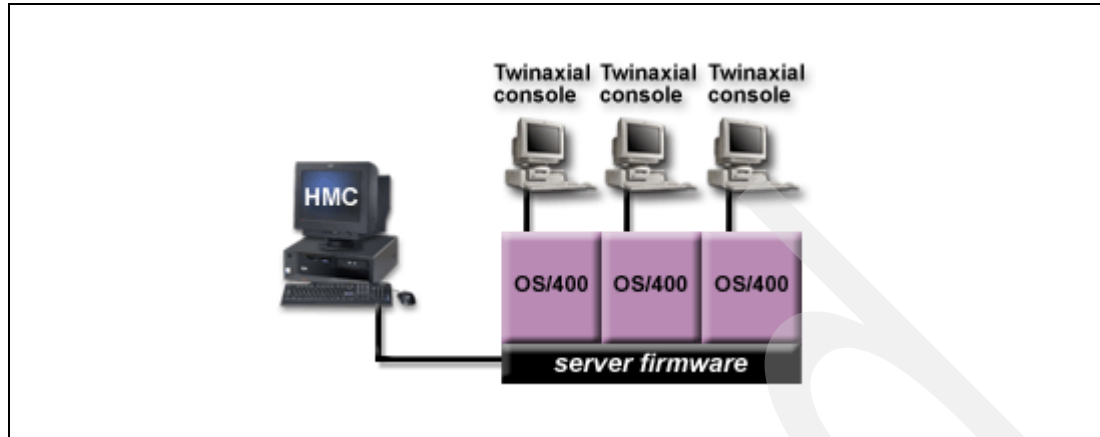


Figure 2-3 Twinax system adapter is needed

2.3 Console for partitioned (LPAR) systems

On System i5 systems, the LPAR structure is not as we have known it on iSeries. Most of the basic concepts are the same, but there are significant changes. System i5 systems have no Primary (controlling) or any Secondary partitions. All partitions are “equal” and independent in terms of their relationship to each other. Linux partitions can still use virtual resources from an OS/400 partition or be entirely independent. AIX V5R2 or later will run in a self contained partition.

The Console mode selection in DST can be overridden by the tagged resource (via HMC) when the partition is built. When tagging the console resource for the partition, you can tag an IOP or an IOA. If you tag an IOA and that resource does not match the console type selected, then the tagged IOA will be used.

For example, if you tagged a WAN IOA and the value for the console mode was a 3 for LAN, the system will try to use Operations Console Direct cable. This is because the TAG for the IOA will override the console mode. If an IOP was tagged and the IOP had a WAN and an Ethernet card, then the console would try to come up on the LAN adapter.

The HMC console is the *only* console that can take the console function away from another console type, and will take it over even if the other console is in a session and doing work.

Primary partitions that are being migrated will require manual intervention to set the console tag. Furthermore, any secondary partition on the source system that uses a console resource in the main system tower will require manual intervention to tag the console resource. Any secondary partition that has a console resource outside of the main system tower should migrate without a problem, provided that the resource is on the system.

If you use HMC as the system console, then you do not need to tag anything, as it is done for you by the HMC configuration wizard.

2.4 Console IOA placement rules for stand-alone systems

In the following sections we list the various rules.

Note: These rules apply to LAN Console and Twinax Console in a stand-alone system. If using the Direct Operations Console, the ECS card/slot will be used rather than any of the other slots.

2.4.1 Model 520/550 CEC

The 520/550 will first look for the console in slot 5. If an appropriate IOA is not there, it will look in slot 2 (which requires a second IOP).

Model 520
SLOT1 - 2nd IOP
SLOT 2 Console if 2nd IOP
SLOT 3 ECS
SLOT 4
SLOT 5 Console
SLOT 6 IOP

2.4.2 Model 570 CEC

The 570 will first look for the console in slot 4 (which requires a second IOP). If an appropriate IOA is not there, it will look in slot 6.

If a GX Adapter (for HSL) is installed, it will block slot 6. Since a Twinax adapter would have to use the same slot as LAN, it cannot be a backup for the LAN in this particular scenario. Direct (Async) will have to be the backup.

Model L4
SLOT 1- IOP
SLOT 2 ECS
SLOT 3 - 2nd IOP
SLOT 4 Console
SLOT 5
SLOT 6 Console or GX adapter (HSL)

2.5 Console IOA placement rules for stand-alone systems, including IXS considerations

Pre-System i5 IXSs that will be migrated from another system must be placed in an expansion tower. They are not supported in the system unit. The new IXS available during 2004 can be placed in the system unit and will use the slots listed in the following tables.

2.5.1 Model 520/550 CEC

The 520/550 will first look for the console in slot 5. If an appropriate IOA is not there, it will look in slot 2 (which requires a second IOP). If an IXS is used, the console must put in slot 2 because the IXS must be in slot 4 and blocks slot 5.

Model 520
SLOT 1 - 2nd IOP
SLOT 2 Console if 2nd IOP
SLOT 3 ECS
SLOT 4 IXS
SLOT 5 Blocked
SLOT 6 - IOP

2.5.2 Model 570 CEC

The 570 will first look for the console in slot 4 (which requires a second IOP). If an appropriate IOA is not there, it will look in slot 6.

If a GX Adapter (for HSL) is installed, it will block slot 6.

An IXS can be put on the second IOP (slot 4) in addition to the console (slot 6).

If you have an IXS and a GX adapter, then you must use Operations Console Direct or HMC as the console.

Model 570
SLOT 1 - IOP
SLOT 2 ECS
SLOT 3 - 2nd IOP
SLOT 4 IXS
SLOT 5 Blocked
SLOT 6 Console or GX adapter

2.6 Connecting to a 5250 console remotely

This section covers connecting to a 5250 console remotely. The remote support for HMC 5250 can use the same SSL configuration as the System Manager Security on the HMC. For more information about configuring System Manager Security, see “System Manager Security” on page 41.

2.6.1 5250 console remotely configured

To connect to a 5250 console remotely, read the following requirements:

- ▶ You must use a 5250 emulator, such as the IBM iSeries Access PC5250 emulator, iSeries Access for Linux emulator, or IBM Personal Communications 5250 emulator on Microsoft® Windows®.
- ▶ If you are using the IBM iSeries Access PC5250 emulator, the required level is Version 5 Release 3 PTF SI13587 or later. To view your service pack level, select **Start** → **Programs** → **IBM iSeries Access for Windows** → **Service** → **Check Service Level**.
- ▶ If you are using the IBM Personal Communications 5250 emulator, Version 5.7 or later is required.
- ▶ If you are using the iSeries Access for Linux emulator, Version 1.4 or later is required.

To connect remotely, do the following steps.

Configure HMC support for your Windows or Linux product

Open the 5250 emulator application on your PC to configure your Windows product, then do the following steps:

1. Create a new session. Figure 2-4 provides an example of how to configure a new session.




Figure 2-4 Create a new session with PC5250

2. Select **Link Parameter**. When the window in Figure 2-5 prompts you, type the HMC host name or IP address in the Host Name field and indicate the port number. Type 2300 into the port number field if you are using non-SSL or 2301 if you are using SSL. Then select **OK** to finish the configuration.

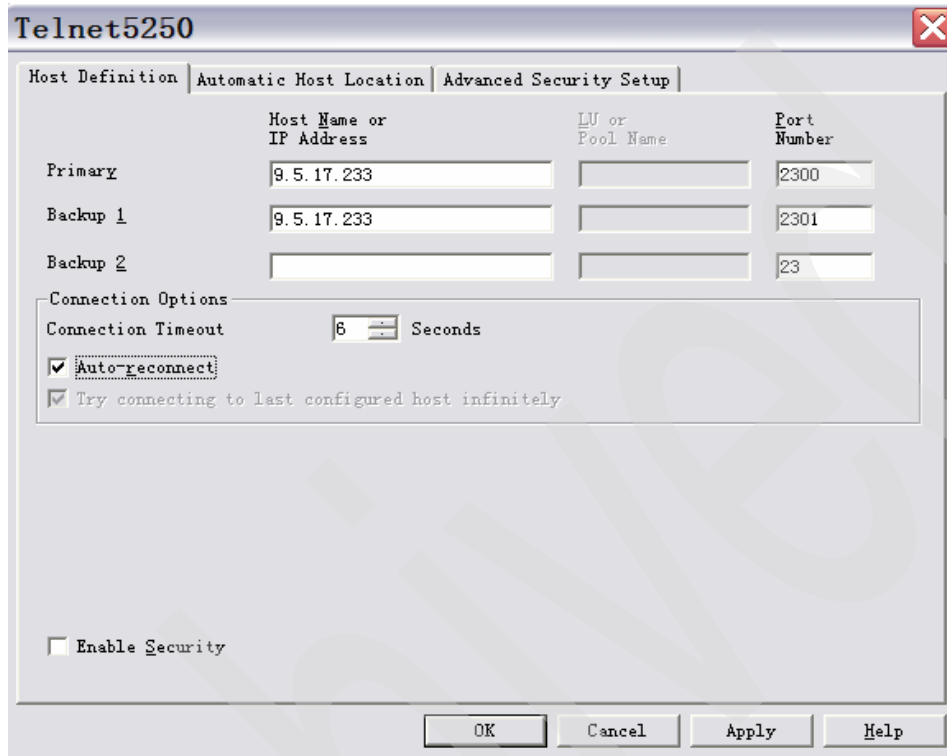


Figure 2-5 Configure 5250 IP address for remote console

3. When the window in Figure 2-6 prompts you, select the correct language and press Enter.

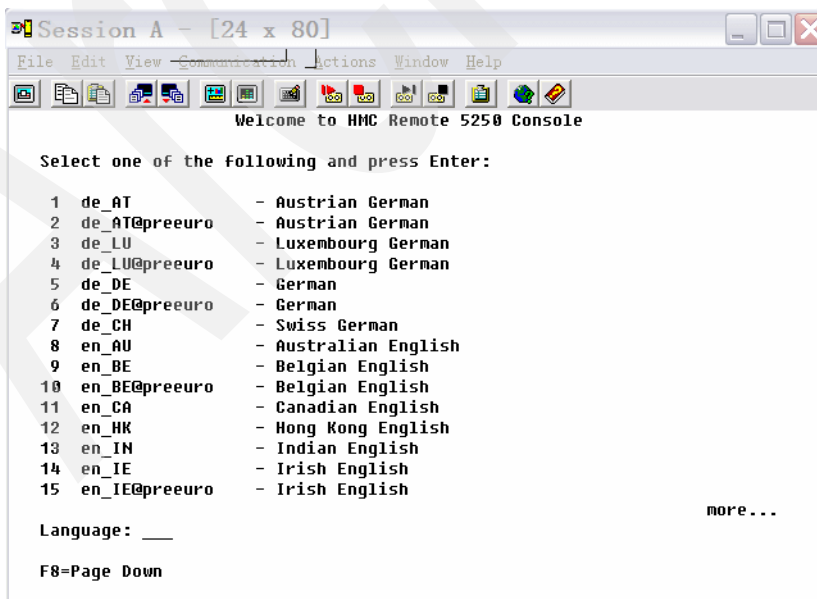


Figure 2-6 Remote 5250 console - select language

4. When the window in Figure 2-7 prompts you, type the correct HMC user ID and password, then press Enter.

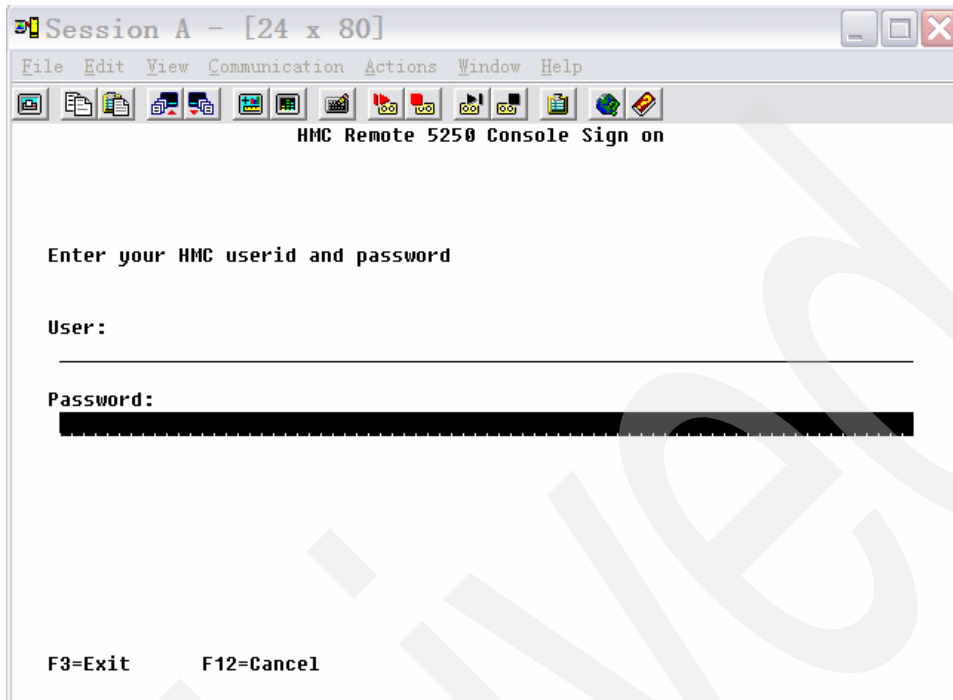


Figure 2-7 Remote 5250 console - HMC user ID and password

5. When the window in Figure 2-8 prompts you, select the management server that you want to access.

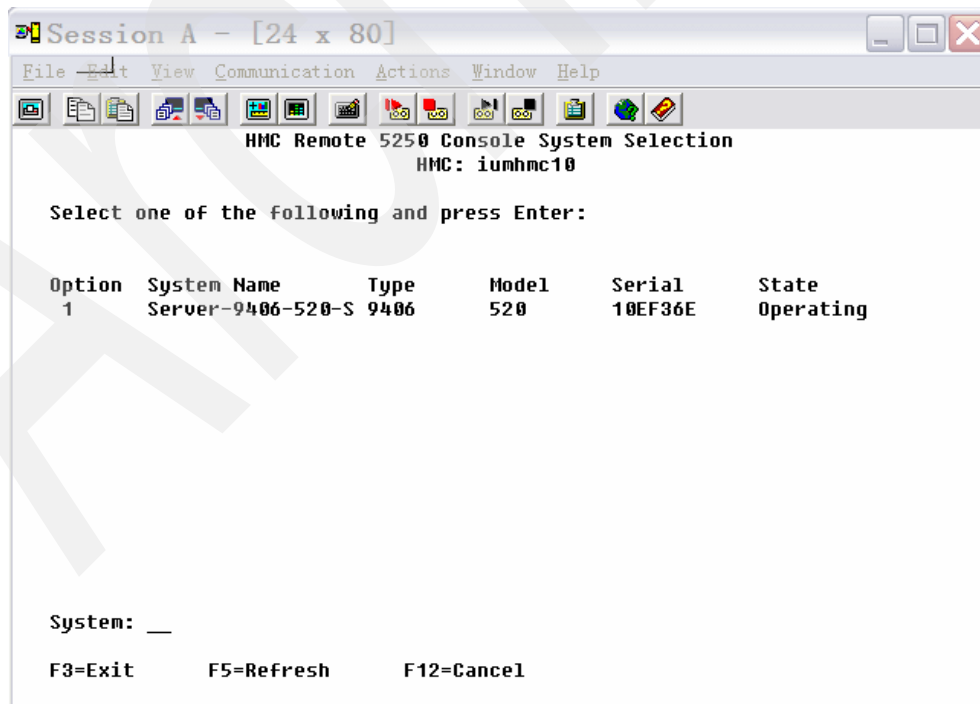


Figure 2-8 Remote 5250 console - select management server

6. When the window in Figure 2-9 prompts you, select connect modes and press Enter.

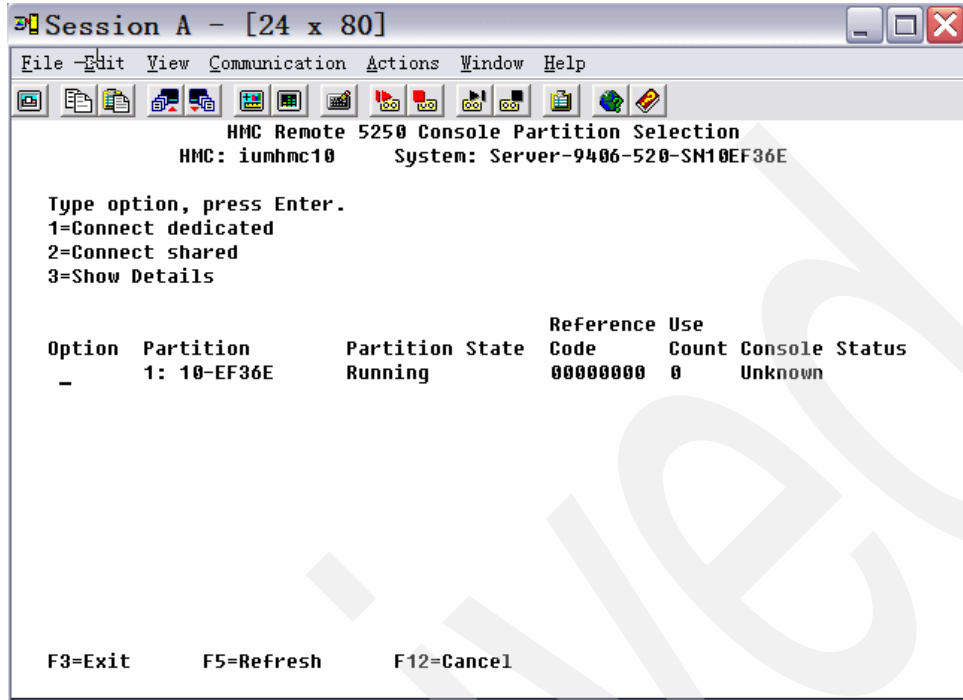


Figure 2-9 Remote 5250 console- connect modes

7. When the window in Figure 2-10 prompts you, type the correct i5/OS user ID and password to manage this system.

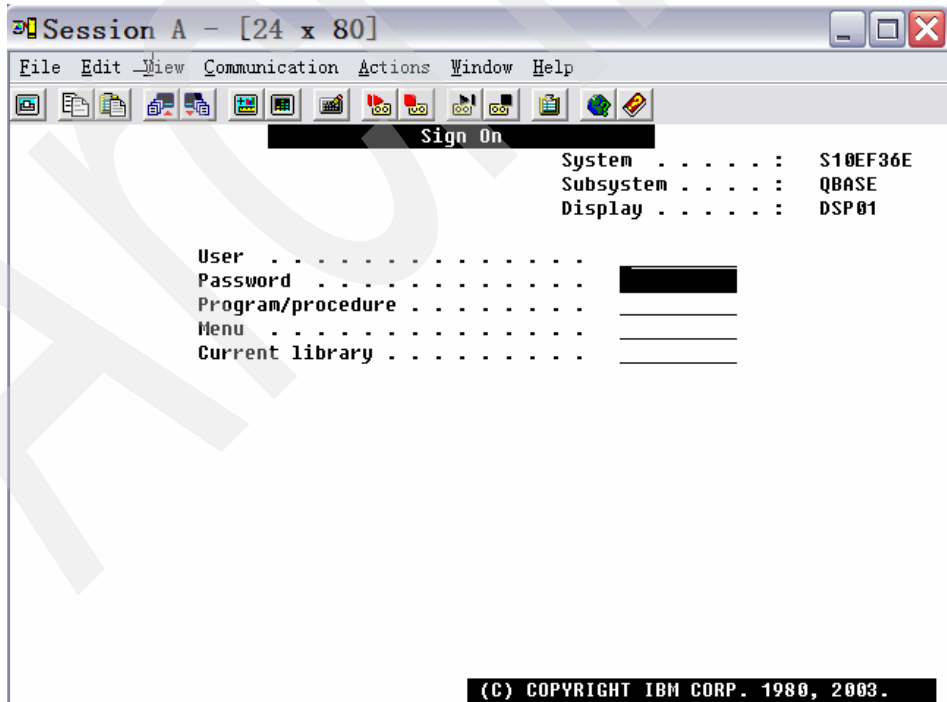


Figure 2-10 Remote 5250 console-sign on screen

Configure your Linux product

To configure your Linux product, do the following steps:

1. Create a new session by using the setup5250 configuration program.
2. In the 5250 Emulator Connection window, type the HMC TCP system name or IP address in the **AS/400 Host Name field**.
3. Select **Advanced 5250 Connection....** The Advanced 5250 Emulator Connection window is displayed.
4. Type **2300** into the **Telnet Port number field**.
5. Type **Emulator User ID** and **Emulator Password** fields.
6. Open the 5250 emulator and select the language you want to see on the interface.
7. Log in using your HMC user ID and password.
8. Select the managed system that contains the logical partition to which you want to connect.
9. Select the logical partition to which you want to connect.
10. If you select **Shared**, you must enter a session key that can be used by another user with another 5250 emulator. If you are the first user, the session key is entered twice to verify accuracy.

Note: If you select **Dedicated**, other users are not able to share the session with you. If you select **Details**, you can view which users are connected to this partition.

System Manager Security

System Manager Security ensures that the HMC can operate securely in client/server mode. Servers and clients communicate over the Secure Sockets Layer (SSL) protocol, which provides server authentication, data encryption, and data integrity. Each System Manager server has its own private key and a certificate of its public key signed by a certificate authority (CA) that is trusted by the System Manager clients. The private key and the server certificate are stored in the server's private key ring file. Each client must have a public key that contains the certificate of the trusted CA.

A Certificate Authority (CA) is a trusted central administrative entity (a local HMC in this situation) that can issue digital certificates to clients and servers (HMC4 in Figure 2-11). The trust in the CA is the foundation of trust in the certificate as a valid credential. A CA uses its private key to create a digital signature on the certificate that it issues to validate the certificate's origin. Others, such as System Manager clients, can use the CA certificate's public key to verify the authenticity of the certificates that the CA issues and signs.

Every digital certificate has a pair of associated cryptographic keys. This pair of keys consists of a public key and a private key. A public key is part of the owner's digital certificate and is available for anyone to use. A private key, however, is protected by and available only to the owner of the key. This limited access ensures that communications that use the key are kept secure.

The owner of a certificate can use these keys to take advantage of the cryptographic security features that the keys provide. For example, the certificate owner can use a certificate's private key to "sign" and encrypt data sent between clients and servers, such as messages, documents, and code objects. The recipient of the signed object can then use the public key contained in the signer's certificate to decrypt the signature. Such digital signatures ensure the reliability of an object's origin and provide a means of checking the integrity of the object.

A server is an HMC you want to access remotely. In Figure 2-11, HMCs 1, 3, and 4 are servers. A client is a system from which you want to access other HMCs remotely. In Figure 2-11, Web-based System Manager Remote Clients A, B, and C, and HMCs 1, 2, and 5 are clients. As shown in Figure 2-11, you can configure multiple servers and clients in your private and open networks.

An HMC can be in multiple roles simultaneously. For example, an HMC can be a client and a server like HMC1 in Figure 2-11. An HMC can also be a CA, server, and client at the same time.




Figure 2-11 System manager security servers and clients

Each server must have a unique private key and a certificate of its public key signed by a CA that is trusted by the clients. Each client must have a copy of the CA's public key.

The following list is an overview of tasks involved in installing and securing the remote client:

1. Configure one HMC as a Certificate Authority (CA).
2. Use this HMC to generate private keys for the servers.
3. Install the private keys on the servers.
4. Configure the servers as secure System Manager servers.
5. Distribute the CA's public key to the servers or clients.

Note: Tasks 3 and 5 are completed by copying the keys to diskette and installing them on the servers or clients.

Archived

HMC overview and planning

This chapter provides an overview of the application functions of the Hardware Management Console (HMC). We also discuss some of the planning needed to install an HMC. The HMC can connect to one or two systems to perform various functions, including these:

- ▶ Creating and maintaining a multi-partitioned environment
- ▶ Running a 5250 console session for each i5/OS partition
- ▶ Detecting, reporting, and storing changes in hardware conditions
- ▶ Powering managed systems on and off
- ▶ Acting as a service focal point for service representatives to determine an appropriate service strategy and enable the Service Agent to call home to IBM
- ▶ Activating additional resources on demand

This chapter is divided into the following sections:

- ▶ HMC concepts, information, and planning
- ▶ Server and partition
- ▶ System manager security
- ▶ Licensed Internal Code maintenance
- ▶ HMC management
- ▶ Service applications

3.1 HMC concepts and initial setup

The main functions of the HMC are to perform logical partitioning functions, service functions, and various system management functions. The partitioning functions and some of the servicing functions, which were previously in the iSeries service tools, will now be performed by functions in the HMC for partitioned systems. The HMC can be used to manage from one to two partitioned systems. Other highlights of the HMC are as follows:

- ▶ The HMC connects to the managed system or systems via an ethernet LAN connection.
- ▶ It runs a Java™ based application running in an embedded OS.
- ▶ The user is only able to access the management applications via a GUI interface.
- ▶ The HMC can be run remotely via a Web-based System Management Remote Client running on a Windows based PC or a Linux based PC.
- ▶ A virtual console terminal can be configured to run on the HMC for each partition, thus reducing the need for extra hardware in each partition. One of these console types is 5250.
- ▶ We recommend that you configure a second HMC for redundancy. This second HMC cannot actively manage the same server at this release.

3.1.1 HMC and Service Processor levels, and information sources

When you are preparing to install or maintain your HMC and/or the services processor on the System i5, you should review the following chart for the corresponding system firmware levels (Figure 3-1).

Some times called	HMC release level	System Firmware			
		SF release level	Base level	Service Packs (SP)	Date base first available
GA7	V5R2	SF240	201	2xx	2006 Feb
GA6*	V5R1	SF235	160	180, 185	2005 Oct
GA5	V4R5	SF230	120	126, 143, 145, 150	2005 Apr
GA4	V4R5	SF225	96		2005 Feb
GA3	V4R3	SF222	71	75, 81	2004 Nov
GA2	V4R2	SF220	43	45, 48, 49, 51	2004 Aug
GA1	V4R1	SF210	24	26, 28, 29	2004 Jun

* i5/OS V5R4 minimum firmware support level

Figure 3-1 HMC and System Firmware levels

Here we list some Web sites containing useful information and the most recent levels of the HMC and System firmware:

- ▶ IBM iSeries™ Hardware Management Console Frequently Asked Questions:
<http://www-1.ibm.com/support/docview.wss?uid=nas148859a914db132a586256f42006003a7&rs=110>
- ▶ Hardware Management Console:
<http://techsupport.services.ibm.com/server/hmc>

- ▶ iSeries Recommended Fixes - Server Firmware: Update Policy Set to Operating System:
http://www-912.ibm.com/s_dir/s1kbase.nsf/c32447c09fb9a1f186256a6c00504227/604992740f846a4986256fd3006029b5?OpenDocument
- ▶ iSeries Recommended Fixes - Server Firmware: Update Policy Set to HMC:
http://www-912.ibm.com/s_dir/s1kbase.nsf/ibmscdirect/E58D7BBF0EAC9A2786256EAD005F54D8

3.1.2 Types of HMC

The HMC runs as an embedded OS on an Intel® based workstation that can be desktop or rack mounted. The embedded OS and applications take over the whole PC, and **no other** applications are allowed to be loaded.

There are some models of HMC that were available for pSeries. These models can be upgraded to run the HMC code and manage System i5 systems. But they cannot manage both Power4 and System i5. The upgrade is simply a scratch install of the System i5 HMC code level.

Whether you opt for a desktop or rack mounted version is personal choice. Customers with space in their rack mounted systems would probably opt for the rack mounted version with the slide-away keyboard and screen:

- ▶ 7310-C03/04 is a desktop HMC.
- ▶ 7310-CR2/3 is rack mounted HMC.

The desktop and rack mounted HMC have a number of standard ports installed USB, Parallel, and so on.

Desktop HMC

The supported desktop models are the 7310-C03/04, and older versions 7315-C01, C02, C03, which can be migrated to System i5 HMC code level. Figure 3-2 shows a picture of the 7310-C03.



Figure 3-2 IBM 7310-C03 Hardware Management Console

On the desktop you can connect a keyboard and mouse to either the standard keyboard, mouse PS/2 style connectors, or to the front side USB ports. You should not use the rear side USB ports. You cannot connect **any other** devices to the HMC. Printers are not supported off the parallel port.

The desktop HMC can use a number of IBM displays as shown in the e-config. The desktop HMC does not have to use the monitor shown in the above figure. There is no ability to add device drivers to the embedded OS. So you should test any proposed OEM display before running in production.

Rack mounted HMC

The supported rack mounted models are the 7310-CR2/3, the older version 7315-CR2 which can be migrated to System i5 HMC code. Figure 3-3 shows a picture of the 7310-CR3.

This picture shows the HMC 7310-CR2 system unit as a standard 1 U unit below the screen/keyboard mounted in a standard 1 U pull-out tray.

The rack mounted HMC does not have standard PS/2 keyboard and mouse ports. You must order the breakout cable shown in Figure 3-6 on page 51. This cable plugs into the large connector to the left of the RJ-45 connectors. This breakout cable terminates in a pair of PS/2 style female keyboard and mouse connectors and a display connector. You can then plug in any standard keyboard or mouse and display.

In Figure 3-3 we have the draw mounted 7316 folding display and slim keyboard. This is a great choice for a dark machine room, where space is restricted. The displays shown in e-config do not include the 7316.




Figure 3-3 7310-CR2 rack mounted HMC with pull-out tray

Note: As shown in the graphic above, the 7310-CR2 is the 1 U bottom unit. The pull-out keyboard and monitor is a **separately orderable feature** IBM 7316-TF2 1U Flat Panel Monitor Console Kit.

3.1.3 Initial setup of the HMC

Depending on the system that is being installed, some of the HMC setup is Customer Set Up (CSU). The IBM CSR installing the system performs some of the HMC setup steps, and the customer performs the remainder of the steps. Remember that disk protection configuration and LPAR configuration are the responsibility of the customer.

When a new system or an upgrade is delivered, it will have a two-page document called “Start here for hardware”. This document is the starting point for setting up the system and the HMC. It guides you to use a CD called the “eServer™ Hardware Information Center”, which will be included with the new system. Alternatively, this same information can be accessed directly from the Web. This is the preferred method; it contains the most recent information:

http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/index.htm

The eServer Hardware Information Center guides the user through cabling up the HMC and then through a checklist to gather information needed to configure the HMC. The information needed includes:

- Network settings for the HMC. These may include the following items:
 - HMC hostname
 - Domain name
 - Enabling DHCP server support on the HMC
 - Enabling DHCP client support on the HMC
 - A gateway IP address
 - Firewall settings
 - Enabling DNS
- Customer contact information, as the HMC will be a service focal point
- Type of connection to connect to the service provider

The Guided Setup wizard guides you through using this information to configure the HMC in very logical steps. To access the setup wizard and a link to the information center, power on the cabled up HMC and sign on with user ID *hscroot* and a password of *abc123*. When the HMC has finished booting up, you will see the Management Environment interface as shown in Figure 3-4.

Click **Information and Setup** → **Launch the eServer Information** to access the Information Center.

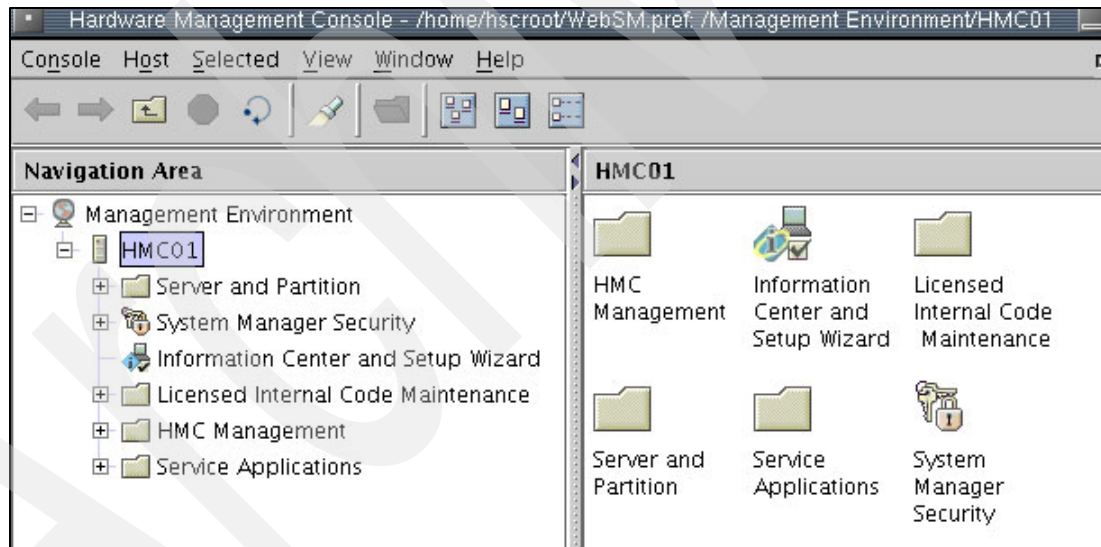


Figure 3-4 The HMC user interface

This System Management environment is divided into two sections. The left hand pane is called the Navigation Area and the right hand pane is called the Contents Area. In the example shown, the name given to the HMC is itschmc1. In the following sections, we give an overview of functions performed under each of the icon headings.

A more detailed description of the guided setup for the HMC is given in Chapter 4, “HMC Guided Setup” on page 105.

3.2 Installing the HMC

In this section we discuss the physical setup of the HMC.

Attention: When installing a new System i5 system, do not power on the system before connecting it to an HMC. The Server processor (SP) on a System i5 system is a DHCP client and will search for a DHCP server to obtain its IP address. If no DHCP server can be found, then the SP will assign a default IP address. If this occurs, you will have to use ASM to manually change the IP setting of the SP.

Both the desktop and rack mounted HMC models come with the embedded OS, HMC code, browse, and eServer Hardware Information Center pre-loaded. Therefore, the HMC can be unpacked and powered up immediately after the networking components are connected.

With a desktop HMC, the additional ethernet LAN card will need to be installed. There is one embedded ethernet adapter, so with this additional LAN card, we have the standard two HMC ports. You cannot use the rear side USB ports. Refer to Figure 3-5.




Figure 3-5 Rear view of desktop HMC ports

Figure 3-6 shows part of the rear of the rack mounted HMC system unit. The HMC ports 1 and 2 are located in the center of the image (one ethernet cable is connected). You can also see the breakout cable, keyboard, display, and mouse.




Figure 3-6 Rear view of rack mounted HMC ports

The vertical USB port to the right of the HMC ethernet ports is not available for use.

The HMC 1 and 2 ports on the rear of the HMC system unit connect to the Service Processor HMC ports 1 and 2. But only **one** cable should be connected at one time. The service process would have major problems if both cables were to be connected between the ports. See Figure 3-7.

Notice that on the rear of the 570, there is a pair of SPCN cables. These are connected to external towers for power control. SPCN cables are connected in a loop on System i5 systems. In previous iSeries and AS/400 servers, the SPCN cables were always daisy chained.




Figure 3-7 Rear of Model 570

Notice the large dark connector on the lower right hand side of the image (below the right-hand fan). This is the system interconnect port for connecting other 570 nodes.

Figure 3-8 shows a closer rear view of a model 570 with the cables identified.




Figure 3-8 Ports on rear of i570

3.3 HMC networking options

When you first install your HMC, you will have a number of networking options available to you. You will have to decide which types of network are best for your environment. The types of networking available are:

- ▶ Private direct
- ▶ Private indirect
- ▶ Private and open

Note: The following examples of HMC networking options are only for single SP on i5. If you plan to connect your HMC to redundant SP(570/595), please see the reference on redundant SP connection topics, in Chapter 8, “HMC duplication and redundancy” on page 245.

3.3.1 Private direct networking

Private direct networking is an ethernet connect between port HMC 1 on the HMC and port HMC 1 on the System i5 system. When you are making this connection, you do not need an ethernet cross-over; a standard Ethernet cable will work. The network is provided by the HMC, DHCP serving, DNS, and Firewall. These components can be established by the tasks included in the Guided Setup wizard. See Figure 3-9.




Figure 3-9 Private direct networking

If you need to make changes to the network, these can be done manually within the HMC interface.

3.3.2 Private indirect networking

A private indirect network shown in Figure 3-10 is effectively the same as private network, but the signals pass through one or many hubs/switches. In Figure 3-10 we show two servers connected to a hub and two HMCs connected to the same hub/switch.




Figure 3-10 HMC with indirect network attach

One HMC is connected to IBM Service and Support. This HMC is actively managing both the servers. The second HMC is connected but redundant, in that it is not actively managing either of the servers, unless you are at GA2 level of HMC/SP.

We do not have any recommendation for the type of hub or switch that can be used in this arrangement. We would suggest that the customer connects the HMC directly to start, and then installs the network devices once they are confident with the HMC operation and management.

3.3.3 Private and open networking

As there are two network connections on an HMC, you can connect it to both a private network and a public network. The HMC shown in Figure 3-11 is the DHCP server to the indirect connected servers.

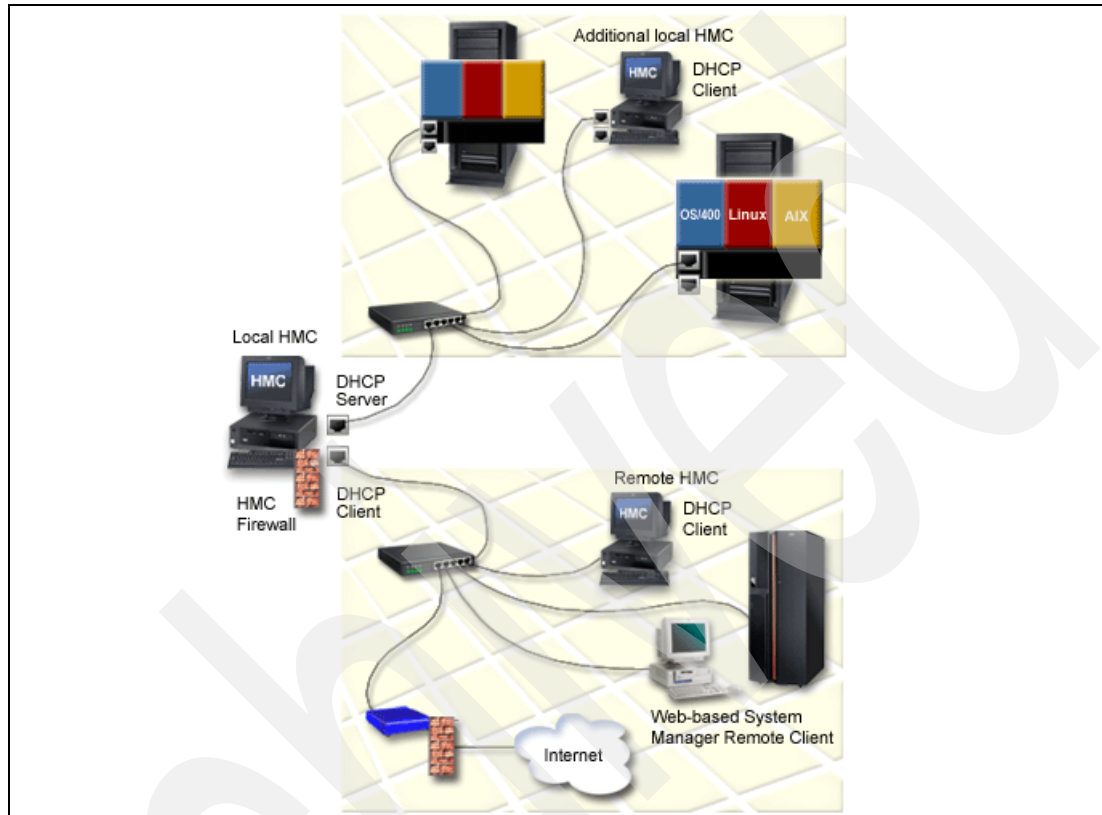


Figure 3-11 HMC connect to both private and public networks

There is also a connection to the local public network as a DHCP client. In this network we have Web-based System Management Remote Clients (WSMRC) and a remote HMC. This arrangement is probably the style most large customers would choose. The “main” HMC is located in the dark machine room. The remote HMC and WSMRC could be located in the bridge or IT department for operational use.

3.3.4 Open networking

One could choose to install both the HMC and the Service process ports onto a public network. This would mean that the network connection from the HMC to the Service Processor is part of the normal public network traffic and subject to these constraints. We would not recommend this type of network connection, and there is no obvious reason to install your HMC this way.

3.3.5 Working with HMC in open networks

To configure an HMC in an open network, you must provide it with the public network information, Domain name, DCHP server, and DNS server. You must also provide the Service Processor with a static IP address in the network. You can work with the Service Processor IP addresses from the Advanced System Manager interface.

3.4 Initial tour of the desktop

You will now be viewing the HMC desktop as shown in Figure 3-12.

The desktop comprises the following components:

Fluxbox desktop. This is a standard Linux desktop supplied with the embedded OS. There is a task bar at the bottom of the Fluxbox desktop.

HMC Window. This panel displays the HMC Management Environment. As shown in Figure 3-12, a managed server has been discovered, but the right-hand navigation bar has not been expanded.

With the mouse, right-click the desktop to reveal a selection panel. This panel allows you to select and start emulators, the HMC console, and browser sessions.

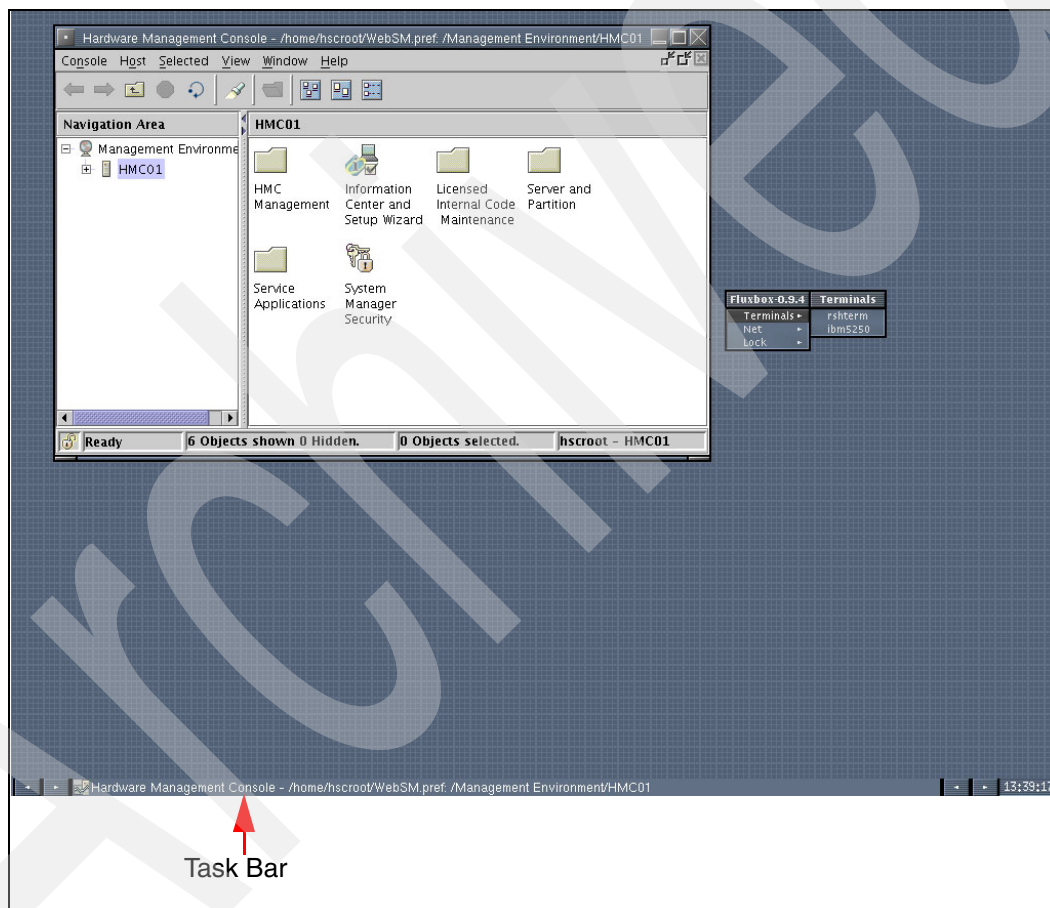


Figure 3-12 Fluxbox Desktop

3.4.1 Terminal menu

Here we discuss the items on the Terminal menu.

rcshell (or restricted shell): This is a command line interface. You can enter the LPAR migration command (see Chapter 9, "Migration of existing LPAR profiles to HMC" on page 251, and for other HMC commands, see Appendix A, "HMC command list" on page 473). While this is a Linux command line, you have a **very** restricted set of available commands you can use.

ibm5250 (5250 Emulator): This is a 5250 emulator session that can be used as the console session for any i5/OS partition you choose to connect to it. You should be aware that a limited number of 5250 emulation sessions should be started. The performance of the HMC could be degraded by starting too many sessions. All sessions traverse the same private LAN and this has limited bandwidth. See Figure 3-13.

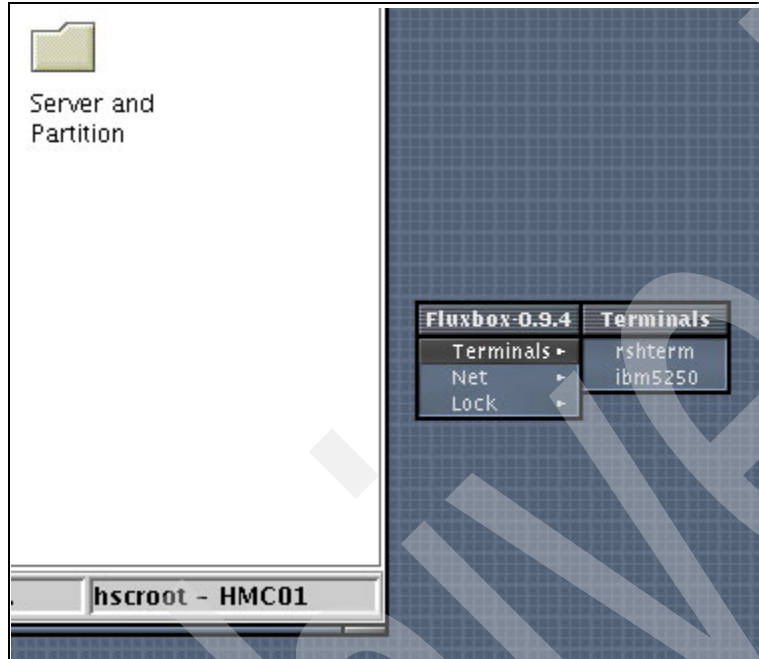


Figure 3-13 Fluxbox Terminals menu selection

3.4.2 Net menu

Here we describe the items on the Net menu.

Opera Browser: This provides a standard browser for accessing the embedded Information Center, the Information Center, and the World Wide Web. See Figure 3-14.

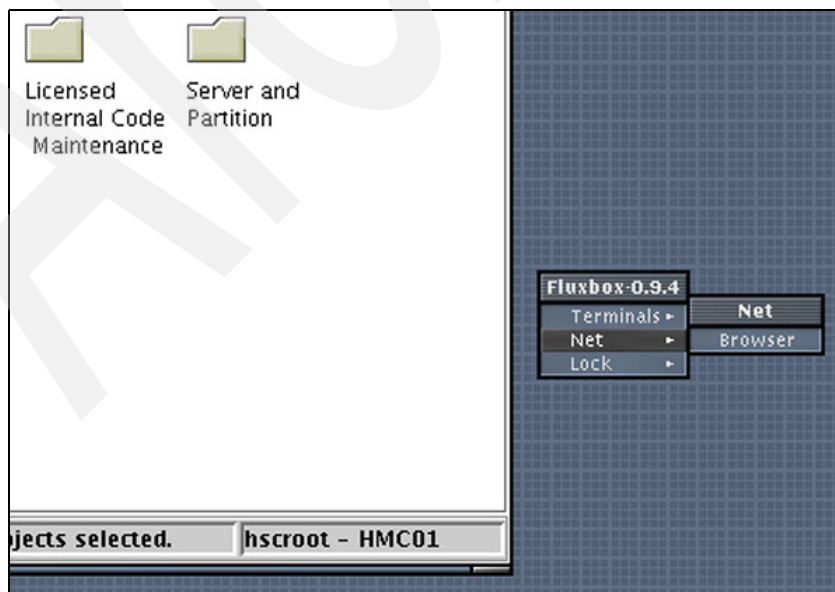


Figure 3-14 Fluxbox Net menu selection

3.4.3 Lock menu

Here we describe the items on the Lock menu.

Lock menu option: This locks the HMC console session. To return to the HMC console, you will need to supply the user ID and password that were used to initiate the previous session. The locked out session is expecting the user ID and password of the previously logged in user. See Figure 3-15.

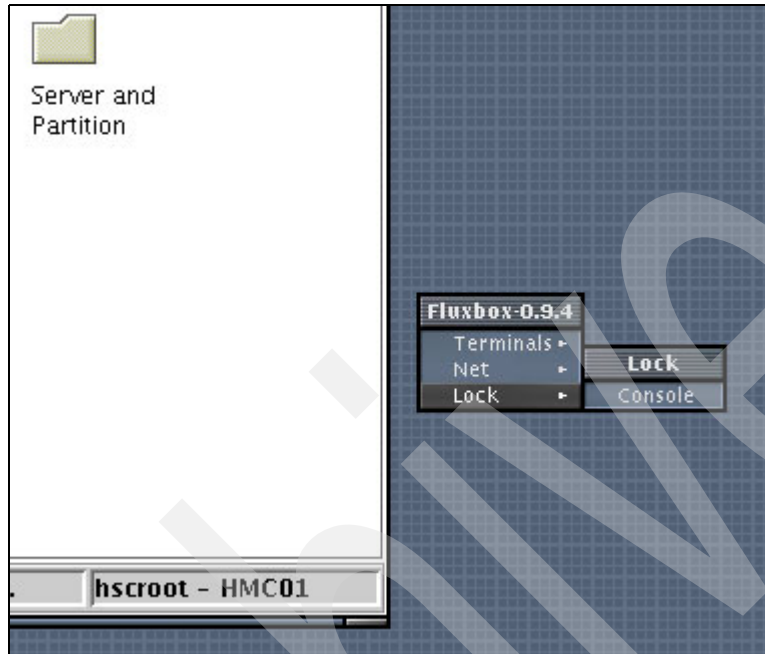


Figure 3-15 Fluxbox Lock menu selection

3.4.4 Shutting down or logging off the HMC

Here we show you how to log off the HMC console and correctly shut down:

1. Select the "console" from the top pull down menu. In this menu you can add and remove hosts, change the desktop theme and fonts, start an `rshterm`, or exit from the HMC, as shown in Figure 3-16.
2. Click **Exit** and the confirmation/menu panel appears. See Figure 3-16.

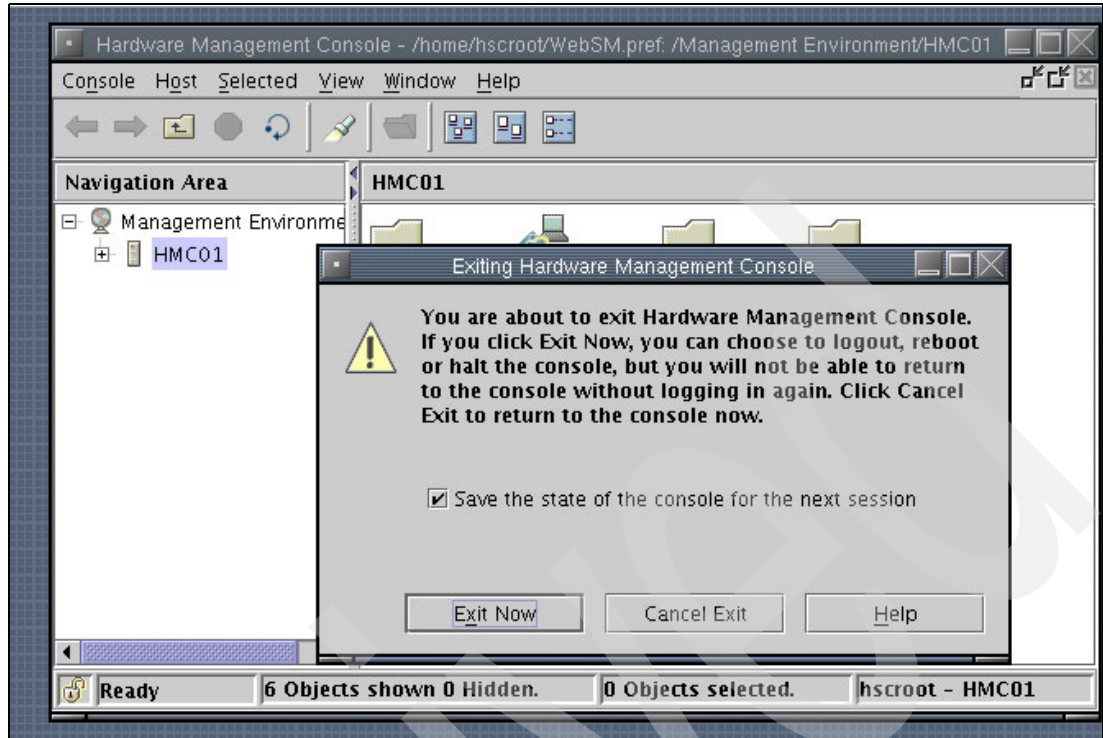


Figure 3-16 Exit the HMC console

3. Click **Exit**, and you are presented with a pull down menu that has Logout as default (Figure 3-17).




Figure 3-17 Exit HMC

4. Click the pull down arrow. You will see the three options, Shutdown, Reboot, and Logout. See Figure 3-18. To shut down the HMC, highlight **Shutdown Console**.




Figure 3-18 Exit pull down menu

5. Click **OK**. See Figure 3-19.




Figure 3-19 Accept shutdown

The HMC will now shut down.

Next we discuss the basic functions and terminology of the HMC.

3.5 Server and partition

If an @server i5 is to be partitioned, then an HMC is needed to create and manage the partitions. The partition configurations are implemented using profiles that are created on the HMC and that are stored on the HMC and on the service processor. A profile defines a configuration setup for a managed system or partition. The HMC allows you to create multiple profiles for each managed system or partition. You can then use the profiles you have created to start a managed system or partition with a particular configuration. There are two types of profiles, a partition profile and a system profile.

Partition profiles

A partition does not actually own any resources until it is activated. The resources that are specified for the partition are stored in the partition profile. One partition can have multiple profiles, and which partition profile is activated depends on what resources the partition will use. Only one partition profile can be active at a time for a partition.

For example, one partition profile may be configured to use 3 processors, 6 gigabytes of memory, and I/O slots 6, 7, 11, and 12 on a particular bus. At another time you may want to activate this same partition using 1 processor, 2 gigabytes of memory, and the same I/O slots. This will obviously have an impact on the total system resources, so careful planning and a clear understanding of the total system resources are required.

An example of multiple partitions needing to share the same slots might be a shared tape drive IOA and IOP.

Partition profiles are not affected when you use the dynamic logical partitioning feature to move resources from one partition to another. This means if you dynamically move a processor or some memory into a partition, then the partition profile will not be updated to reflect that change, and the next time the partition is powered off/on, it will revert back to the configuration as it was created in the partition profile. Any I/O resources that are dynamically moved out of a partition to another one, will stay with that partition even if the original owning partition is powered off and back on. The partition will not try to take back the I/O resource during the power off and power on in this partition, if it has already been allocated to another partition.

System profiles

Using the HMC, you can create and activate often-used collections of predefined partition profiles. This list of predefined partition profiles is called a system profile. The system profile is an ordered list of partitions and the profile to be activated for each partition. The system profile is referred to when the whole system is powered off and back on. Depending on which system profile is chosen when the system is powered on, this will determine how the system resources will be allocated out to each partition.

For example, a company may want to switch from using 12 partitions to using only 4. This would mean deactivating the 12 partitions and then activating a different system profile that defines only 4 partition profiles. So the process of partition activation is grouped to one task.

With this concept of profiles in mind, we will now look at the functions available under the server and partition icon.

3.5.1 Managed system properties

If the server management icon in Figure 3-20 is expanded as shown in Figure 3-20, the contents pane shows that, in this example, we have one system (R5-595i) being managed by this HMC.

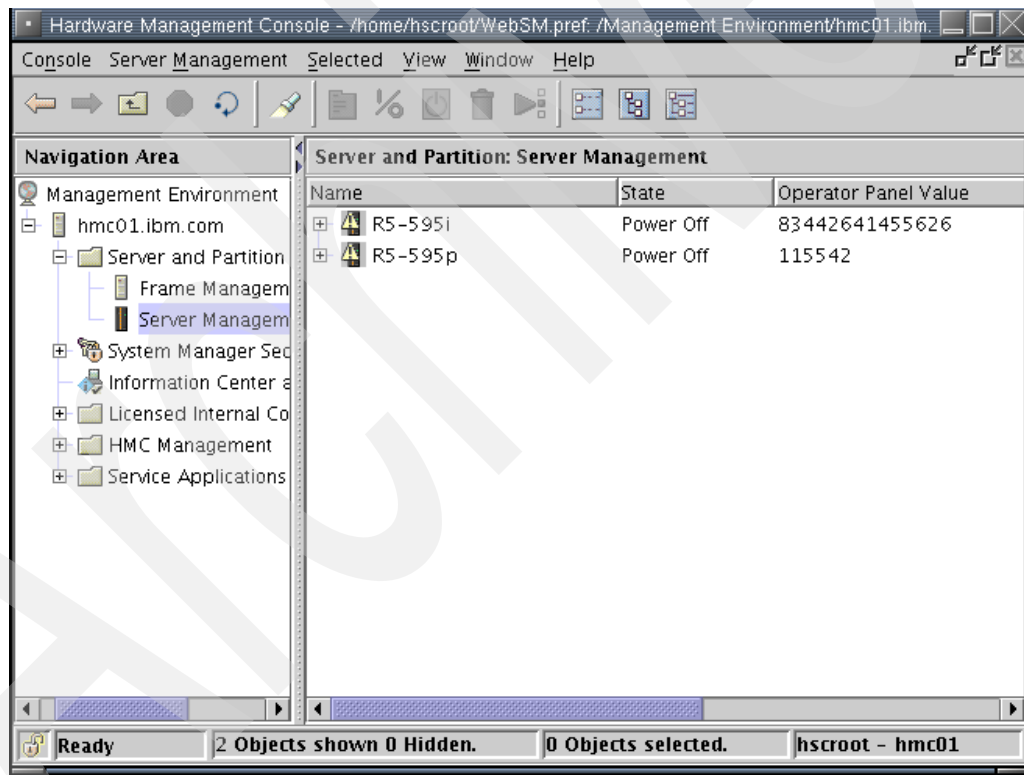


Figure 3-20 Server and partition starting point

The managed system name is R5-595i. Notice that there is no plus sign next to the “Partitions” object; this managed system has had no partitions defined. The next few screen captures show the information that is available in properties for the system R5-595i.

To access the Properties panel, right-click **R5-595i** → **Properties**.

An alternative method to access the same information is: Click **R5-595i** → **Selected** → **Properties**.

System properties - General

These properties relate to general information about the server.

The Name of the processor is a shipped value that can be user defined. The Serial Number and Type/Model are fixed by manufacturing. Type/Model may be changed by an MES upgrade.

State is the value shown on the overview panel and indicates the current server status. Possible values for “State” are shown in Table 3-1.

Table 3-1 Possible values for State

Value	Description
Operating	The server is initialized and is working.
Power off	The server is powered off.
Initializing	The server is powered on and is initializing. The initialization time may vary depending on the hardware and the partition configuration of the server.
Pending authentication	The system is waiting for the password to be authenticated.
Failed authentication	The service processor password is not synchronized with the HMC password
Error	The operating system or the hardware of the managed system is experiencing errors.
Error - Terminated	Power On, dump in progress.
CoD Click to Accept	Power On, operational and waiting for CUoD. Click to Accept.
Powering Off	Power Off in progress.
Standby	The managed system is powered on using the Power on Standby option. It will stay in this state until a partition or system profile is activated. You can create and modify profiles while the managed system is in this state.
Incomplete	The HMC failed to get all the information from the managed system.
No connection	The HMC is not connected or the handshake with the managed system failed.
Recovery	The save area in the Service Processor is not synchronized with the HMC.

The service processor version is an important value and should be checked against the current supported level when the server is shipped. Service processor levels should be maintained on a regular basis.

Maximum number of partitions (254) is just an indication of the maximum. The true number of partitions a model can be capable of depends on the number of processors and the amount of resources available to create partitions.

The value of Service Partition displays the name of the service partition if one has been allocated. In our example, we have not yet created a partition, so a service partition has been assigned. A service partition is an i5/OS partition that has been designated to report and monitor errors when the HMC undergoing maintenance or is not available. If there are updates that affect the service processor, they will be in the form of MHxxxx PTFs. If your HMC is not available and you want to apply these PTFs, you must use the ASM interface to change the update policy from HMC managed to i5/OS managed. On an System i5 system only, an i5/OS partition can be a Service Partition. On a p5 server, either an AIX or a Linux partition can be a service partition.

The Power Off policy check box does exactly what it says. If checked, the server will power down after all the partitions are powered off. See Figure 3-21.

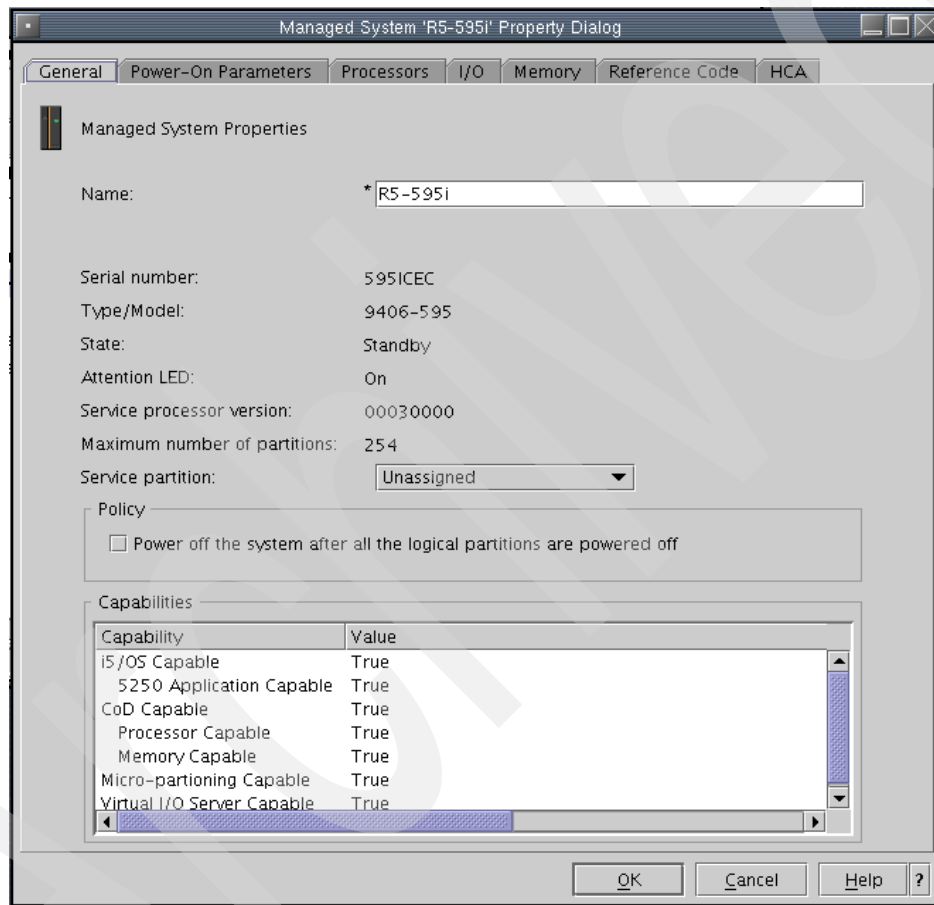


Figure 3-21 Managed system - General properties

Some of the information is obvious, but the Service processor version and service partition are new concepts. Depending on the model of the @server, the four capabilities will be listed in the bottom pane and may be true or false. The four capabilities are:

- ▶ CoD capable. True if the managed system has capacity upgrade on demand capabilities.
- ▶ OS400 capable. True when the managed server is capable of running OS400:
 - Processor capable. True when the managed server has CoD enabled and processors are available.
 - Memory capable. True when the managed system has CoD enabled and memory is available.

See Figure 3-22 for an example of managed system IPL properties.

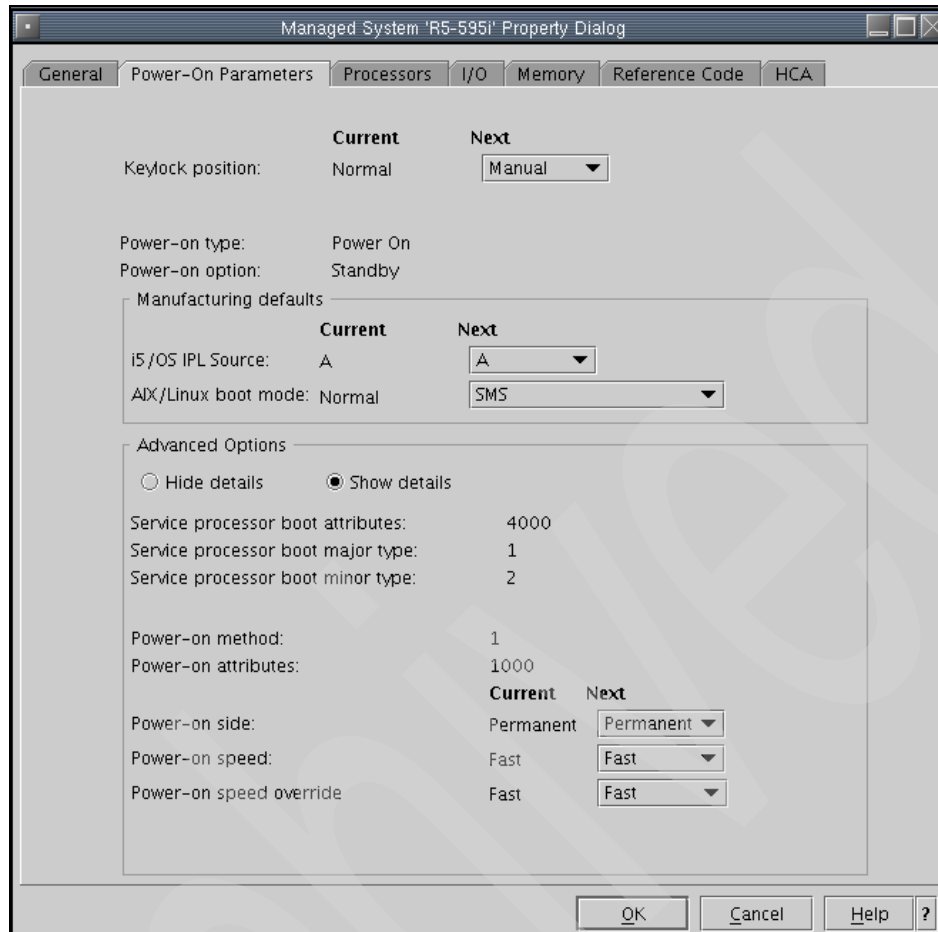


Figure 3-22 Managed system - IPL properties

The Power On parameters tab (Figure 3-22) shows information related to how the partition will be booted up or powered on. The drop down boxes for Power On have the following options:

- ▶ Keylock position — Normal or Manual. These are the same condition states as the existing 8xx or earlier model keylock states.
- ▶ Power-on type — Fast or slow. This indicates how the system will boot up and whether it will IPL the full partition hardware or not. Fast can be a very speedy IPL, a few minutes.
- ▶ Power-on option — standby. This is an important condition and allows the server to power on without starting or sensing the bus structure. When migrating LPAR configs, you **must** start in standby and have no new hardware attached, or you will lose your LPAR bus numbering.
- ▶ i5/OS IPL source — A, B, C, D. Again, these are the same as the existing i5/OS IPL states, A (permanently applied PTFs), B (temporary applied PTFs), C (service mode), and D (IPL from external media, CD, DVD etc.).
- ▶ AIX/Linux boot mode — normal.
- ▶ Advanced options — show or hide by clicking radio buttons. This will then display the following additional information.

- ▶ The following three attributes reflect physical hardware boot conditions, not the partition:
 - Power-on type — permanent or temporary; these are similar to the i5/OS IPL A and B values.
 - Power-on speed — fast or slow; this indicates how much hardware/microcode checking is done during IPL.
 - Power-on speed overrides — fast or slow; indicates how many hardware diagnostics are carried out during the power-on phase.

System Properties - Processors

Figure 3-23 shows a very useful informational panel. This is where you can see how all of the system processors are allocated. You can also see how the processors are configured, dedicated, or shared.

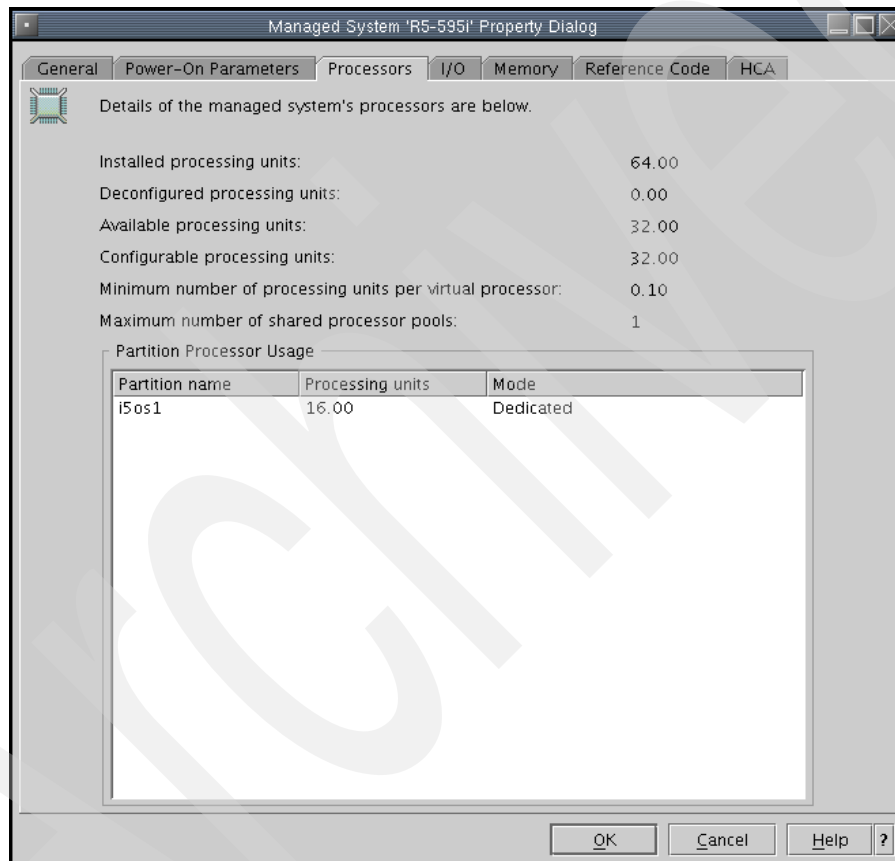


Figure 3-23 Managed system - Processor properties

The processor properties page shows the total number of available processors on the system which are the unassigned processors that are available. The configurable processors are those that are already configured in partitions and include all the possible processors on the system, including any Capacity Upgrade on Demand (CoD) processors not currently implemented. The total available processor units is the number of processors on the managed system, excluding any CUoD processors not implemented, times the minimum processor units per processor.

The Processor usage pane shows how the processors are currently divided up across the partitions. As we have no partitions created on the example server, there are no values in this area.

System Properties - IO resources

When expanded, Figure 3-24 shows the I/O resources across the managed system.

The example shown has two units or frames, a 5094 expansion tower, and a 5088 IO expansion tower. Unit 5094 contains bus 23 to 25. We have not shown the buses in the other unit. The cards in the card slots are shown as the generic card types. For example, card slot C05 in bus 25 contains a PCI I/O controller and the type indicates that this is a #2844 card. Note that the slots/cards are unassigned, as they are not currently being used by a partition.

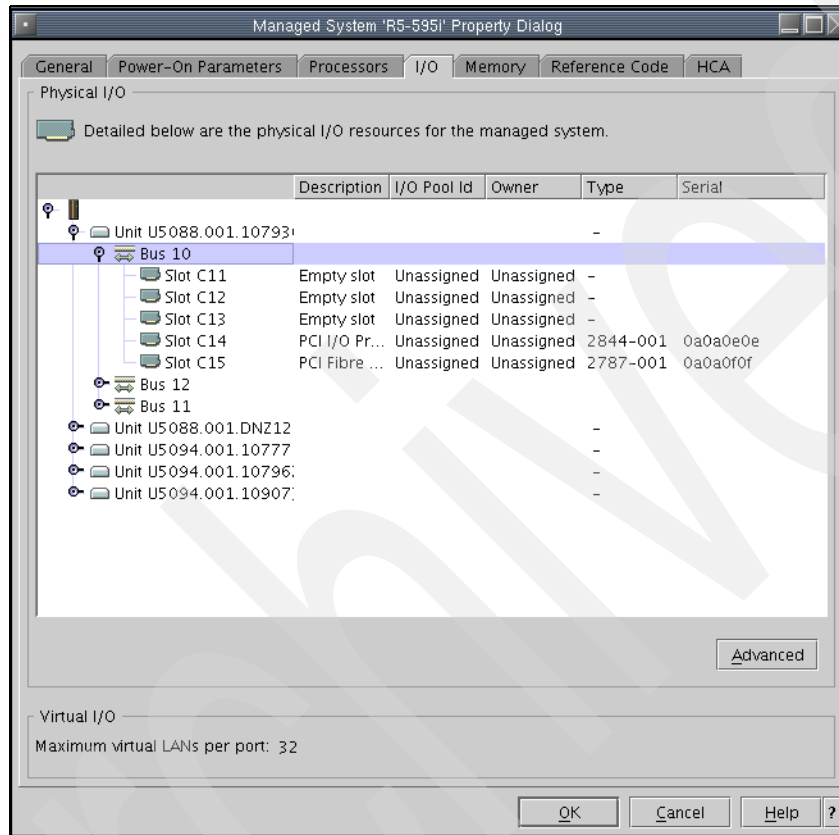


Figure 3-24 Managed system - I/O slot properties

The Hypervisor (Figure 3-25) is code that runs on the managed system and allows the three operating systems (AIX, i5/OS, and Linux) to run on the one physical system. The Hypervisor will only report the Custom Card Identification Number (CCIN), if it can recognize the card. Certain cards need a driver loaded into the card for it to report the CCIN number, and hence they may not report it to the Hypervisor. For i5/OS partitions, a rack configuration listing will still be able to be run from the partition to display the I/O resources and their CCIN numbers. An example of a CCIN number is 2849, which is the CCIN for a 10/100 ethernet card. This is important to note for that partition.

The I/O slot view will only drill down to the card slot positions. It will not show any I/O attached to the cards in the slots, and hence will not show any disk drives, tapes drives, CDs, or DVDs.

Important: When planning to partition a new system that has no operating system on it, or when adding new hardware to an existing system, this new hardware will initially only be seen by the HMC, and the IO resources pane will only show a view of the hardware down to the card slot positions.

Therefore, if you need a more comprehensive view of the I/O resources showing where the disk drives, tape drives, CDs, or DVDs are, prior to loading an operating system, then you may have to consider drawing a diagram of the system showing the locations.

Our suggestion is to **always** use the LVT and supplement it with a sketch of the system. You could photocopy out the system layouts from the i5 System Builder. Fill in the locations of the devices and indicate to which partition they are allocated.

Once the partition is created and the OS is loaded, hardware assigned to that partition can be viewed from hardware management programs.

The column labelled Pool ID designates the I/O pool that the card slot is assigned to. The I/O pools are used in conjunction with clustering and allow systems in the same cluster to pool their I/O in a common pool. See Figure 3-25.

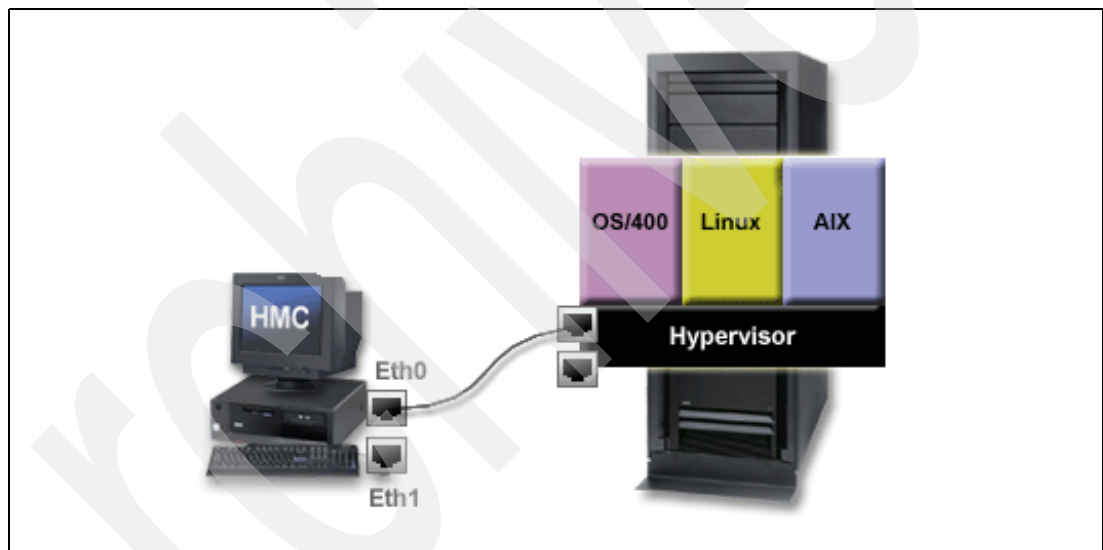


Figure 3-25 Hypervisor runs under the partitions

System Properties - Memory resources

The memory properties are shown in Figure 3-26. The available memory is the unassigned memory that is available to allocate to partitions. The configurable memory is the total amount of memory on the managed system.

The memory region size determines the smallest increments in which the memory can be allocated to a partition. There is a minimum amount of memory that must be allocated to a partition for it to be able to run, and this will be either 128 MB or 256 MB depending on the value that has been set for the maximum memory to be used in a partition. Any increase in the memory size over this minimum would be done in increments of 16 MB, if the memory region size was 16 MB as shown in Figure 3-26. The bottom pane shows the amount of memory allocated to each partition.

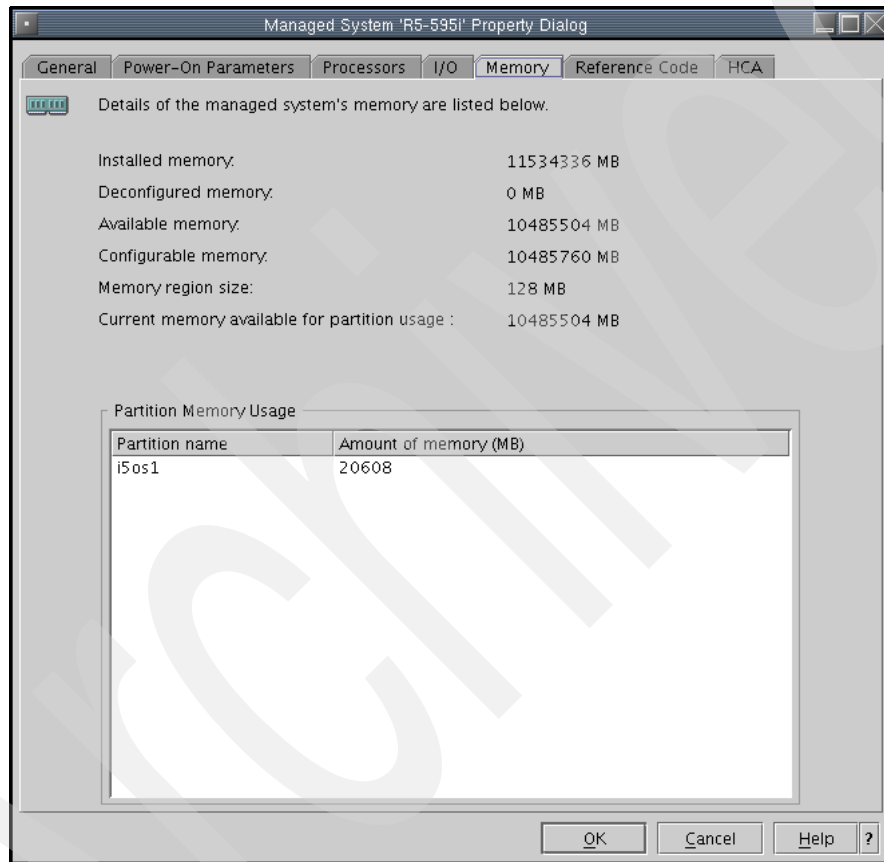


Figure 3-26 Memory resources display

System Properties - System Reference Codes (SRCs)

The tab under properties shows the SRCs for the managed system (Figure 3-27). This will show up to the last 25 System Reference Codes (SRC) as reported through the Hypervisor to the HMC. There is an option to show the details of each of these SRCs.

The pull down list allows the user to change the number of the SRC history to be displayed. Highlighting a particular SRC code and selecting “Details” will display more information on the code, when available.




Figure 3-27 Managed system SRC panel

System Properties - Highlight a Host Channel Adapter (HCA)

The last tab under properties shows the HCA for the managed system (Figure 3-28). This page will highlight and display the channel adapter's current partition usage.

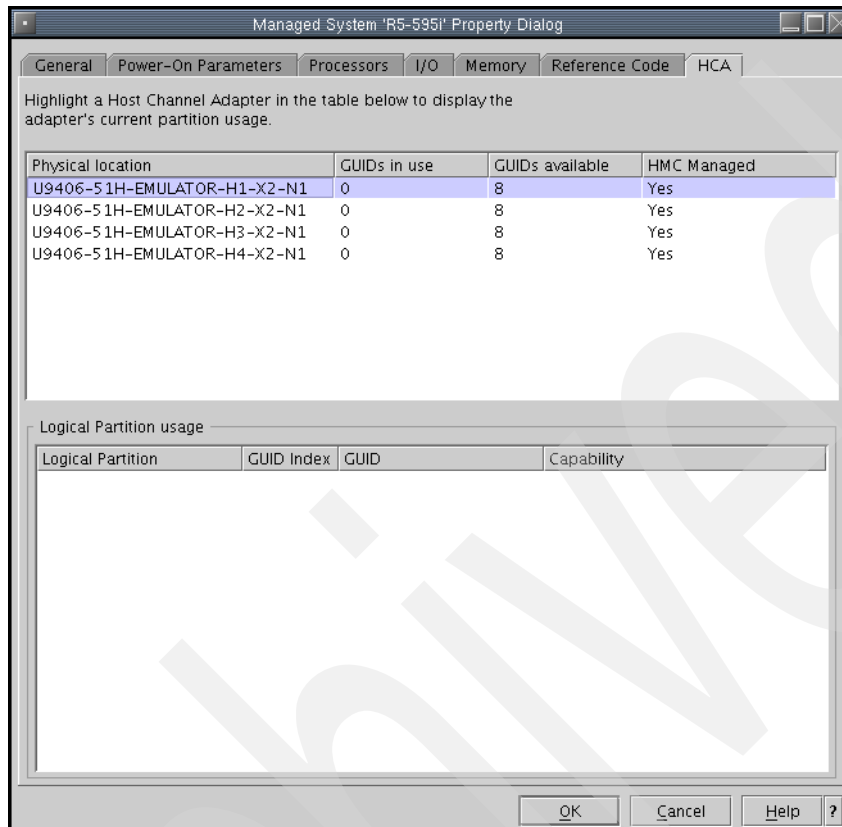


Figure 3-28 Managed system HCA panel

3.5.2 Other system-wide options

In this section we look at the other options that are available system-wide by right-clicking the managed system name. Figure 3-29 shows the other options.

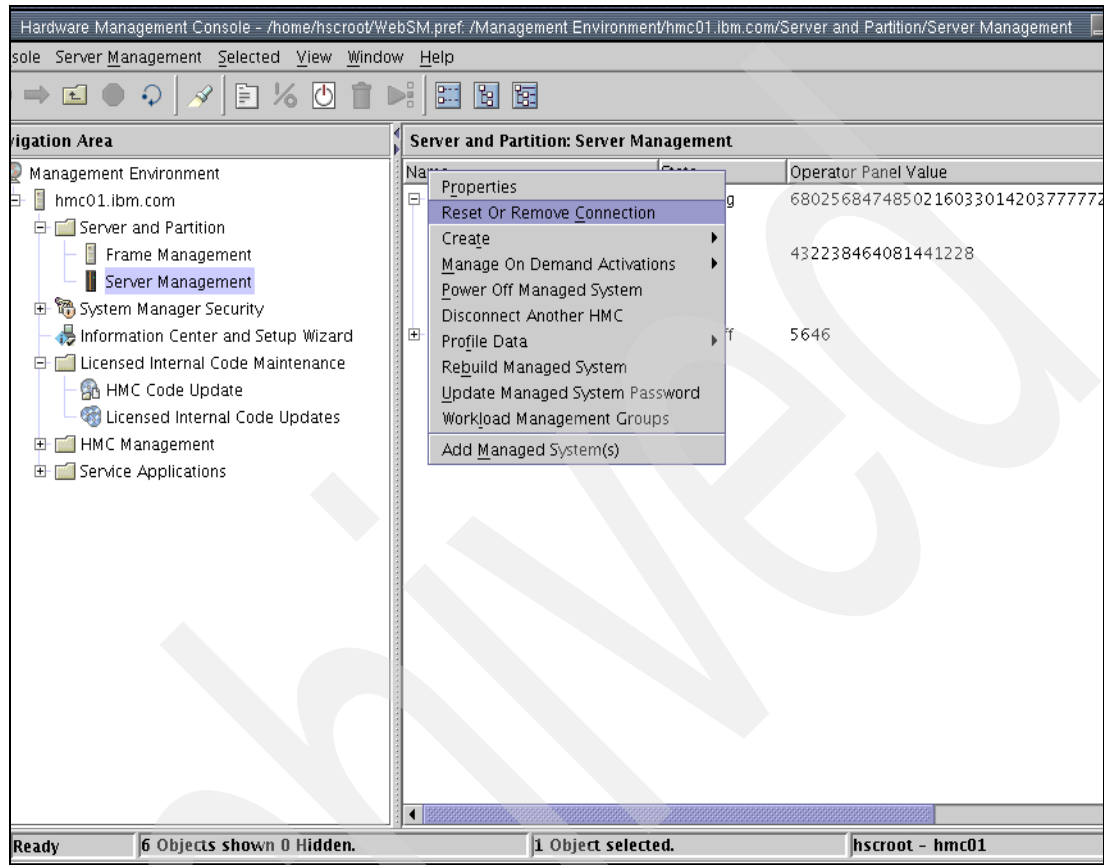


Figure 3-29 System wide functions

The following list provides information on each selection:

- ▶ The properties selection has already been discussed in previous sections of this chapter.
- ▶ Reset or Remove connection — This task works on the actual server itself. If for some reason the server and the HMC lose contact, a Reset would attempt to restore contact. A Remove would remove the managed server from this HMC sphere of operation.

An example of a where a Reset might be used is if the HMC and server were placed on a public network and this network failed. When the network was restored, the HMC detail would probably say “No connection”. You would use a Reset to attempt to recover.

An example of a Remove might be if an HMC was no longer required to control a server, if the HMC was a redundant HMC controlling an additional server, while the “normal” HMC was unavailable for a short period. Once the normal server returned to service, you would use the Remove to delete the managed server information from the redundant HMC.

- ▶ The Create option is taken to create a new logical partition. This is covered in detail in Chapter 5, “Partition creation using the HMC” on page 139.
- ▶ The Capacity on demand option is provided as the management panels for the CoD function.
- ▶ The power off option will power off the entire managed system, which includes all the partitions. This option needs to be used with caution.

Very Important: Powering off an i5/OS partition from the HMC should only be used if all other options fail.

You should use the Power Down System (PWRDWNSYS) command from i5/OS. The Power Off partition or Power Off server options in HMC are equivalent to Function 8 or hitting the white power button twice on an existing 8xx server. This will not allow the server to shut down in an orderly manner.

You should be especially careful if you have integrated Windows servers. The time to shut down the Windows server™ must be taken into account when you power down the system. There is a new Network Server Description parameter in i5/OS V5R3 Shut Down Timer (SHUTDTIMO) that specifies how much time Windows servers are give to power down. This tries to prevent you powering down the server to quickly and causing the integrated Windows servers to crash.

- ▶ Disconnecting another HMC forces a second HMC to release its lock on the server. See Figure 3-30.

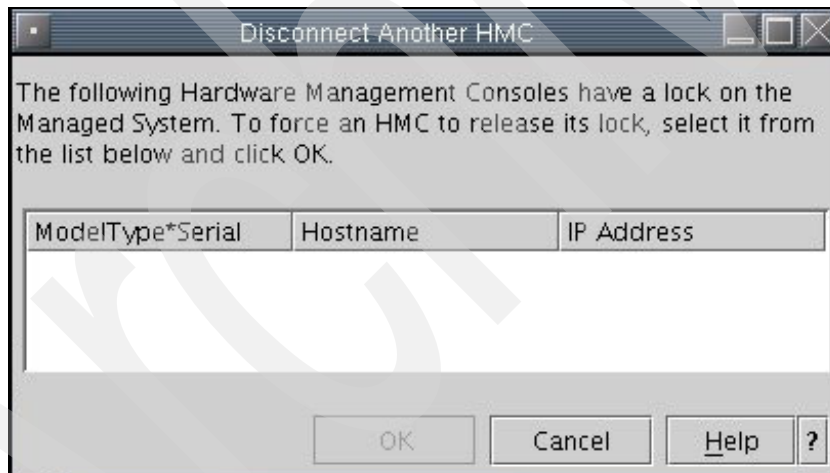


Figure 3-30 Disconnect another HMC

- ▶ Profile data is the system profile, logical partition, and partition profile information for the managed server that is highlighted. Within this option, you can:
 - Initialize all the profile data, which deletes all existing profile data.
 - Restore the profile data file if it has been lost or damaged (see Figure 3-31):
 - Full restore from selected backup
 - Backup priority — merge current profile and backup
 - Managed system priority — merge current profile and backup

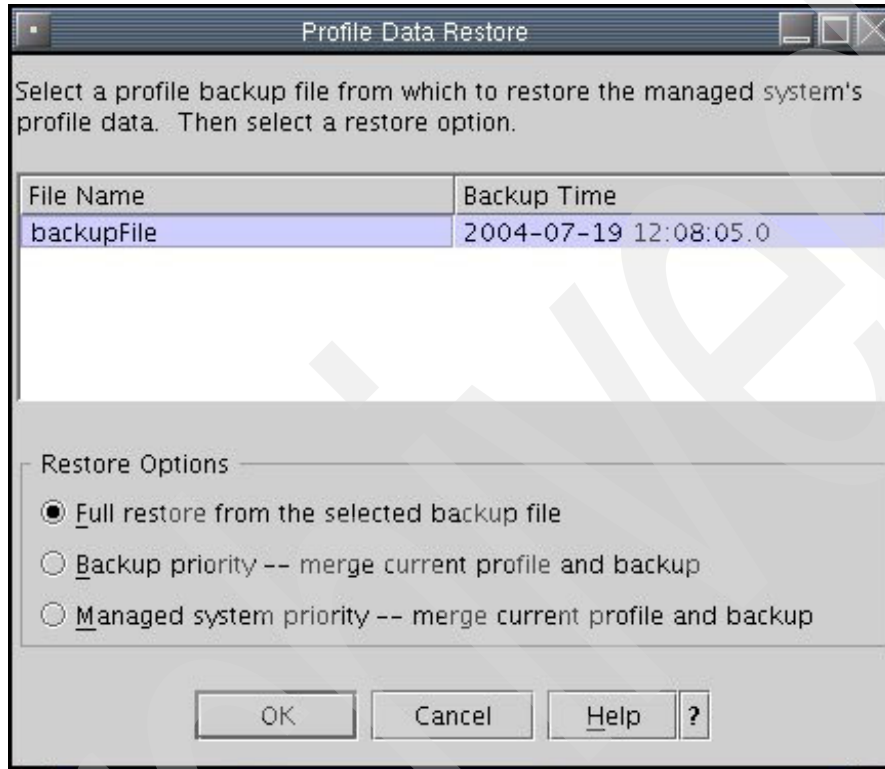


Figure 3-31 Restore profile data

- Back up profile data to disk.
- Remove profile data, selectively delete profile information.
- ▶ The option to rebuild the managed system is used to refresh the managed system information. Rebuilding the managed system is useful when the system's state indicator in the Contents area is shown as Recovery. The Recovery indicator signifies that the partition and profile data stored in the managed system must be refreshed.
- ▶ Update managed system password. This is the password that you would use to log onto the ASM interface. The pre-defined ids and default passwords are:
 - Admin / admin
 - General / general

You may have already been prompted to change these.

3.5.3 Available partition options

We have covered the functions that are available from a system wide perspective; now we take a look at the functions available at a partition level. To show the functions available for each partition, right-click the partition.

Alternatively, select the i5/OS partition and then right-click to see the options available as shown in Figure 3-32.

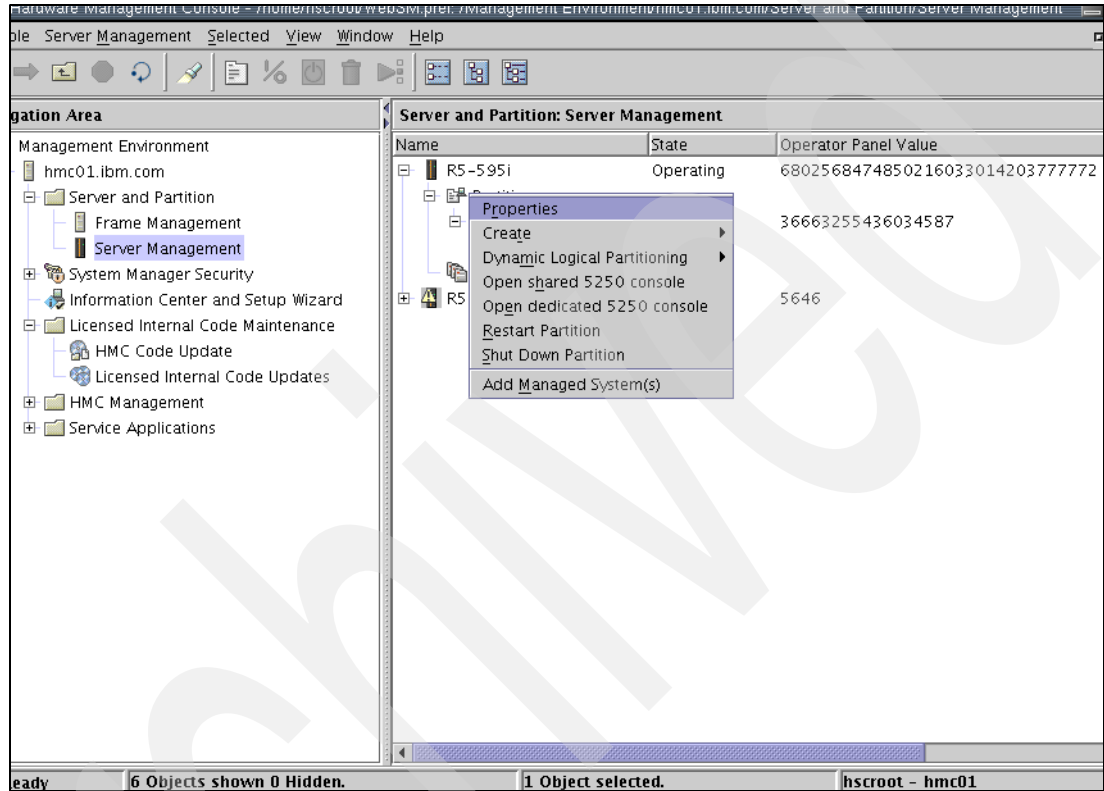


Figure 3-32 Functions available on a partition

Many of the windows under properties are self explanatory, so only windows with new functions or terminology will be discussed.

Create Partition Profile

Use this option to create an additional profile over the selected partition.

Work with Dynamic Logical Partitioning Resources

This takes you to further selections as shown in Figure 3-33. These allow you to work with IO, processor, memory, and virtual adapter resources.

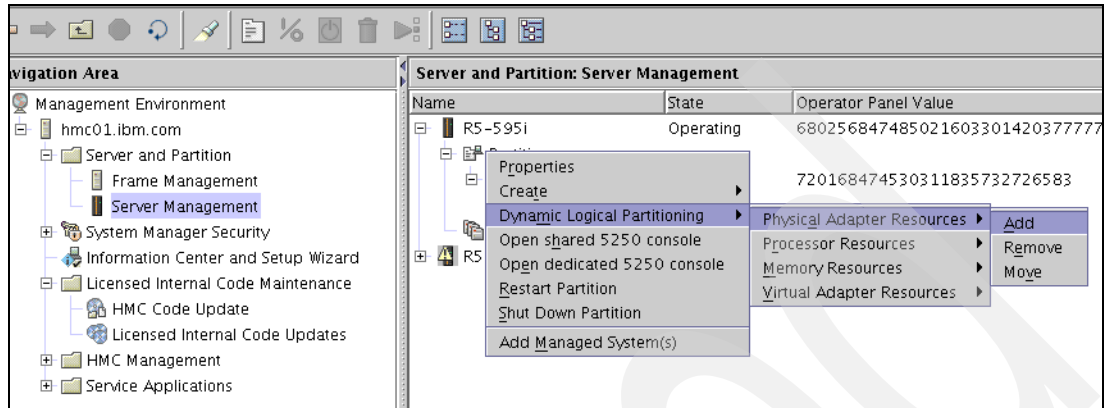


Figure 3-33 Dynamic Logical Partitioning selections

You have the option to add, remove, or move resources with the exception of Virtual Adapters, where you can only add or remove adapters.

Open Terminal Window

This selection allows you to start a 5250 terminal session with a managed server.

On the HMC, highlight the *partition* that you want to access. Then *right-click* the prompt selection menu that looks like Figure 3-34, and shows where you can select a 5250 console session.

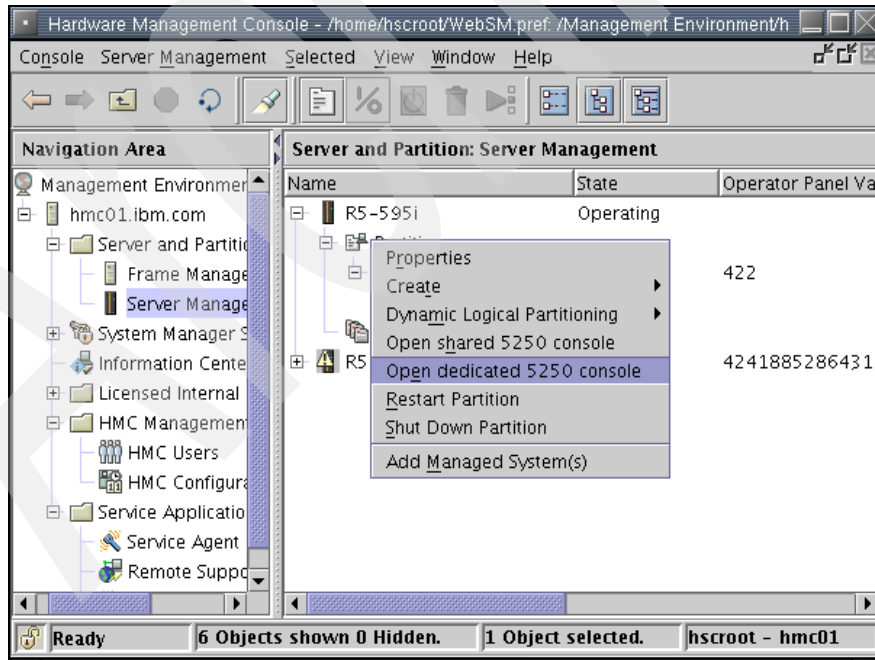


Figure 3-34 Open 5250 Console with HMC

You can choose **Open shared 5250 console** or **Open dedicated 5250 console**. Figure 3-35 shows you an example to select **Open dedicated 5250 console**.

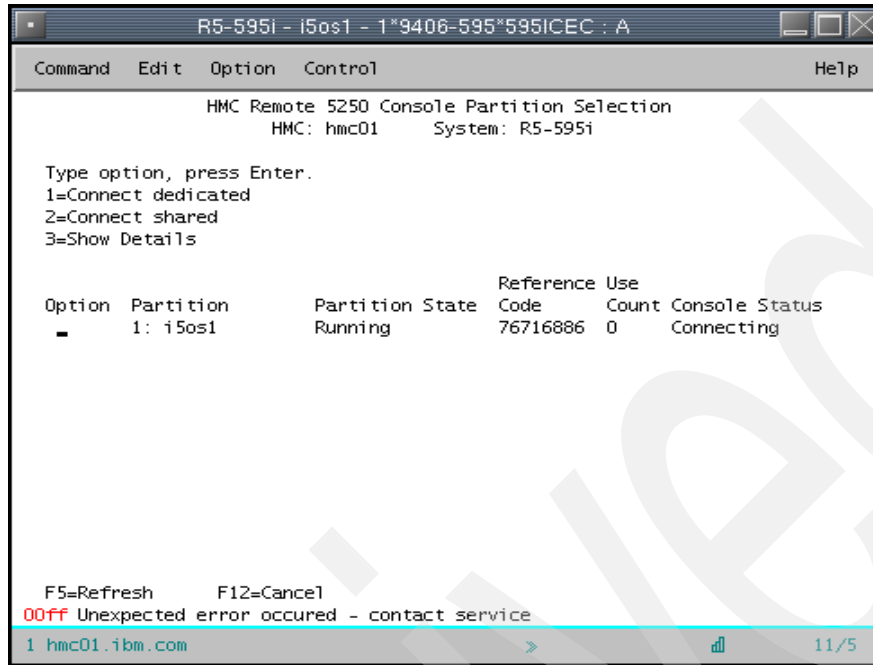


Figure 3-35 Open 5250 dedicated Console

You may need to wait a short time until the console status shows *Connecting*, then the sign on screen prompts as shown in Figure 3-36.

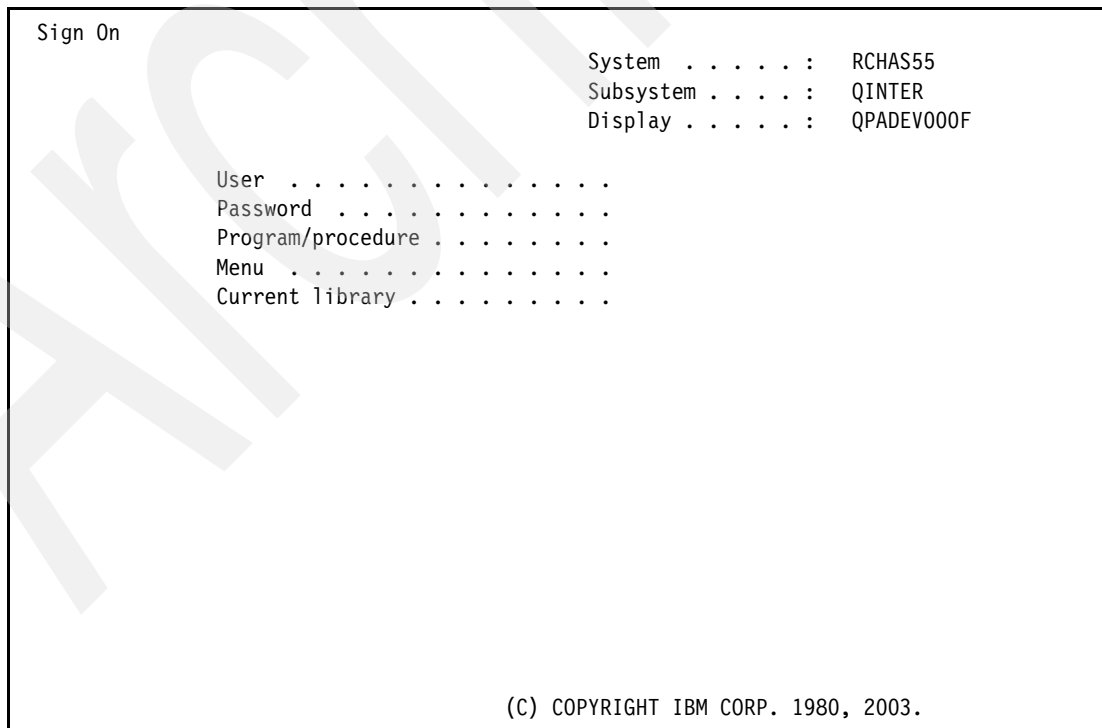


Figure 3-36 5250 console-sign on screen

Figure 3-37 shows where you select **Open shared 5250 console**.

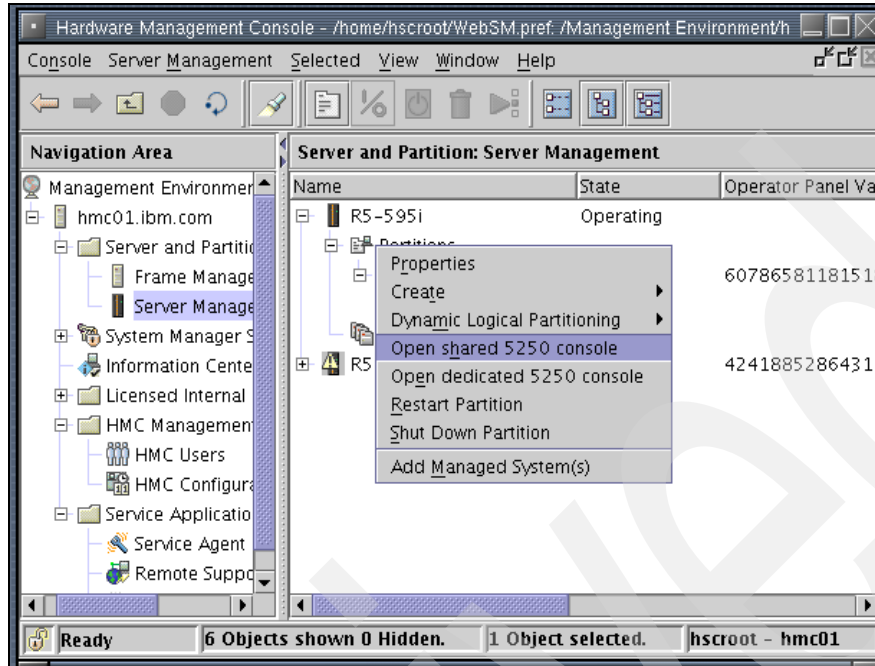


Figure 3-37 Open shared 5250 console

After you select **Open shared 5250 console**, the next panel prompts as shown in Figure 3-38, asking you to enter the session key, then press Enter.

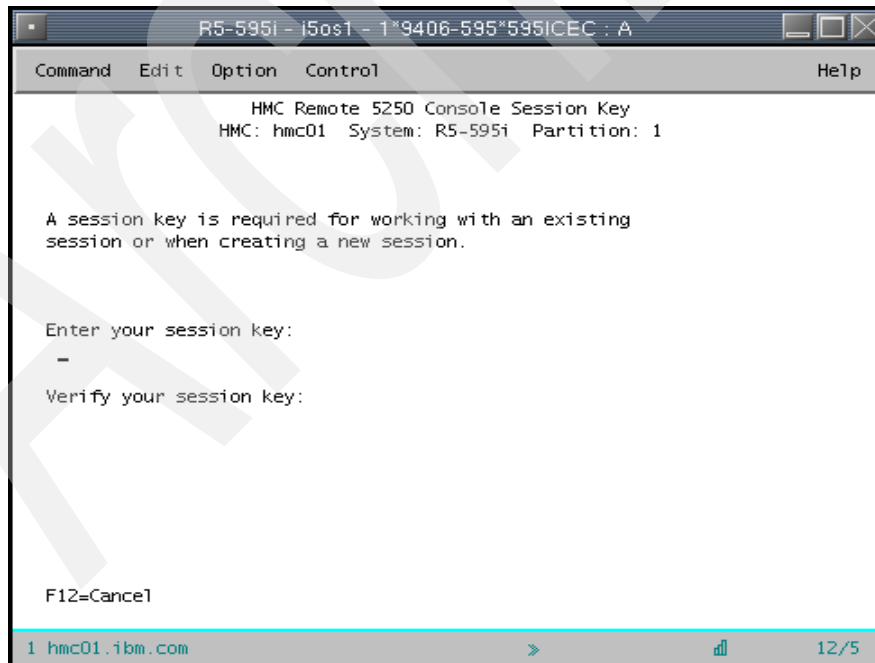


Figure 3-38 Open HMC 5250 share console -enter session key

Then a panel like the one in Figure 3-39 shows the shared 5250 console connection status. You may need to wait a short time until the console status shows *Connecting*, then the sign on screen prompts as shown in Figure 3-36 on page 76. In Figure 3-39 you can press F12 if you want to cancel the sign on.

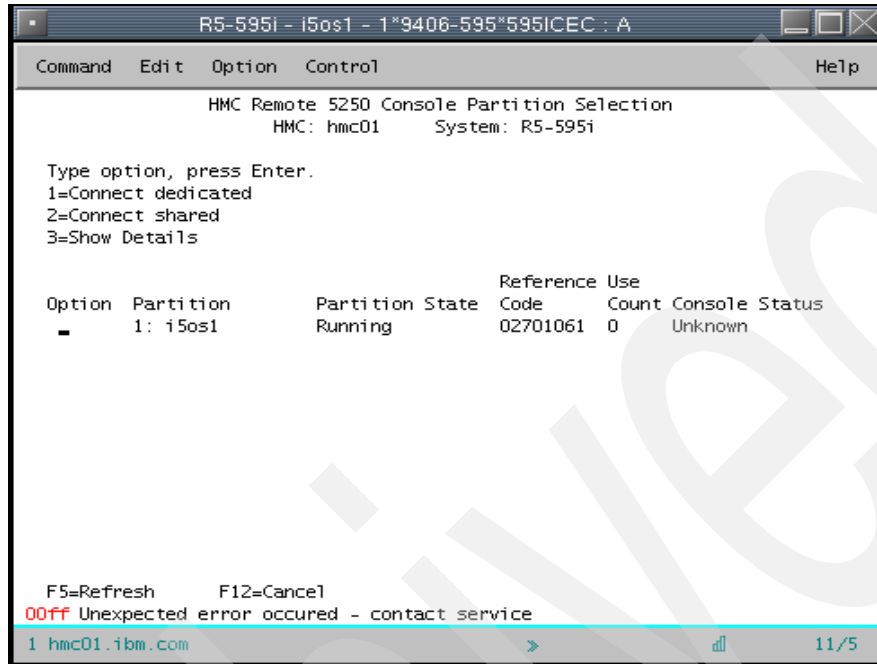


Figure 3-39 Open HMC share console

Then the next panel prompts as shown in Figure 3-40; it shows management system status.

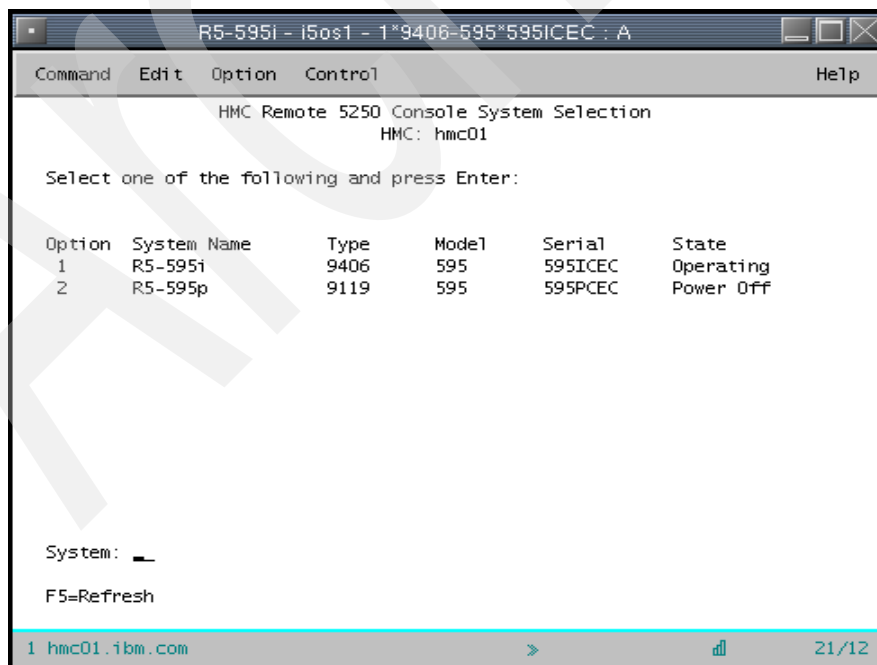


Figure 3-40 Open 5250 console-cancel

In this panel you can select **Command** → **New 5250** to open another 5250 console. Figure 3-41 shows an example to open a new 5250 session.

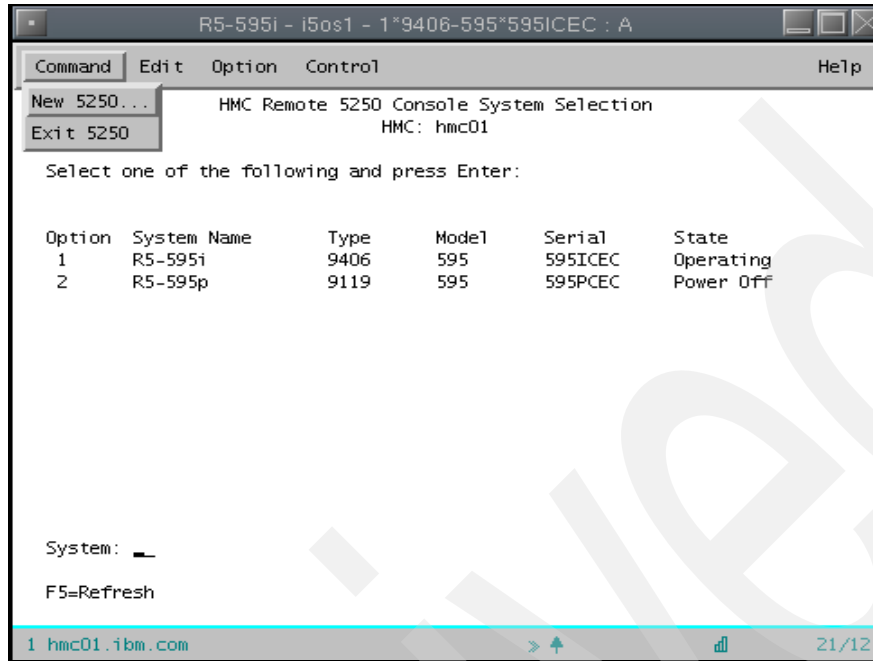


Figure 3-41 Open shared 5250 console a new session

Then the next panel prompts as shown in Figure 3-42, asking for a new 5250 host name. Enter the name and press **OK**.

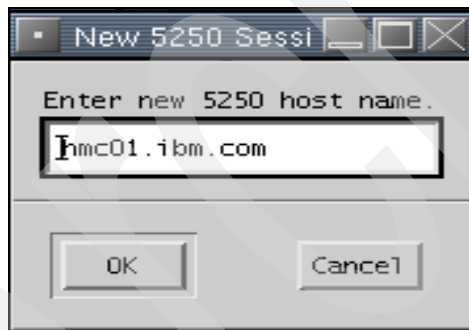


Figure 3-42 Open shared 5250 console -new 5250 host name

Then the next panel prompts as shown in Figure 3-43, asking for the HMC userid and password to open a new 5250 session.

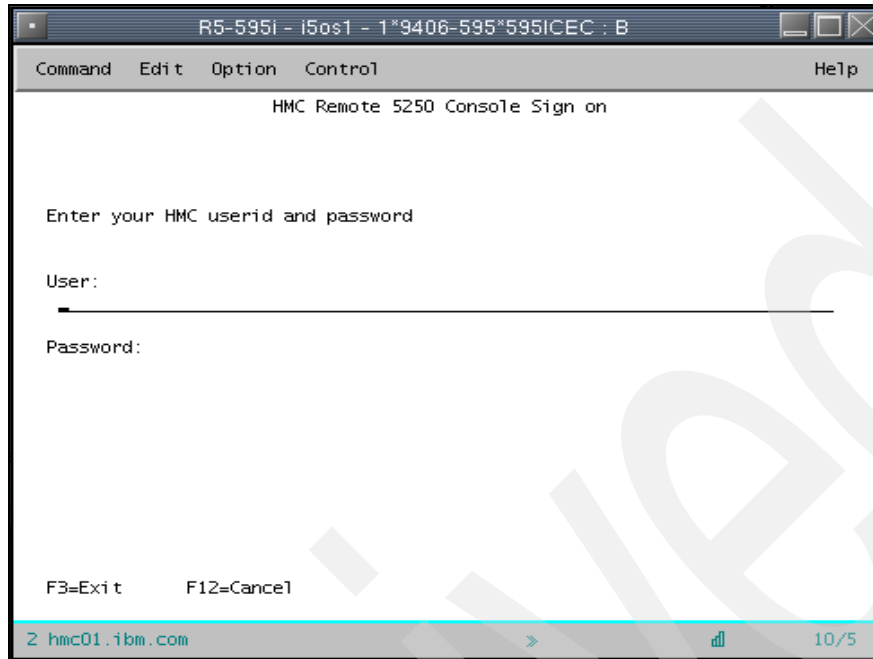


Figure 3-43 Open shared 5250 console -HMC user ID and password

After you have entered a valid HMC user ID and password, then the next panel, which looks like Figure 3-44, shows a management server on another window. Select the management system that you want to access, then press Enter.

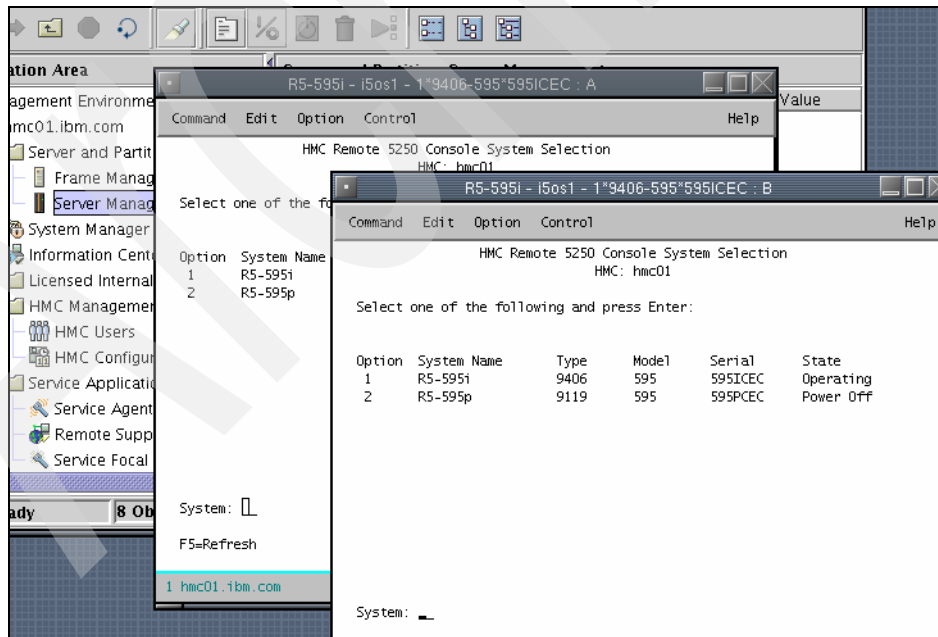


Figure 3-44 Open shared 5250 console new session

Then the next panel, which looks like Figure 3-45, shows session connection status. You may need to wait a short time until the console status shows *Connecting*, then the sign on screen prompts as shown in Figure 3-36 on page 76.

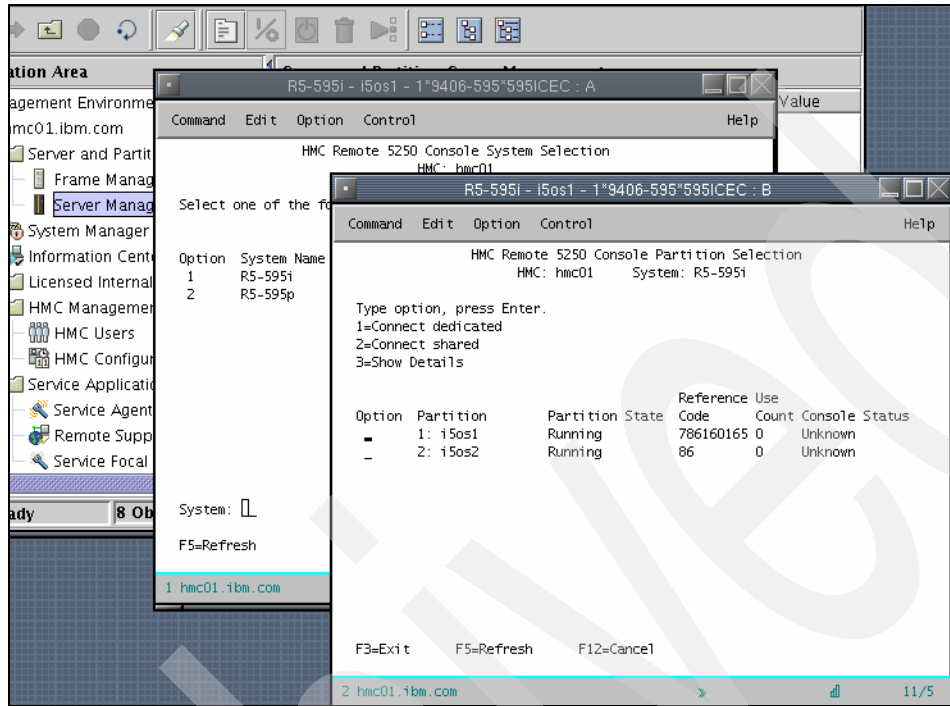


Figure 3-45 Open HMC shared 5250 session-connection status

Next, Figure 3-46 shows opening multi-5250 consoles or terminal sessions. The session's title is named as follows: "*management server name-partition name-partition number - hardware feature: A/B/C (the sequence of partition 5250 session)*". For example:

R5-595i-i5os1-1*9406-595*595ICEC:A

Note: The title's names are the same as before version 4.5

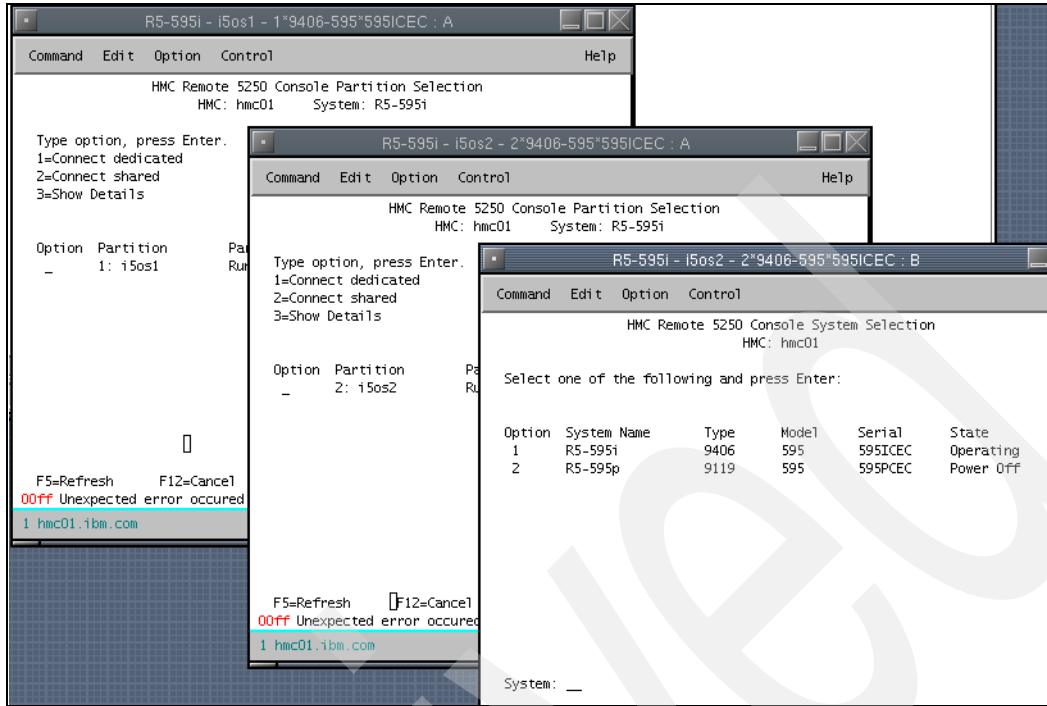


Figure 3-46 Open HMC multi-5250 sessions

The following steps show you another way to access 5250 screen. In the HMC desktop, *right-click* the prompt menu, then select **Terminal** → **IBM5250**. Figure 3-47 shows an example to open a 5250 terminal.

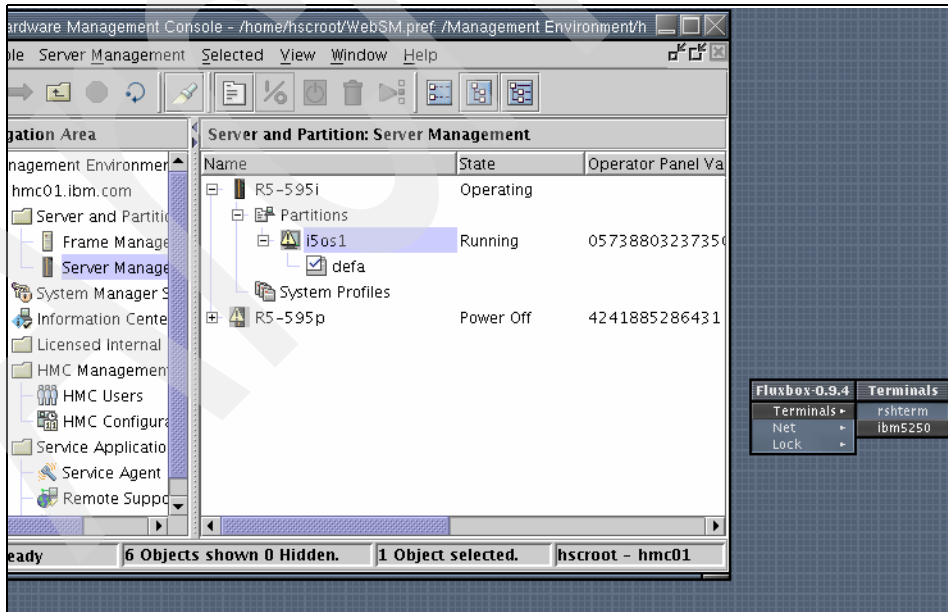


Figure 3-47 Open 5250 Terminal

Then the next panel prompts as shown in Figure 3-48, asking you to set up a 5250 session.

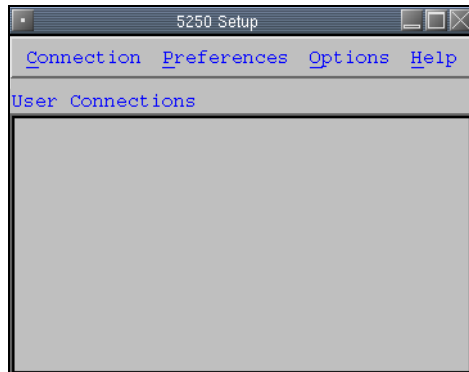


Figure 3-48 Open 5250Terminal -setup

You can select **Preferences** → **5250** as in Figure 3-49, to check the preference parameters.




Figure 3-49 Open 5250 Terminal -setup preference

Then the next panel prompts as shown in Figure 3-50. You can change it or use the default value, then click **OK**.

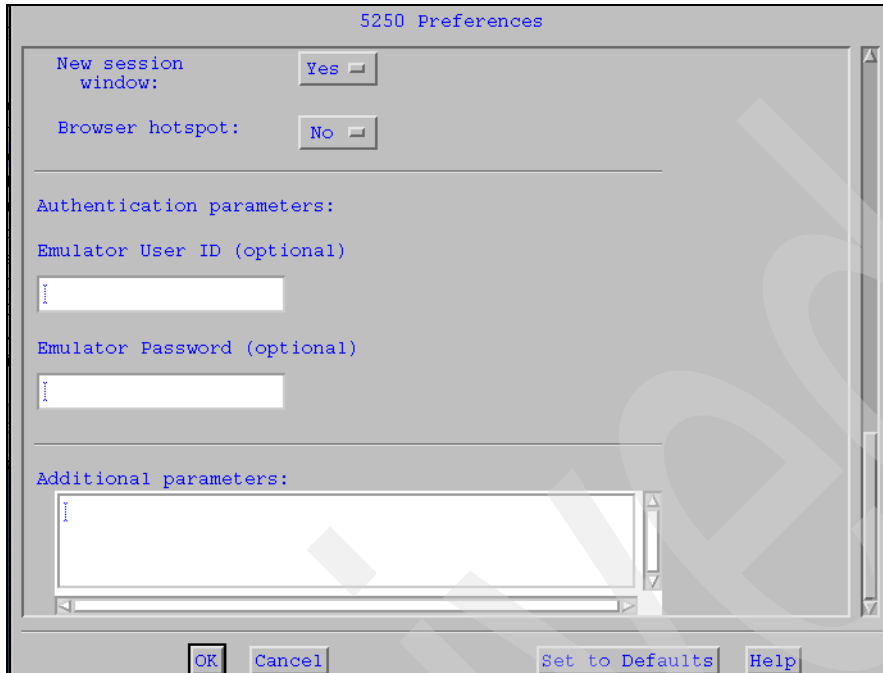


Figure 3-50 Open 5250Terminal - setup preference parameters

In Figure 3-48 on page 83, select **Connection** → **New**, then the next panel prompts as shown in Figure 3-51. Enter the *Connection Description* and correct *i5/OS Host Name or IP Address*, then click **OK**.

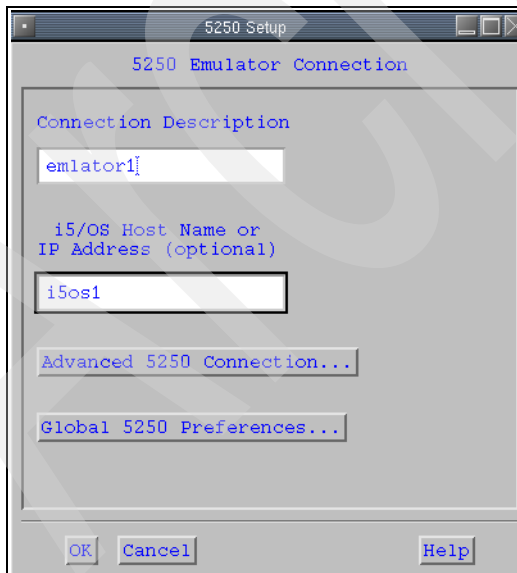


Figure 3-51 Open 5250Terminal setup connection

Then the next panel, which looks like Figure 3-52, shows a new connection added.

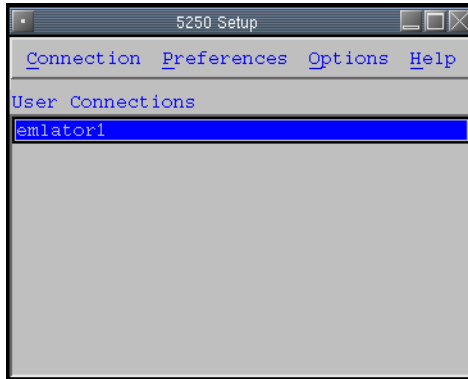


Figure 3-52 Open 5250Terminal new connection

Then, highlight the connection and select **Connection** → **connect** from the menu. Figure 3-53 shows an example to select a connection.




Figure 3-53 Open 5250Terminal connect

Then the sign on screen prompts as shown in Figure 3-36 on page 76.

Restart Partition

As explained in the text shown in Figure 3-54, this option should only be used with caution. Restarting a partition will result in an abnormal IPL of i5/OS, which is not what you want to do. Only use this option if all else fails, or under the direction of Support personnel.

As the text goes on to explain, this would primarily be used for an i5/OS hang situation. Again, this is something that does not happen often, and would result in a call to Support for assistance before using this option on a production system.

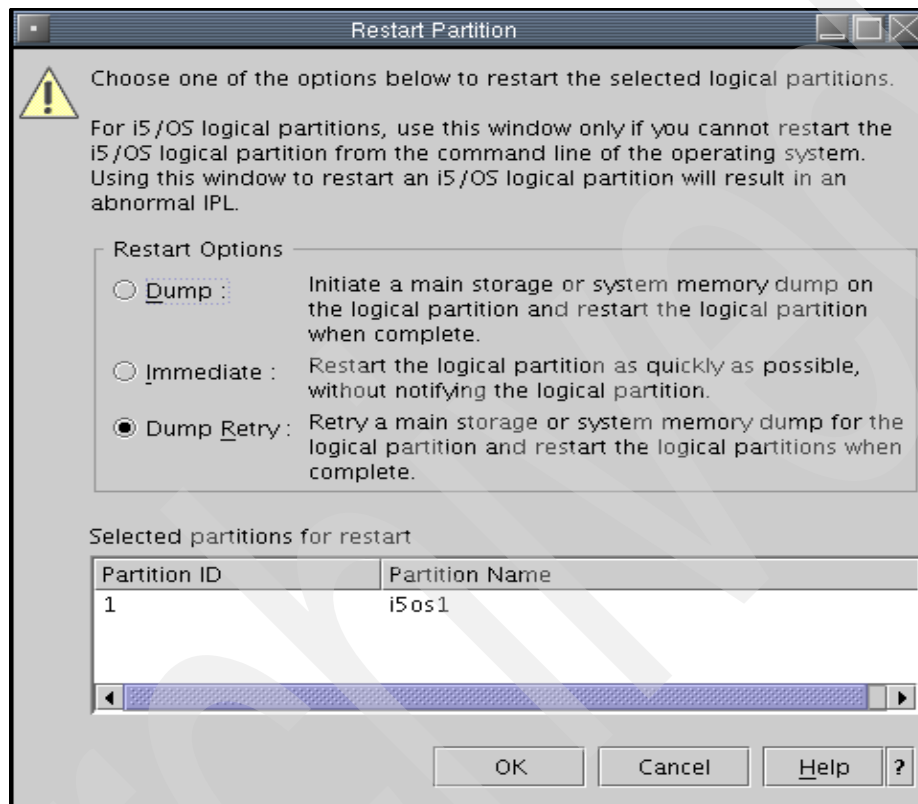


Figure 3-54 Restart partition selections

Shut Down Partition

This option is equivalent to a power failure (or pulling the power plug) on the server, or as if the partition were a separate server, experiencing a power failure.

When you select this option, you are presented with a selection panel as shown in Figure 3-55, and this offers similar options to those available with the Power Down System command — Delayed and Immediate power down:

- ▶ **Delayed shut down** — Shuts the system down in a predetermined time. If the partition has not shut down during the time period, the shut down will occur anyway and cause the next IPL to be classified as abnormal.
- ▶ **Immediate shut down** — Shuts the system down but does not allow any time for safe shut down of tasks and services, and is equivalent to a power failure. This option will result in an abnormal IPL.

Again, these options are only to be used in emergency. They do not replace the i5/OS commands.

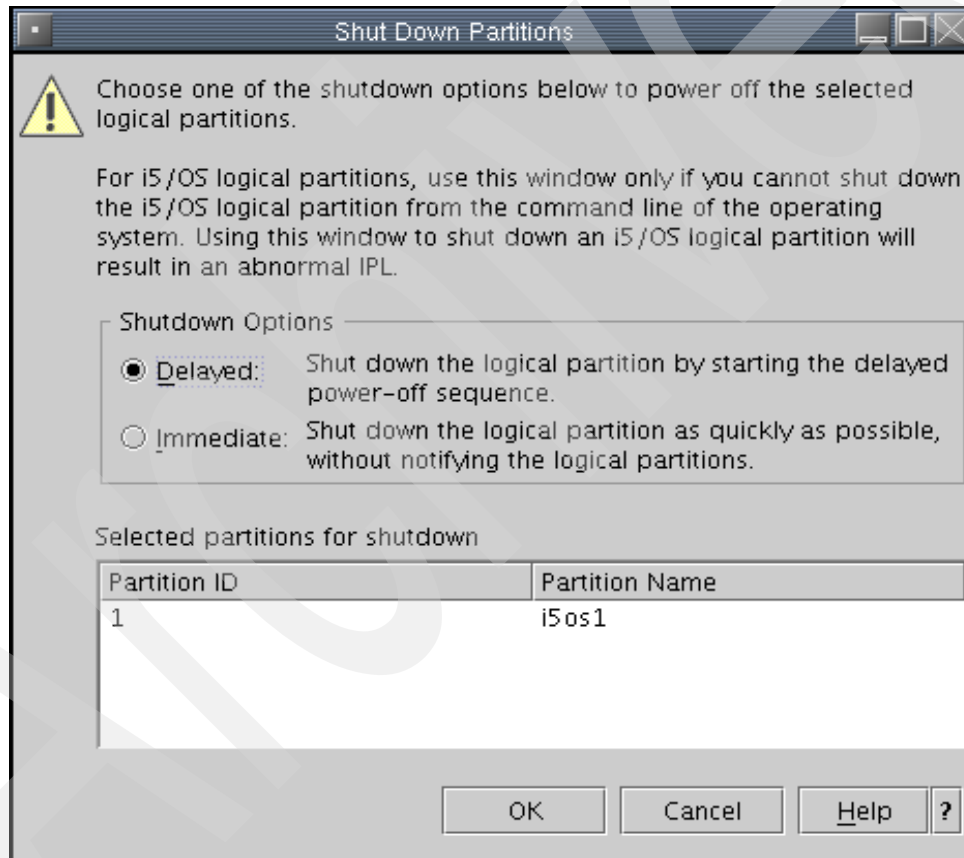


Figure 3-55 Shutting down a partition

Partition properties - General tab

This provides basic information about the operating system running in the partition and its current state, as we can see in Figure 3-56. You can see the partition ID and profile name. You can also see that the partition is running i5/OS and the version — in our example, the OS version is depicted as a series of zeroes (0s).

You can display the general properties by selecting the partition and right clicking the partition.

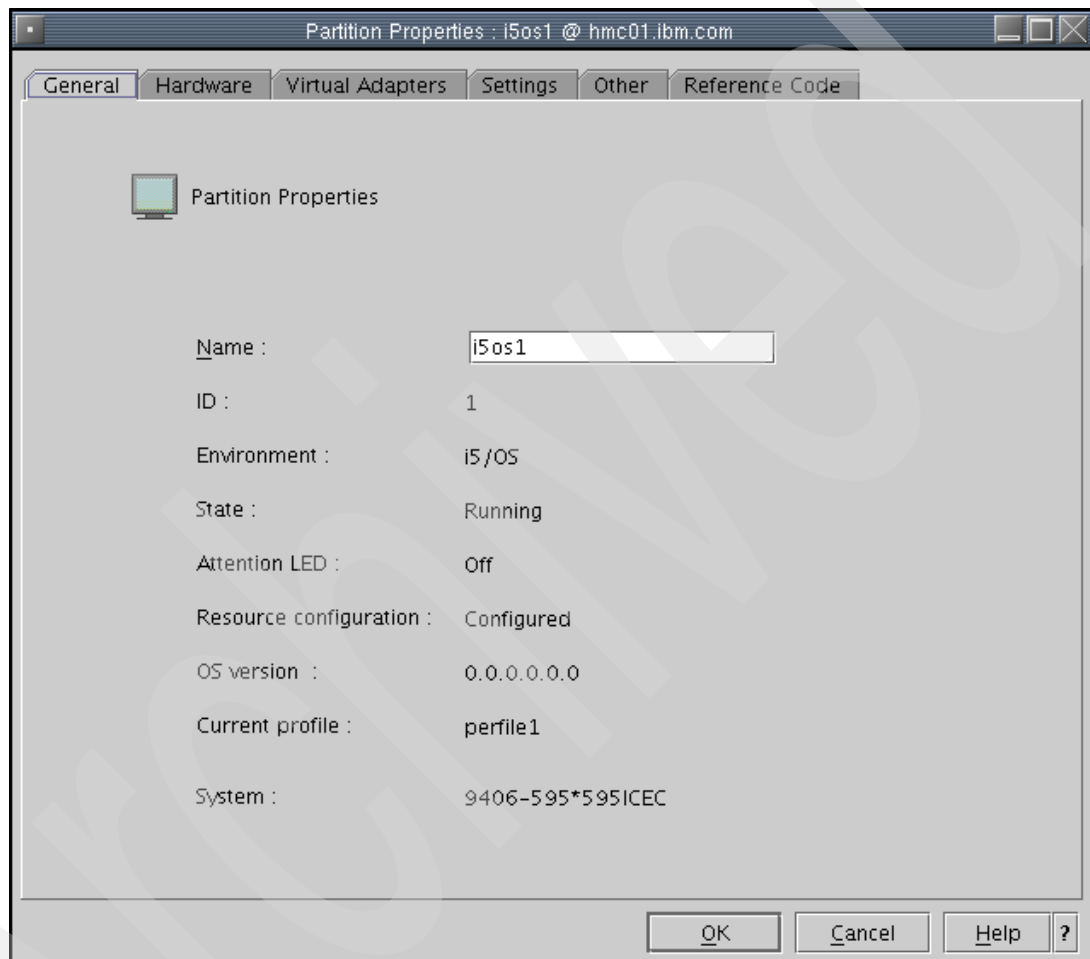


Figure 3-56 Partition properties general tab

Partition properties - Hardware tab

Under the Hardware tab, the slot level detail can be drilled down to by clicking the “twistie” on the left hand side of the pane. This is similar to the i5/OS options to display hardware resources. But it does not go down to the detail of device level, for example, no disks are displayed. You must manually view the disk or look in the OS running on the system. See Figure 3-57.

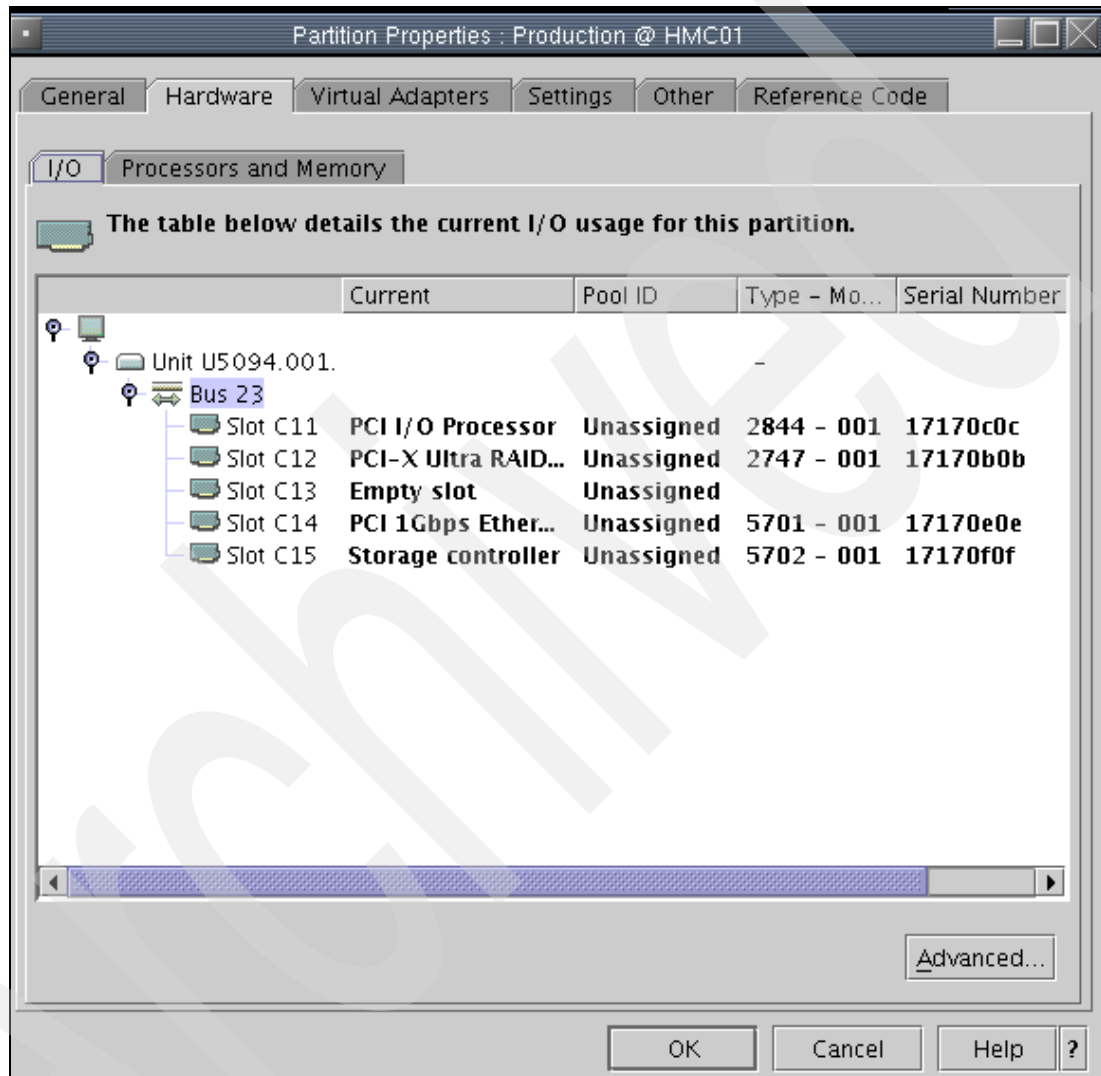


Figure 3-57 Partition properties tab - hardware IO view

Hardware IO - Advanced options button

By clicking the **Advanced Options** button, you will see the panel shown in Figure 3-58.

The IO pool options will allow the devices selected to participate in a switchable configuration.

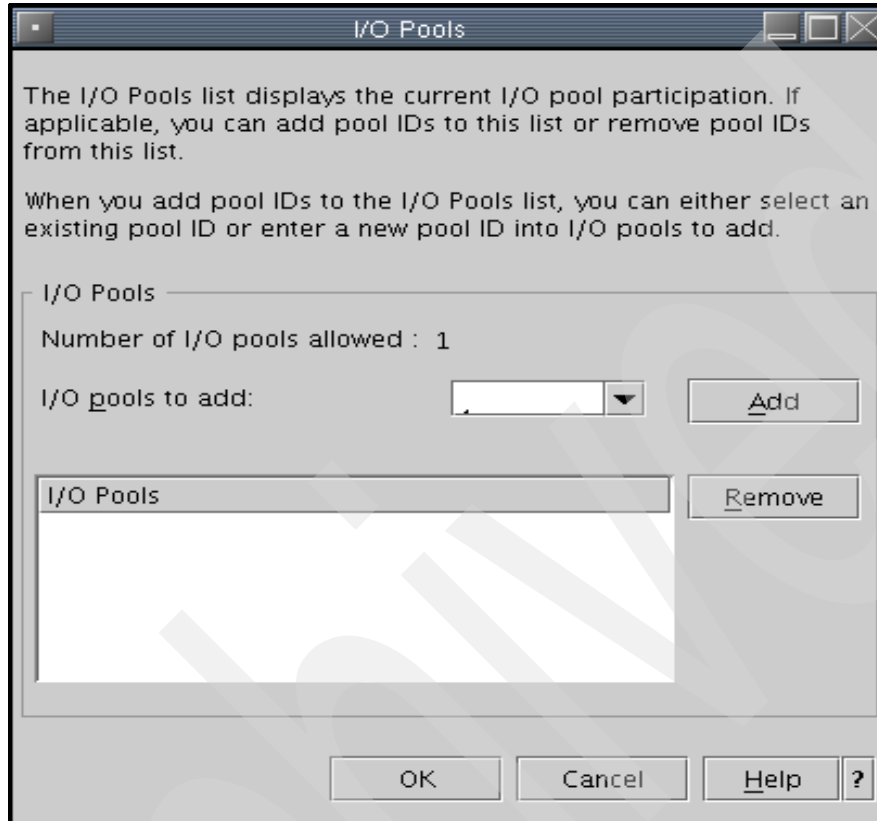


Figure 3-58 Partition properties - hardware IO, advanced

Partition properties - Processors and Memory

The panel displayed in Figure 3-59 shows a similar panel to the System properties, and shows the resources associated with the partition being viewed.

In the panel shown, we have created this partition as shared. The shared partition indicator is showing “enabled”.

The resources can only be viewed, they cannot be changed or deleted. You would need to use the DLPAR functions to add or reduce the processing or memory.

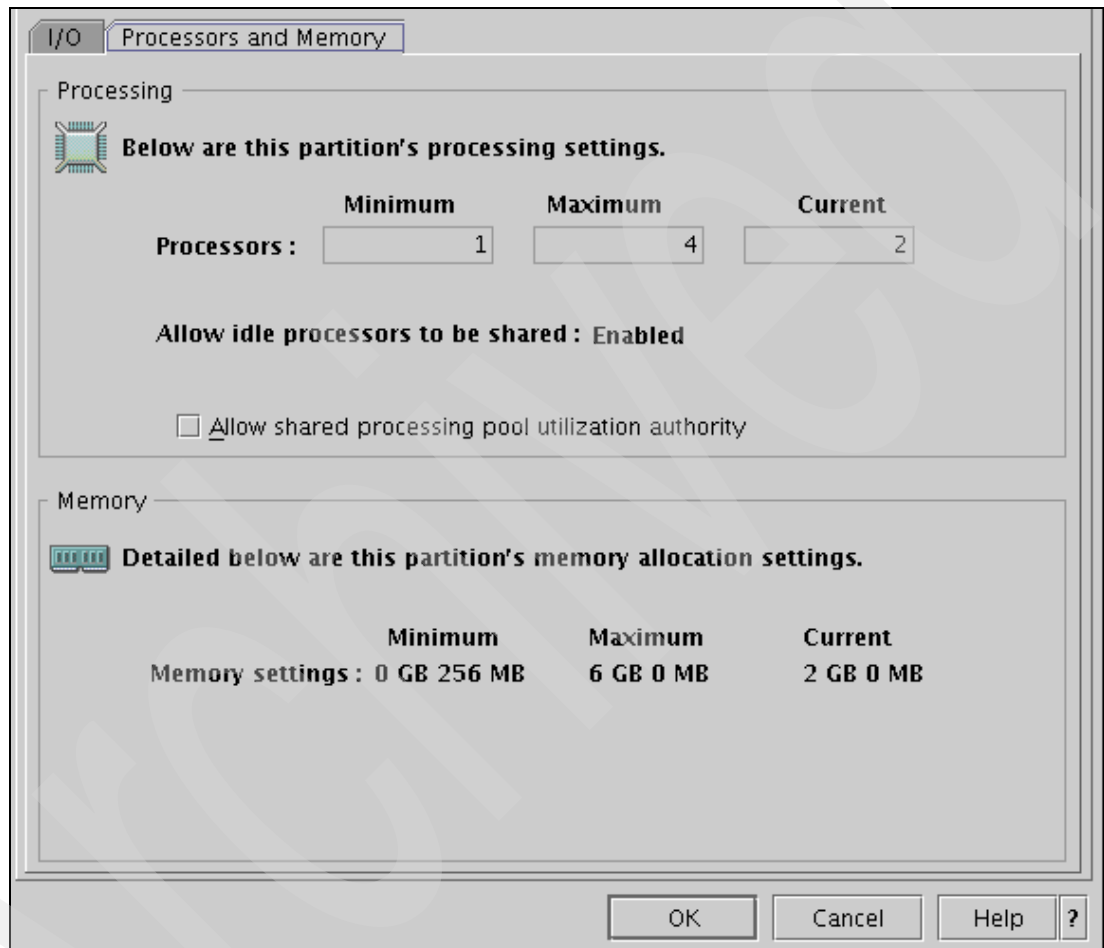


Figure 3-59 Partition properties - hardware processors and memory

Partition properties - virtual devices

There are three tabs within this panel, virtual ethernet, virtual serial, and virtual SCSI. The panel shown in Figure 3-60 is the general panel. Any adapter type configured will be shown.

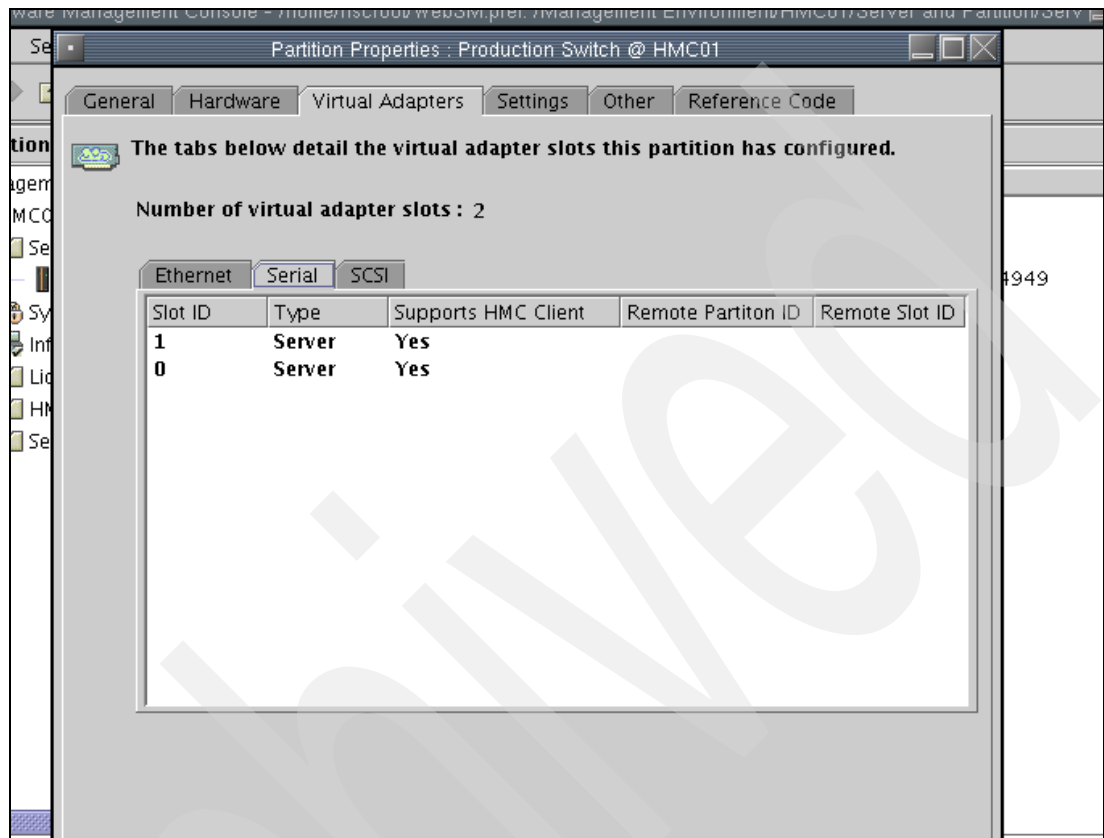


Figure 3-60 Virtual Serial display

If you select any of the radio buttons (ethernet, serial, or SCSI) you can add or remove adapters.

In Figure 3-61 you can also see that we have selected the first serial adapter and displayed its properties. You can change the access/usage of the virtual port.

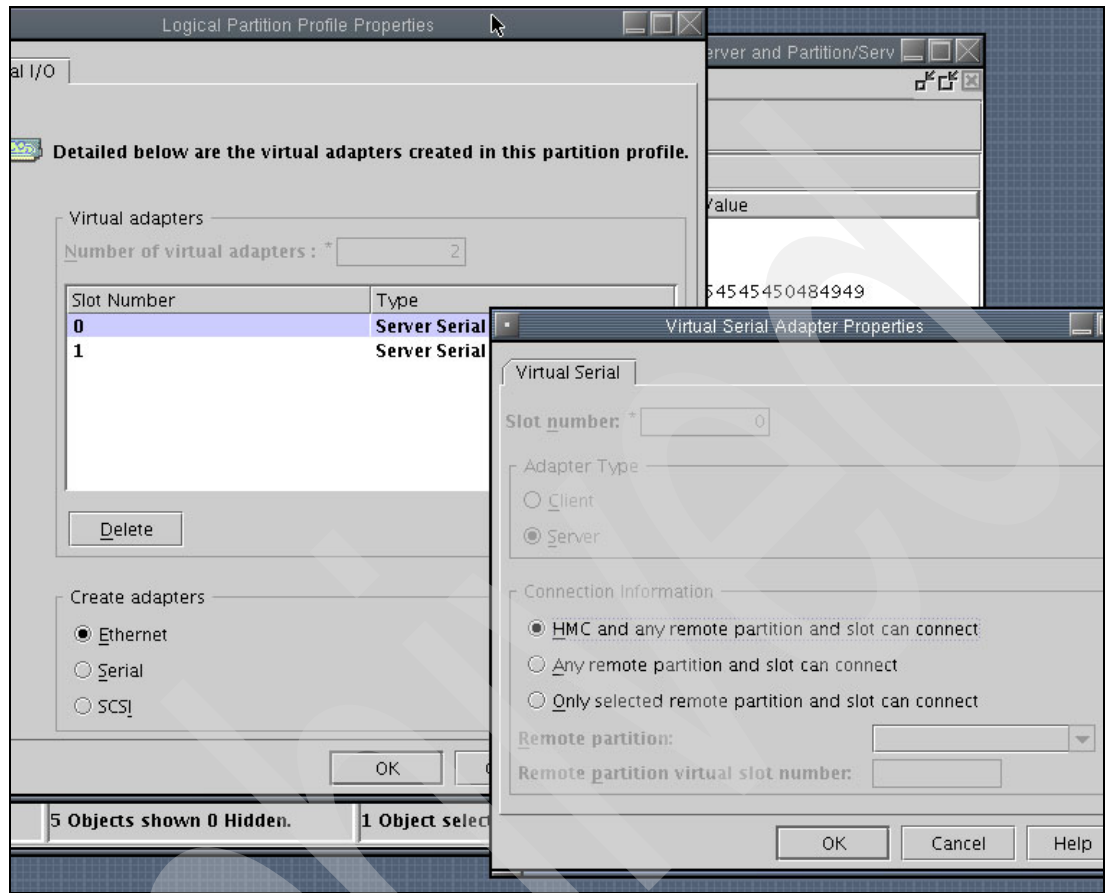


Figure 3-61 Virtual Adapter properties

In Figure 3-62 we have added an ethernet adapter by selecting the ethernet radio button and then clicking the **Create** button.

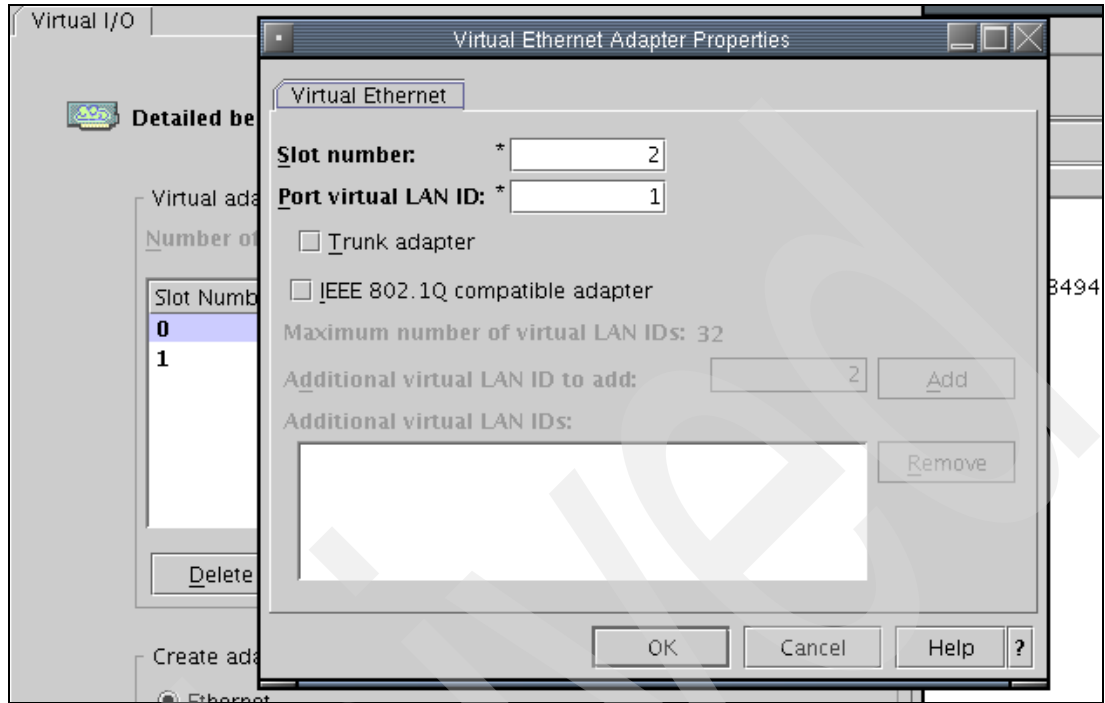


Figure 3-62 Add virtual ethernet adapter

Once you have clicked the **OK** button, you will be returned to the Virtual Adapters main panel, (Figure 3-63) and the new virtual ethernet adapter will have been added.

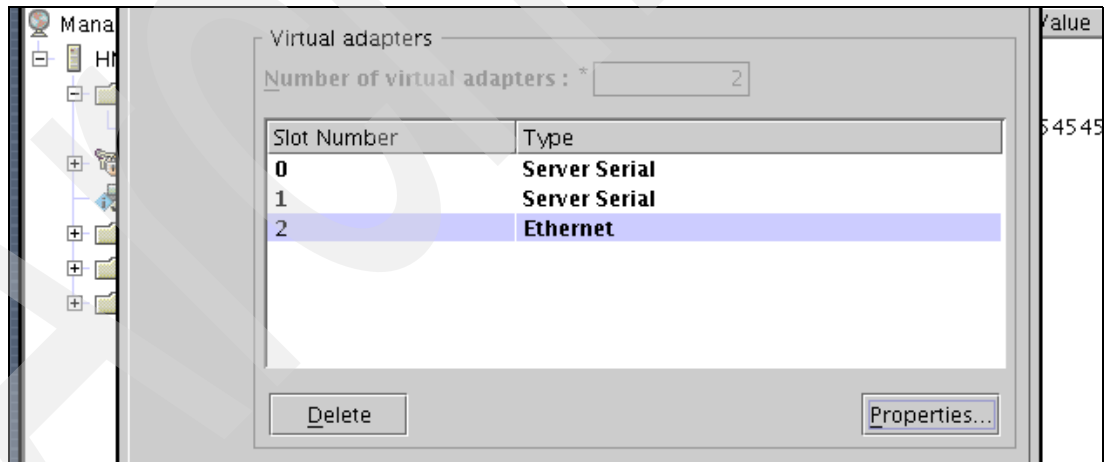


Figure 3-63 Virtual ethernet added

This function is equivalent to the iNav VLAN panel in the physical system properties or the green screen VLAN selection panel in the SST Configure Logical Partitions function. Neither of these functions is available on an System i5 system. These functions are still available on other servers that support i5/OS V5R3 and logical partitioning.

The panel shown in Figure 3-64 is how VLAN was configured on 8xx servers. Simply clicking the check boxes with the same number on two or more partitions creates a virtual LAN connection between the partitions.

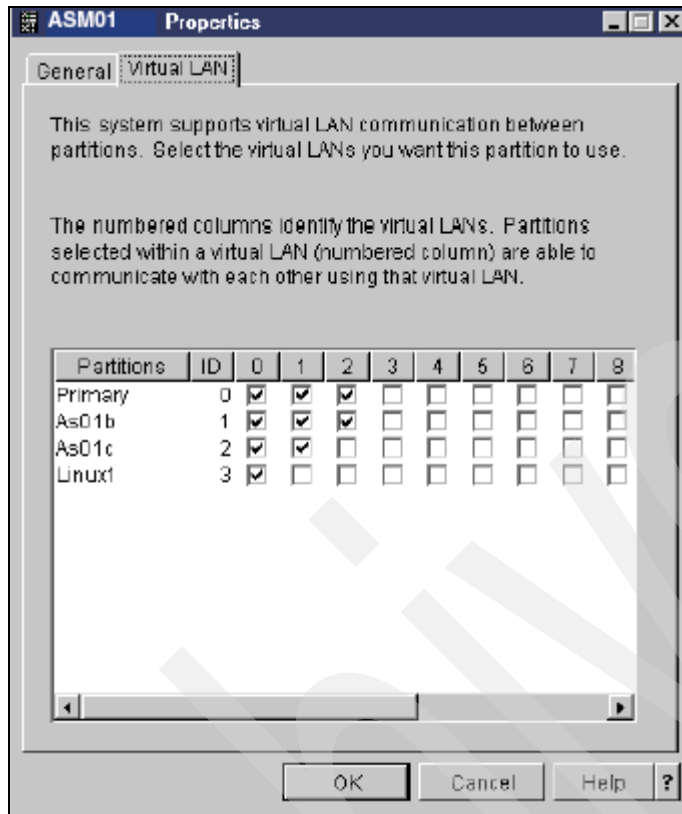


Figure 3-64 OS/400 Virtual Lan configuration on 8xx servers

If the LPAR migration tool is used to move partition configuration information from a 8xx server to a new System i5 system, the VLAN information will be migrated along with the partition information.

Partition properties - Settings

The Settings panel is where the IPL source, keylock, and automatic reboot functions can be set. This is the functionality that used to be performed by the primary partition on the earlier implementations of partitioning. This panel reinforces the fact that there is no longer a concept of a primary partition. The HMC performs these functions on the new System i5 system. See Figure 3-65.

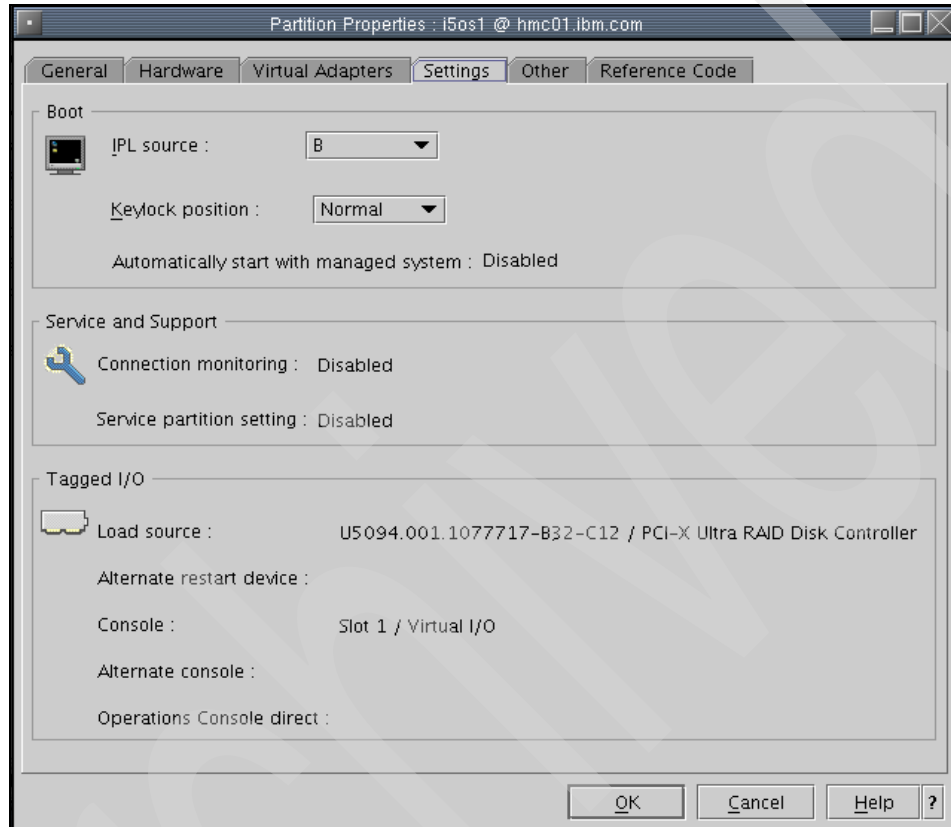


Figure 3-65 Partition properties Settings tab

Partition properties - Other

When we view the Misc. (Miscellaneous) tab, we are shown current information on Work Load Management groups, Power Controlling partitions, and Communications.

There are no Work Load Management groups (WLM) are currently defined. If they were defined, they could provide an alert function for possible system problems or out of band conditions. At present these will only alert a user ID. In the future they could make system decisions and dynamically change resources. See Figure 3-66.

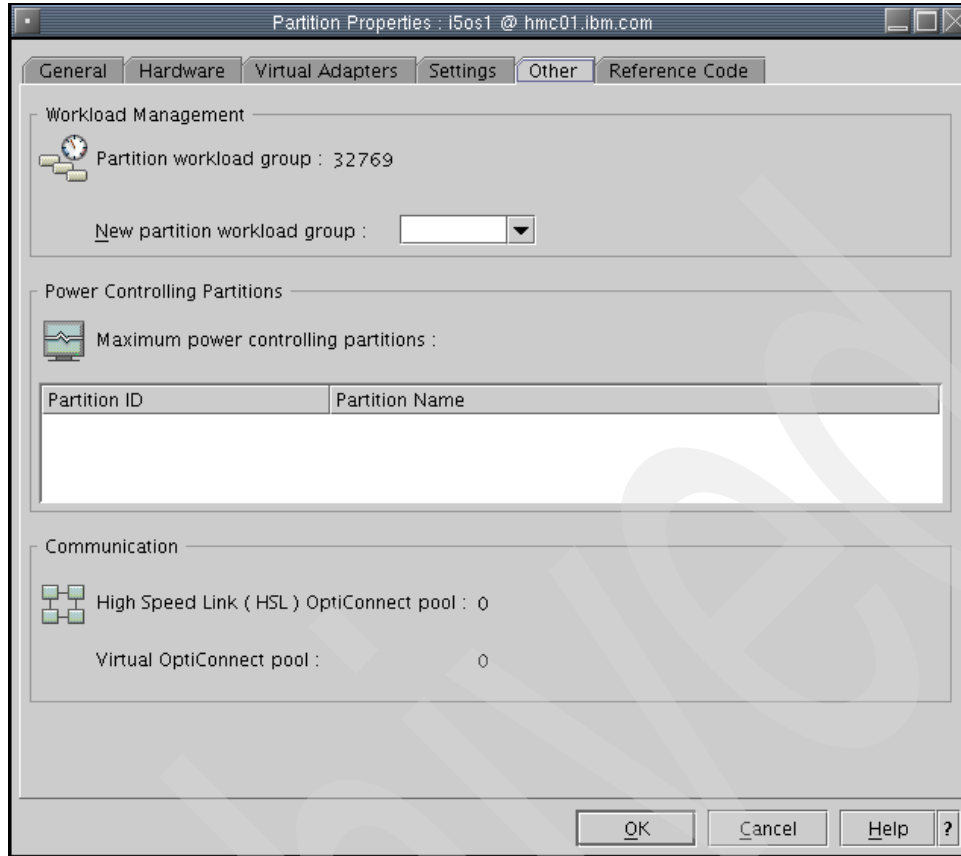


Figure 3-66 Partition properties miscellaneous tab

The HSL opticonnect and virtual opticonnect are the same as implemented in previous versions of partitioning; the power controlling partition is a new concept.

The power controlling partition is used in the context of i5/OS partitions hosting Linux. An i5/OS partition can host the console and the disk for the Linux partition and hence needs to control the power for the hosted Linux partition. This panel shows the Linux partitions for which this i5/OS partition is the power controlling partition.

3.5.4 HMC configuration

Click **HMC Management** → **HMC Configuration**.

This option allows you to perform the following configuration functions for the HMC:

- ▶ Customize console date and time.
- ▶ View console events.
- ▶ Customize network settings.
- ▶ Test network connectivity.
- ▶ Schedule operations.
- ▶ Enable or disable remote command execution.
- ▶ Enable or disable remote virtual terminal.
- ▶ Change current locale.

Customize date and time

You use this option to change the date and time and the time zone.

View console events

This is a log that allows you to view recent HMC activity. Each event is time stamped and the events that are logged include:

- ▶ When a partition is activated.
- ▶ When a system is powered on.
- ▶ When a user logs on.
- ▶ When a partition is shut down.

Customize network settings

Like the date and time option, the network settings would have been set up when the guided setup wizard was run. This option allows you to display or change those network settings if you need to, due to network changes. The network settings that can be displayed or changed include:

- ▶ Hostname and domain name for the HMC.
- ▶ Whether the HMC will run on a private or public network.
- ▶ Media speed.
- ▶ Whether the HMC will run as a DHCP server. This would be when the HMC is on a private network.
- ▶ Whether the HMC will run as a DHCP client. This could apply if the HMC was on a public network and it wasn't assigned a static address.
- ▶ Firewall (filter) settings to filter applications that can access the HMC. For example, you may want only certain clients to have webSM access to the HMC so they could be allowed access by their IP address.
- ▶ Whether the HMC is DNS enabled or not.
- ▶ Routing information, including a default gateway address.

Test network connectivity

This option allows the HMC to "ping" a TCP/IP address to check for connectivity between the HMC and another host on the network.

Schedule operations

This option allows you to schedule a backup of the HMC critical data. You can schedule the backup to occur once or on a repeated schedule. You have the option to repeat the backup hourly, daily, weekly, or monthly. The critical data that is backed up is the same as that described in 6.8, "Licensed Internal Code Updates" on page 211.

Note: Only the most recent backup image is stored at any time on the DVD-RAM.

Enable or disable remote command execution

This option allows you to enable or disable the ability to run remote commands to the HMC from a remote client using the SSH protocol. PuTTY would be an example of a remote client using the SSH protocol. The commands that can be executed are restricted.

Enable or disable remote virtual terminal

This option allows you to enable or disable the ability to run a remote virtual terminal on the HMC. An example would be using a 5250 client on a remote PC to access a OS400 partition. The remote client connects to the HMC by using the HMC's IP address. Multiple remote 5250 virtual terminals can be connected to the one partition. At the same time, care should be taken, as commands typed on one virtual terminal will display on the other terminals connected to the same partition. This function is useful for remote support.

Change current locale

This option allows you to change the locale (language) used by the HMC. The default is a US English locale. Some of the other optional languages available are:

- ▶ Spanish
- ▶ French
- ▶ Hungarian
- ▶ Japanese
- ▶ Italian
- ▶ Russian
- ▶ Slovak
- ▶ German

Note that in some languages, certain words and phrases may not translate. Locale can be changed at any time and only requires an HMC reboot.

3.6 Service applications

The service applications are used to service the managed systems from a central point. The applications are divided into four options, as follows:

- ▶ Problem determination
- ▶ Inventory Scout services
- ▶ Service focal point
- ▶ Remote service

3.6.1 Problem determination

The problem determination option allows the service support organization to view the HMC internal code and determine how to fix a software problem. Only service support personnel can access this option.

3.6.2 Inventory Scout services

The Inventory Scout is a tool that surveys managed systems for hardware and software information. Inventory Scout provides an automatic configuration mechanism and eliminates the need for you to manually reconfigure Inventory Scout Services. Depending on the levels of your HMC and partition software, you might be required to manually configure partitions that you create in order to perform Inventory Scout tasks.

The Inventory Scout collects the Vital Product Data (VPD) from the hardware resources in the managed system or systems that the HMC is managing. For the Inventory Scout to collect all the information accurately, all the managed system partitions must be active. The Inventory Scout collects information such as a resource type and serial number, its part number, its operational status, and other VPD, depending on the resource type.

This VPD is sent weekly to a database at IBM by a scheduled job. The initial collection of data will send all collected VPD, but any subsequent transmission will only send what has changed since the last transmission to IBM. The information sent to IBM is a valuable aid for IBM Remote Technical Support personnel when solving problems on the managed systems, as they will have an accurate profile of what the resources are on the system without having to connect to the system.

There are three options available under Inventory Scout services:

- ▶ Inventory Scout profile configuration
- ▶ Collect VPD information
- ▶ Restart inventory scout daemon

Inventory Scout profile configuration

Running this option runs a wizard which guides you through setting up the Inventory Scout profile. This should only be needed to run if the initial setup wizard for the HMC was not run or if a new AIX partition has been created since the initial setup wizard was run on the HMC.

- ▶ Collect VPD information — This collects the VPD to diskette if required.
- ▶ Restart inventory scout daemon — Option to restart the inventory scout daemon.

3.6.3 Service Focal Point

The Service Focal Point application is used to help the service representative diagnose and repair problems on partitioned systems. Service representatives use the HMC as the starting point for all service issues. The HMC groups various system management issues at one control point, allowing service representatives to use the Service Focal Point application to determine an appropriate service strategy.

Traditional service strategies become more complicated in a partitioned environment. Each partition runs on its own, unaware that other partitions exist on the same system. If one partition reports an error for a shared resource, such as a managed system power supply, other active partitions report the same error. The Service Focal Point application enables service representatives to avoid long lists of repetitive call-home information by recognizing that these errors repeat, and by filtering them into one error code.

The options available under Service Focal Point are as follows:

- ▶ Repair serviceable event
- ▶ Manage serviceable events
- ▶ Install/add/remove hardware
- ▶ Replace parts
- ▶ Service utilities

Repair serviceable event

This option allows the user or the service representative to view a serviceable event and then initiate a repair against that service event.

In the following paragraphs, we give an example of the steps taken to view an event and initiate a repair.

Click **Service Focal Point** → **Repair Serviceable Event** and select the managed system (in Figure 3-67 the system name is called “Unknown”).

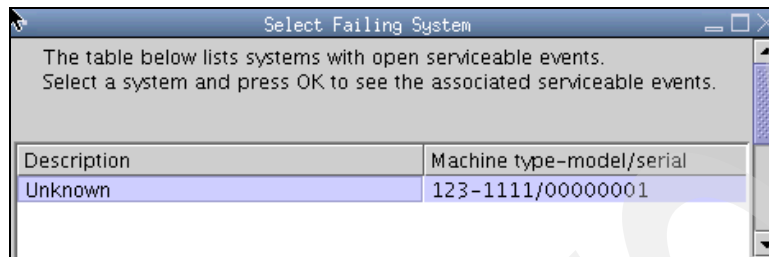


Figure 3-67 Example of selecting a managed system with a serviceable event

To view the serviceable event, highlight the managed system and click **OK**.

Figure 3-68 shows an example of a serviceable event.

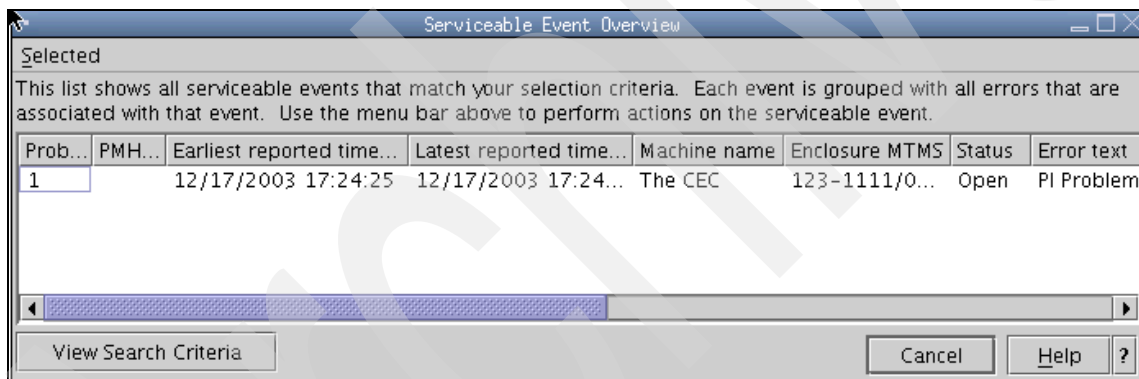


Figure 3-68 An example of a serviceable event

To look at the details of the serviceable event, highlight the serviceable event, then click **Selected** → **View details**.

The Serviceable Event Details panel opens, showing extended serviceable event information, including the following values:

- ▶ Status.
- ▶ Earliest original time stamp of any managed object.
- ▶ i5/OS or AIX error log ID (The Linux system error log does not place entries into Service Focal Point.)
- ▶ Should this error ever get called home?
- ▶ Error was called home.
- ▶ Pointer to extended error-data collection on the HMC.

The window's lower table displays all of the errors associated with the selected serviceable event. The information is shown in the following sequence:

- ▶ Failing device system name
- ▶ Failing device machine type/model/serial
- ▶ Error class
- ▶ Descriptive error text
- ▶ Details

To initiate a repair on the serviceable event, highlight the serviceable event, click **Selected** → **Repair**.

This allows you to select the search criteria to view the serviceable events that fit the search criteria:

Click **View Search Criteria**.

Figure 3-69 shows the possible search criteria that can be changed. The term MTMS stands for Machine Type Machine Serial. An example would be 9406-12345.

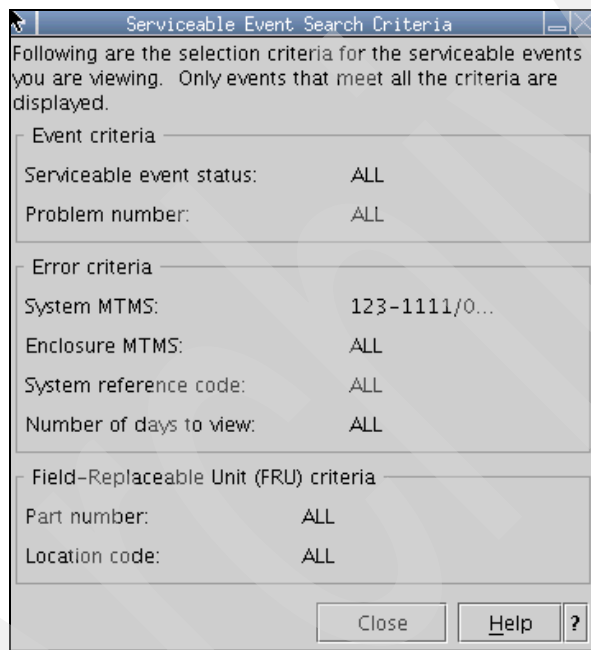


Figure 3-69 An example of a serviceable event search criteria panel

The selection fields allow you to narrow down which particular serviceable event you want to look at. If in doubt, leave the defaults, and the serviceable events will be listed in chronological order.

Manage Serviceable Event

This option is very similar to the last option (repair serviceable events).

Click **Manage Serviceable Event**.

The first panel that comes up is the same as Figure 3-69, which is the search criteria panel. This will display one or several serviceable events, depending on your search criteria.

Figure 3-70 shows the options that are available for a serviceable event.

Highlight a service event, and click **Selected**.

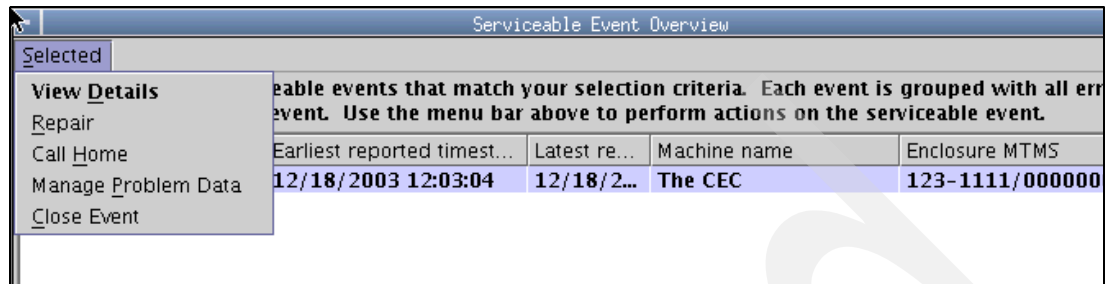


Figure 3-70 Options available to manage a serviceable event

The view option and repair option have been covered in the section on repair serviceable events. The three extra options are:

- ▶ Call home. You can force the HMC to report the serviceable event to IBM.
- ▶ Manage problem data. This option allows you to view, save to DVD, or call home to IBM, data associated with the serviceable event. The particular data is selectable from a drop down menu and includes dumps, logs, PTF listings, etc.
- ▶ Close event. This option allows you to close a serviceable event with comments and the name of the person closing the event.

Note: You must close a serviceable event after it has been serviced to ensure that if a similar error is reported later, it is called home. If an old problem remains open, the new similar problem is reported as a duplicate. Duplicate errors are neither reported nor called home to a service center. Close a serviceable event when the partition that reports the error is active. Closing the event causes the new status of the serviceable event to be correctly sent to the partition.

Install/add/remove/hardware

This option guides the user or service personnel through the steps to install, add, or remove hardware.

Exchange parts

This option guides the user or service personnel through the steps to replace a failed part.

Service utilities

There are a number of utilities available under this option. To select a utility:

Highlight the managed system, click **Selected**, and select one of the available utilities. The available utilities are as follows:

- ▶ Create serviceable event. This is a good option to report a problem that has not been created automatically. You write a description of the problem and click **Request service**, which will dial home to IBM and report the problem. There is also an option to test automatic problem reporting which will transmit a test problem to IBM to test the connection to IBM.
- ▶ Launch ASM menu. Use this option to access the service processor menus via a browser interface.

- ▶ **Manage Dumps.** Use this option to show any dumps associated with a service event. You can save them to DVD, send them to IBM or delete them.
- ▶ **Edit MTMS.** Use this option to modify the MTMS or the configuration ID of a selected enclosure.
- ▶ **System attention LED.** Use this option to look at the status of a system or partition system attention LED. You can select to turn off the attention LED.
- ▶ **Identify LED processing.** Use this option to turn on or off the identify LED for a selected enclosure or a selected enclosure's FRUs.
- ▶ **Lamp test.** Use this option to blink the LEDs in a selected system.

Another option within the service utilities is the ability to change the LAN "health check" settings. LAN "health check" is a function that allows for the LAN on which the HMC is running can be monitored. If the LAN that the HMC is on, fails, then the i5/OS partition that is designated as the service partition will wait a specified amount of time, then call home to IBM to report the problem. The settings for determining when a service call will be placed for the LAN failure can be set as follows. Click **Service utilities** → **Highlight the managed system** → **Actions** → **Surveillance settings**.

The settings that can be changed are:

- ▶ Number of disconnected minutes considered an outage.
- ▶ Number of connected minutes considered a recovery.
- ▶ Number of minutes between outages considered a new incident.

If the initial guided setup wizard for the HMC was run, then these settings would have already been set up.

3.6.4 Remote support

Remote support enables connectivity to IBM from the HMC and to the HMC from IBM. The Remote Support Facility (RSF) must to be enabled to allow Service Agent (SA) to connect to IBM to transmit the inventory of the managed system to IBM. Enabling remote support also allows for the reporting of problems to IBM via the HMC. Remote support must be enabled if IBM needs to connect to the HMC for remote servicing. This remote servicing is always initiated from the managed system end rather than the IBM end for security reasons.

The options available for remote support are as follows:

- ▶ **Customer information.** Used to enter the customer contact information such as address, phone numbers, contact person.
- ▶ **Outbound connectivity settings.** The information required to make a connection to IBM from the HMC for problem reporting and Service Agent inventory transmissions.
- ▶ **Inbound connectivity settings.** The information needed for IBM to connect to the HMC for remote service.
- ▶ **E-mail settings.** This option is used to set a notification by E-mail when the HMC reports a problem to IBM. The user defines what E-mail address will receive the notification.
- ▶ **Remote support requests.**
- ▶ **Remote connections.**



HMC Guided Setup

This chapter provides an overview of the Guided Setup function included in the Information Center component of the HMC. This chapter is divided into the following sections:

- ▶ Guided Setup planning and checklist
- ▶ User ids and authority
- ▶ HMC Networking setup
- ▶ HMC Service setup

4.1 HMC Guided Setup

The HMC Guided Setup wizard guides you through the main tasks needed to help you set up and tailor the many functions of the HMC. The wizard will launch automatically the first time the HMC is started. Using this wizard is the simplest way to configure your HMC.

Before using the Guided Setup wizard, you must understand the main concepts of HMC and decide which functions are relevant to your environment. The Guided Setup will enable you to configure the following functions:

- ▶ Set date and time.
- ▶ Change the passwords of the predefined user IDs for hscroot and root.
- ▶ Create new user IDs.
- ▶ Configure network settings.
- ▶ Customer contact information for support and services.
- ▶ Configure connectivity for service-related activities.
- ▶ Configure Service Focal Point.

Use the planning checklist below to help you prepare ahead of time.

4.1.1 HMC Guided Setup wizard checklist

Table 4-1 is a checklist of information that should be collected before the HMC installation and setup tasks are carried out.

Table 4-1 Guided Setup checklist

Task	Description	Comment
	HMC code level	Help pull down - About HMC
	Service processor code level	ASMi menu, configuration
	HMC language locale	HMC configuration
	Locale and time zone	Where HMC will reside
	HMC ids/passwords	hscroot and root
	HMC roles	Set up task authority
	HMC network configuration	Whether you want private or public network
	Domain, DNS, IP address range	Get network administrator involved
	Ethernet adapter configuration	Review installed adapters
	Service connection type	Dialup, VPN, or passthru
	Service provider information	IBM id, modem numbers
	HMC backup — media	Buy DVD-RAM disks
	HMC backup — networked	Locate suitable network drive

4.1.2 Using the Guided Setup wizard

This section walks you through an example of setting up an HMC via the Guided Setup wizard. Ensure that you have completed the HMC Guided Setup wizard checklist before continuing with the next section.

Important: If you cancel or exit the Guided Setup wizard at any time *before* you click the **Finish** button, all your inputs will be lost.

Once you have been through the Guided Setup wizard, you cannot “rerun” this function. Any adjustments to setup will need to be made through normal HMC functions.

Set up the language locale

Before you start the Guide Setup wizard, you should decide whether you want to use the wizard and/or the HMC in one of the 28 supported languages or the default of US English. When you change the locale for the HMC, the help text also changes to the new language selected.

Figure 4-1 shows how to change the language locale. Once you accept the new locale, you must reboot the HMC, but you do not have to shut down.

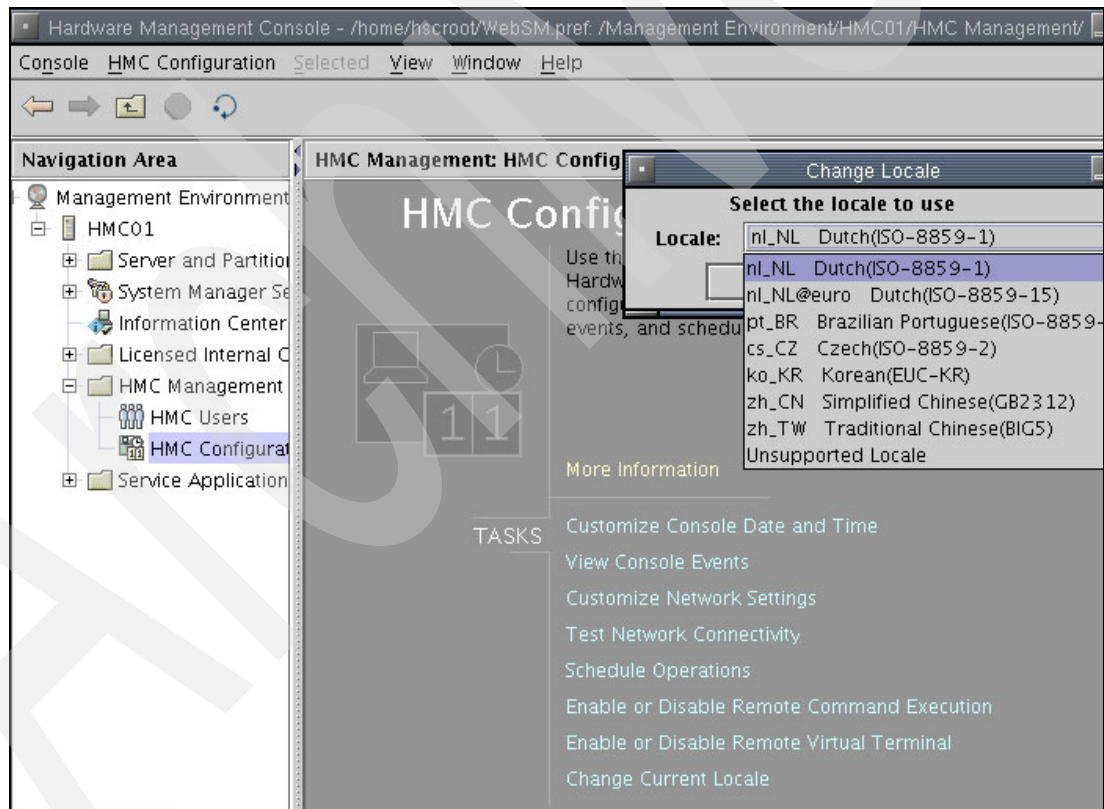


Figure 4-1 Change the language locale

Figure 4-2 shows the HMC with the first panel of the Guided Setup wizard in French.

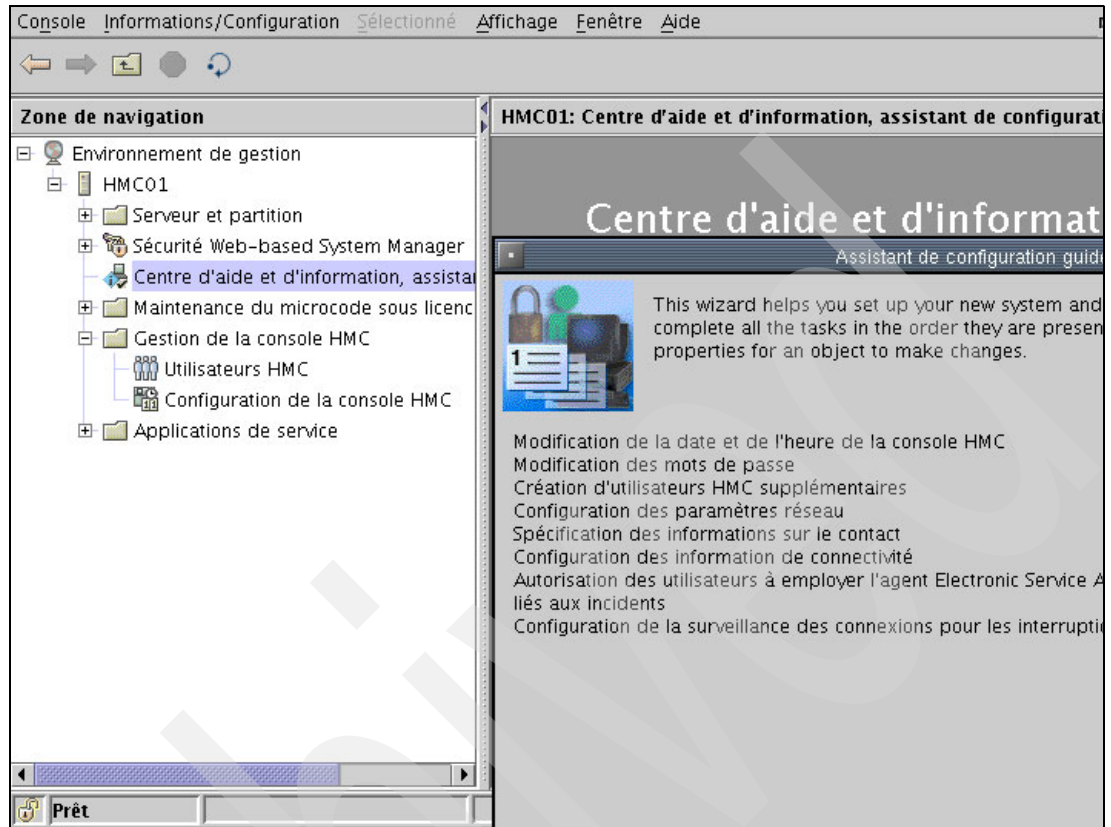


Figure 4-2 HMC Guided Setup with French locale

You can flip-flop between language locales at anytime, but you must reboot the HMC. If you use this function and cannot read the language, you must just remember where the change locale option is located.

Note: You can change the locale to Japanese, but you cannot enter information in DBCS. In some languages, not all words are translated.

Launching the Guided Setup wizard

The Guided Setup wizard should launch automatically when you first start the HMC. However, the wizard can also be launched from the HMC desktop by using the following steps:

1. Login in to the HMC with the default system administrator user ID *hscroot* and the password *abc123*.
2. In the Navigation area, click **Information Center and Setup Wizard**.
3. In the Contents area, click **Launch the Guided Setup Wizard** as shown in Figure 4-3.

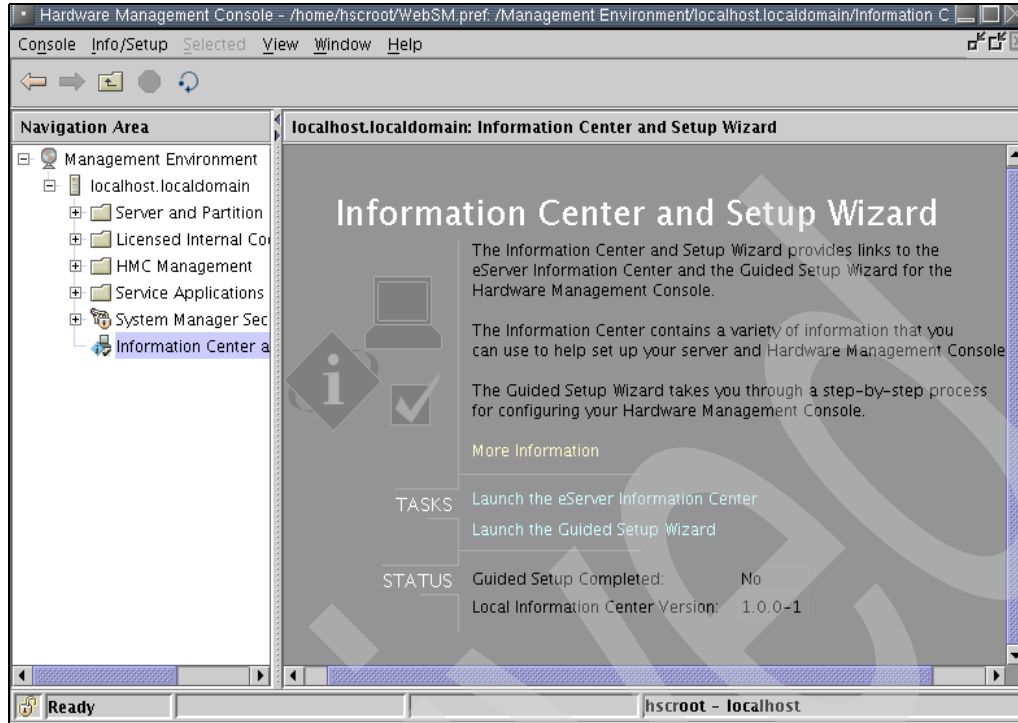


Figure 4-3 HMC Guided Setup

4. The Guided Setup Wizard Welcome page appears (Figure 4-4). Click **Next** to continue with the wizard.



Figure 4-4 HMC Guided setup welcome page

5. On the Guided Setup wizard - Change HMC Date and Time panel (Figure 4-5 on page 110), enter the correct date/time and time zone for your environment. This is typically the time zone of the server, assuming the HMC is local to the machine. For remote machines, you must decide which is the correct time zone for your environment.

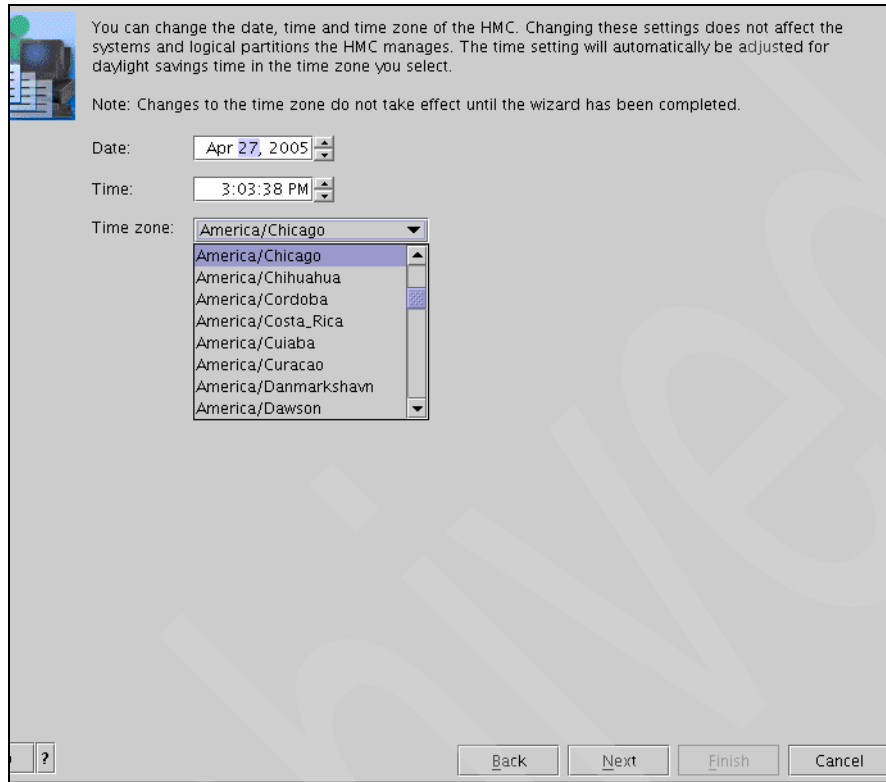


Figure 4-5 HMC Guided setup - Time and date setting

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

- The Guided Setup Wizard - Change hscroot Password panel is now displayed as shown in Figure 4-6. Enter the current *hscroot* password (normally this should be the default password of *abc123*) and then the new password you would like. The hscroot user ID is the i5/OS QSECOFR equivalent profile for the HMC, this user ID has full rights to all functions available on the HMC.

You need to change the predefined password for the default user ID of hscroot. The original password is published in the documentation. For sign-on security, it needs to be changed immediately.

1

Enter the new password.

User ID: hscroot

New password:

Re-type new password:

Role: hmcsuperadmin

Help ? Back Next Finish Cancel

Figure 4-6 Guided Setup wizard - Change hscroot password

You should change the hscroot id's default password of *abc123*, as most of the HMC documentation refers to this password and poses a security risk to your environment. Click **Next** to continue with the Guided Setup wizard.

7. The Change root Password panel is now displayed as shown in Figure 4-7. The root user ID is used by the authorized service provider to perform maintenance procedures and cannot be used to directly log in to the HMC. Enter the current root password (normally this should be the default password of *passwd* - where 0 is the number zero rather than the letter o. Enter the new password you would like for the root user ID.



Figure 4-7 Guided Setup wizard - Change root password

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

8. The Create additional HMC users panel is now shown (see Figure 4-8). You can now optionally create new HMC users at this stage. In our example, we decided to create a new hscoper user ID with a role of hmccoperator to allow our operation staff access to HMC and work with partitions. See 7.3, “HMC User Management” on page 238 for further information on creating users and their roles.

You can also skip this section and create users manually later on if you prefer.

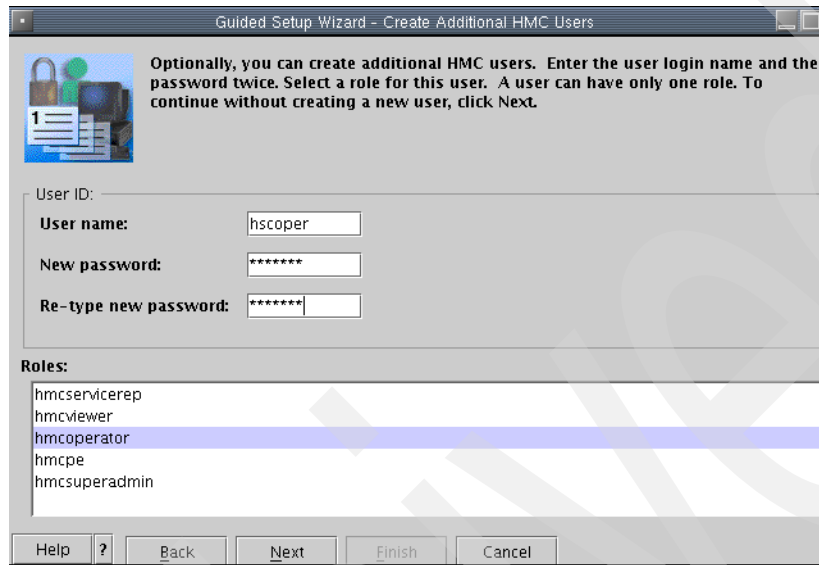


Figure 4-8 Guided Setup wizard - Create additional HMC Users

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

9. The Guided Setup wizard - Create additional HMC users panel will appear (see Figure 4-9).

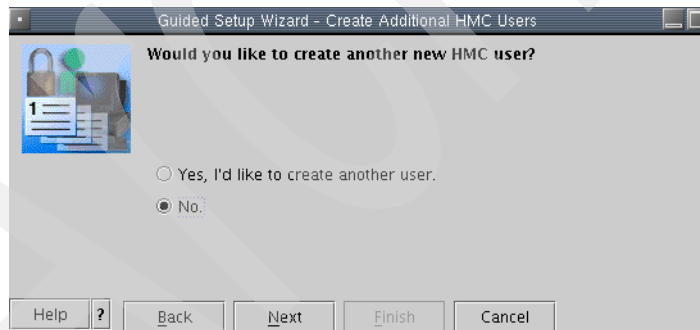


Figure 4-9 Guided Setup wizard - Create more additional HMC Users

If you need to create more additional users, then select the **Yes** radio button, otherwise select **No** and click **Next** to continue.

10. This completes the first part of the Guide Setup wizard. The Guided Setup - The Next Steps panel is displayed (see Figure 4-10).

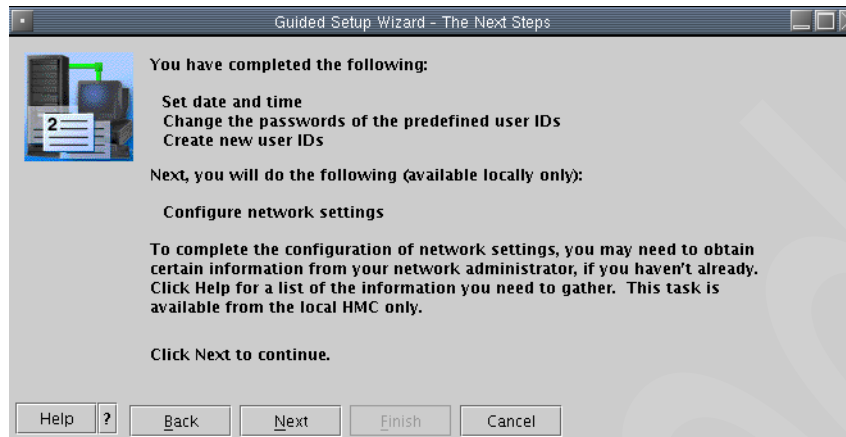


Figure 4-10 Guided Setup wizard - The next steps

The next section will configure the HMC network settings. You will need to have planned your network environment for HMC before continuing with these next tasks. You should use the values entered in the HMC Guided Setup wizard checklist.

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

11. The Guided Setup Wizard - Configure Network Settings panel appears as shown in Figure 4-11. In our example we show how to configure the HMC for both a private network and an open network. We use the first ethernet network card (eth0) in the HMC for the private network configuration, and the second ethernet card (eth1) to connect to our open network.

Enter a host name for your HMC, as well as your domain name and a description for the HMC.

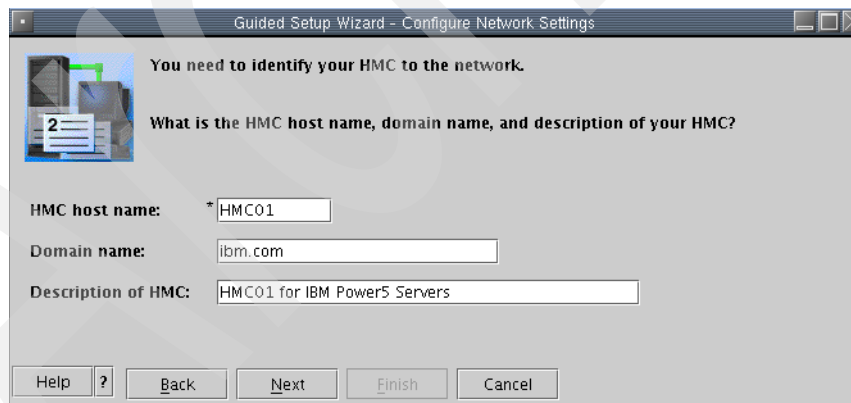


Figure 4-11 Guided Setup wizard - Configure network settings

Note that the ethernet ports are called eth0 and eth1, but the HMC ports are labeled HMC 1 and HMC 2.


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

12. The Guided Setup - Configure DNS panel is now shown (see Figure 4-12).

A DNS server is a distributed database for managing host names and their IP addresses. By adding a DNS server IP address to our HMC will 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, and the order that the addresses are entered will be 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**.



The screenshot shows a window titled "Configure DNS" with a help icon and a "2" in a blue box. The text reads: "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." Below this, it says: "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." The question "Do you want to use DNS?" is followed by two radio buttons: "Yes, I want to use DNS by defining DNS servers." (which is selected) and "No, I do not want to use DNS." Below the "Yes" option is a section titled "DNS server search order:" containing a list box with one entry: "DNS server address: 9.5.6.200". To the right of the list box are "Add" and "Remove" buttons. At the bottom of the window are "Help", "Back", "Next", "Finish", and "Cancel" buttons.

Figure 4-12 HMC Guided Setup wizard - Configure DNS

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

13. Now the Guided Setup - Specify Domain Suffixes panel is shown (see Figure 4-13).

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 will be 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 4-13 Guided Setup wizard - Specify Domain Suffixes

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

14. The Guided Setup Wizard - Configure Network Setting panel is then displayed (Figure 4-14). In our example we see two LAN adapters available (eth0 and eth1), however, you may only see one adapter in your HMC system.

We will configure eth0 for a private network and then will return to this panel to configure eth1 for an open network. The private network will be used to connect to our managed systems and other HMC's.

The second LAN adapter will be used to connect to our existing open network. This interface can be used by Web-based System Management Remote Client clients to remotely connect to the HMC and run HMC functions from client desktops such as Windows or Linux. We can also use this second interface to connect to IBM for service functions via VPN rather than using dial up connections.

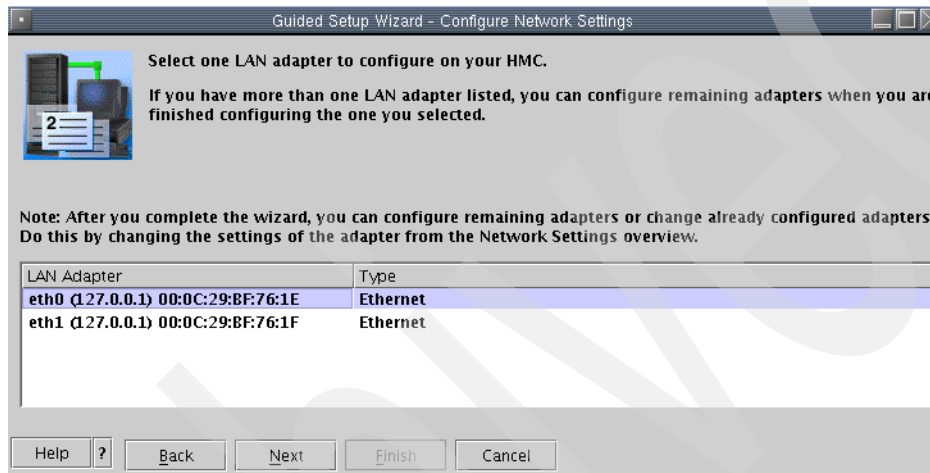


Figure 4-14 Guided Setup wizard - Configure Network Settings

We select LAN adapter eth0 and click **Next** to continue with the Guided Setup wizard.

15. The Guided Setup Wizard - Configure eth0 display is then shown (Figure 4-15). You can normally leave the LAN adapter speed at automatic detection for the initial setup. However, you can set each adapter to your actual network speed if you know this.

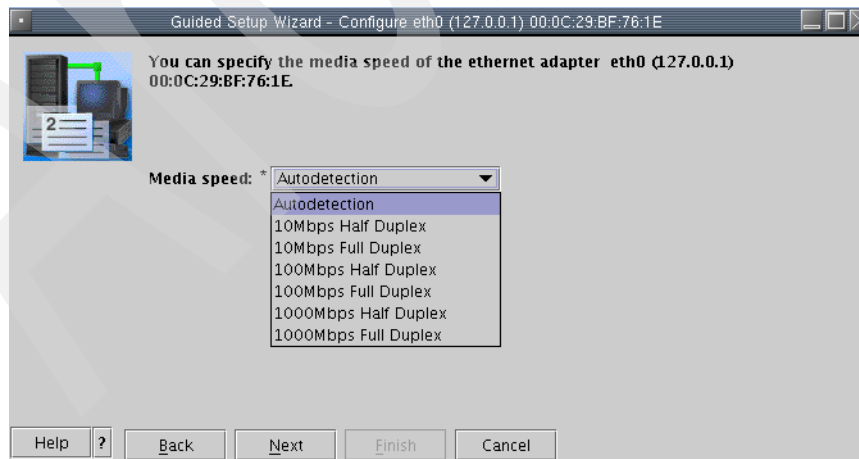


Figure 4-15 Guided Setup Wizard - Configure LAN adapter eth0

We select **Autodetection** and click **Next** to continue with the Guided Setup wizard.

16. The Guided Setup Wizard - Configure eth0 panel is now shown (Figure 4-16). As previously mentioned, we are setting the first LAN adapter to be our link to our private network of HMCs and managed systems.




Figure 4-16 Guided Setup wizard - Configure eth0

We select the Private network radio button and click **Next** to continue.

17. Now the Guided Setup Wizard - Configure eth0 panel appears (Figure 4-17). As this is the first HMC on our private network, we have to define the HMC as a DHCP server.

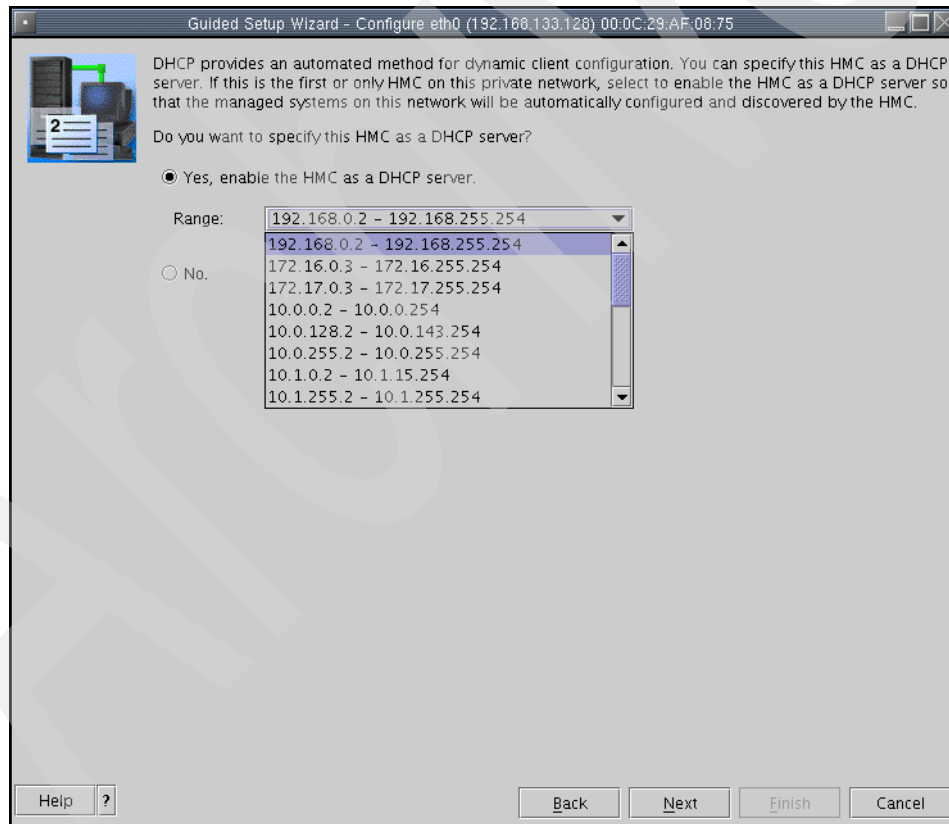


Figure 4-17 Guided Setup wizard - Configure eth0 DHCP

The HMC provides DHCP services to all clients in a private network. These clients will be our managed systems and other HMC's. You can configure the HMC to select one of several different IP address ranges to use for this DHCP service, 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 a choice of standard nonroutable IP address ranges that will be assigned to its clients. The ranges we can select from are:

- 192.168.0.2 - 192.168.255.254
- 172.17.0.3 - 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 - 10.0.143.254
- 10.1.255.2 - 10.0.255.254
- 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
- 9.6.24.2 - 9.6.24.254
- 9.6.25.2 - 9.6.25.254

The 9.6.24.2 - 9.6.24.254 and 9.6.25.2 - 9.6.25.254 ranges are two special range of IP address that can be used to avoid conflicts in cases where the HMC is attached an open network which is already using nonroutable addresses.

The HMC LAN adapter will be assigned the first IP address out of the range selected. In our example, we select the 192.168.0.2 - 192.168.255.254 range, so our HMC will be given IP address 192.168.0.2. Any other client (HMC or managed system) will also be given an address from this range.

The DHCP server in the HMC uses automatic allocation, which means that each managed system will be reassigned exactly 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 will reassign each client with same IP address as before. When a managed system starts, it will try to contact the DHCP service to obtain its IP address. If the managed system is unable to contact the HMC DHCP service then, the managed system will use its last given IP address.

We select the range 192.168.0.2 - 192.168.255.254 and click **Next** to continue.

18. The Guide Setup Wizard - configure eth0 panel is now shown (Figure 4-18). We can specify that one of our LAN adapters can act as a gateway device to our open network (if required).



Figure 4-18 Guided Setup wizard - Default gateway

In our configuration LAN adapter, eth1 will be our open network device, so we will set that card to be our default gateway device later on.

In our example eth0 is the private network LAN adapter, so we can just click **Next** to continue.

19. The Guided Setup Wizard - Configure Network Settings panel is now displayed (Figure 4-19). This completes the network configuration of the private network interface eth0. We can now proceed with the configuration of the second network card (eth1) for our open network.

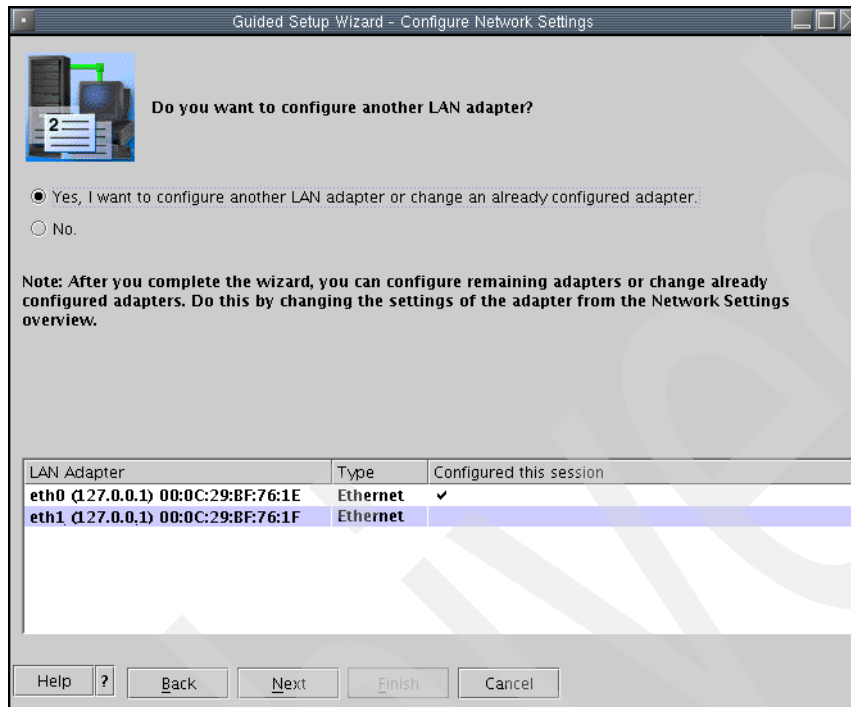


Figure 4-19 Guided Setup wizard - Configure Network Settings

Select the **Yes** radio button if it is not already flagged. The second ethernet card should be highlighted in grey (as in Figure 4-19). Click **Next** to continue with the Guided Setup wizard.

20. The Guided Setup Wizard - Configure eth1 display is then shown (Figure 4-20). As before, we can leave the LAN adapter speed at automatic detection for the initial setup. However, you can set each adapter to your actual network speed if you know this.

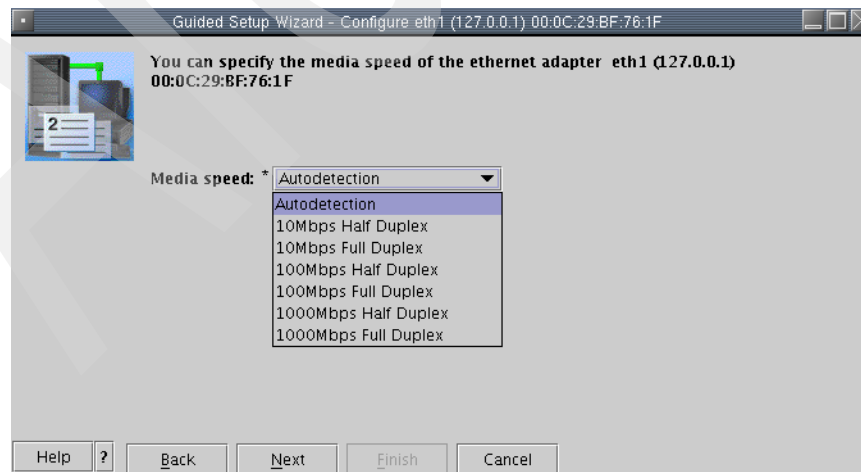


Figure 4-20 Guided Setup - Media speed for eth1

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

21. The Guided Setup Wizard - Configure eth1 panel is now shown (Figure 4-21 on page 122). This time we select the Open network radio button and click **Next** to continue.

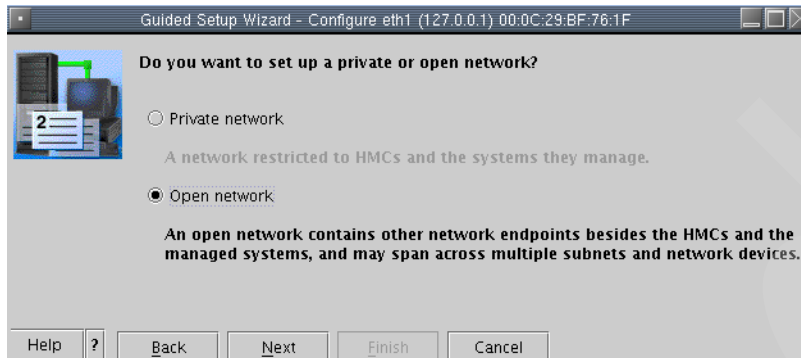


Figure 4-21 Guided Setup - Configure eth1 open network

22. The Guided Setup Wizard - Configure eth1 panel is now shown (Figure 4-22). You can configure the eth1 interface to use a fixed IP address or obtain one automatically from your open network DHCP server.

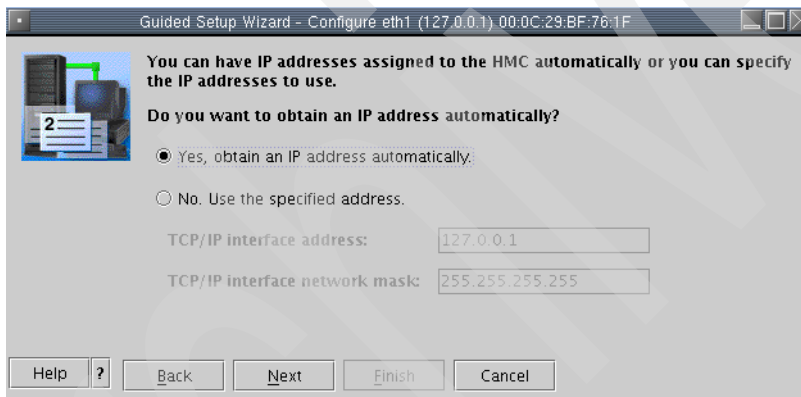


Figure 4-22 Guided Setup - Configure eth1 IP address

In our example, we select the **Yes** radio button to allow the eth1 interface to automatically obtain an IP address from our open network DHCP server. Click **Next** to continue with the wizard.

23. The Guide Setup Wizard - Configure eth1 panel is now shown (Figure 4-23). We can specify that the eth1 LAN adapter act as a gateway device to our open network. Enter your gateway IP address in the Gateway address field. If you don't know your gateway address, then you should contact your network administrator.




Figure 4-23 Guided Setup - eth1 gateway selection

In our example we enter our gateway address of 9.5.6.1. and click **Next** to continue.

24. The Guided Setup Wizard - Configure HMC Firewall for eth1 panel is now displayed (Figure 4-24). Usually there is a firewall that controls outside access to your company's network. As the HMC is connected to the open network we can also restrict remote access to this device by using the HMC built in firewall. There are various applications that run on the HMC, which can be secured in order to protect the HMC from unauthorized remote access.



Figure 4-24 Guided Setup - eth1 firewall

Select the **Yes** radio button to configure the HMC firewall settings and click **Next** to continue.

25. The Guided Setup Wizard - Configure HMC firewall panel appears next (see Figure 4-25 on page 124). In the top panel (current applications) are listed all the available applications that are on the HMC. In the bottom pane (Applications allowed through firewall) are all the applications available to the open network through the HMC firewall. You can decide to remove applications completely from the firewall by selecting the relevant application from the bottom panel and clicking the **Remove** button.

Applications such a WebSM.name, which allows users to access the HMC by installing a remote client on their desktop PC are by default blocked. You can allow applications to pass through the firewall by selecting them from the top pane and clicking **Allow incoming** or **Allow incoming by IP address**. The function, Allow incoming, will enable all remote clients access to the selected application. The other option, Allow incoming by IP address, will only authorize specific remote client's IP addresses to have access to that application.

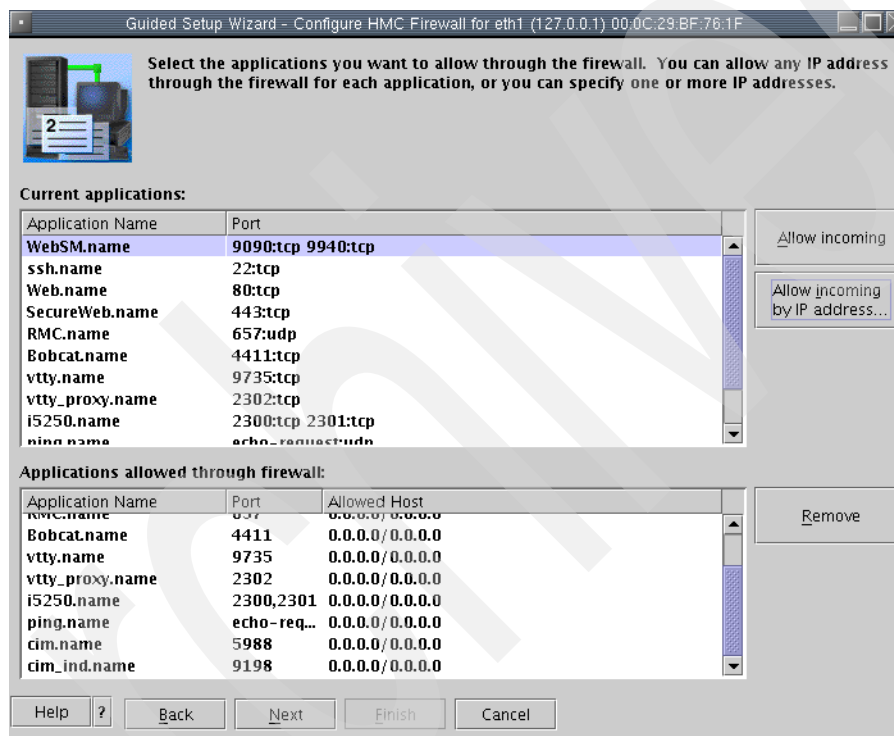


Figure 4-25 Guided Setup - eth1 firewall websm

In our example we select the WebSM.name application and click **Allow incoming by IP address**.

The Configure HMC firewall for eth1 panel is then displayed (Figure 4-26). You can add clients IP addresses by entering an IP address and network mask in the fields provided and clicking the **Add** button. You can repeat this process for other IP addresses, as each application can have multiple authorized IP address entries. You can also remove any unwanted clients by selecting the relevant IP address and clicking the **Remove** button.

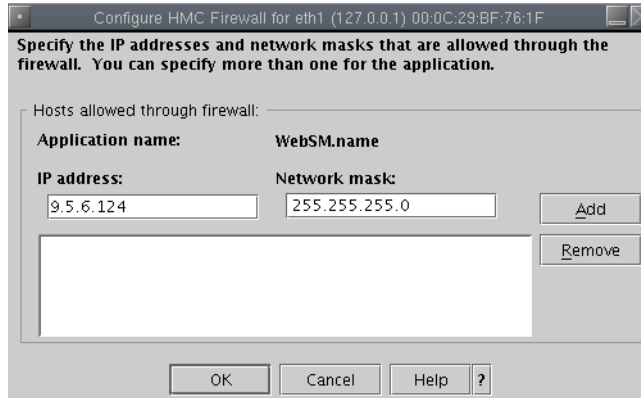


Figure 4-26 Guided Setup - eth1 firewall by IP address

In our example we enter the remote IP address 9.5.6.124 and mask 255.255.255.0 and click **Add** and then click **OK**. When we return to the HMC firewall configuration panel, we click **Next** to continue with the Guided Setup wizard.

26. The Guided Setup Wizard - Configure Network Settings panel is shown (Figure 4-27). If you have more network adapters available you can configure them now by selecting the relevant adapter and selecting the **Yes** radio button.




Figure 4-27 Guided Setup - Final network configuration panel

As both our network adapters have now been configured, we select the **No** radio button and click **Next** to continue.

Note: It may take a few minutes before the next panel appears.

27. The Guided Setup - The Next Steps display is shown (Figure 4-28). This completes the network configuration section of the Guide Setup wizard. We now continue with the next part of the wizard, which enables the service and support functions within the HMC.

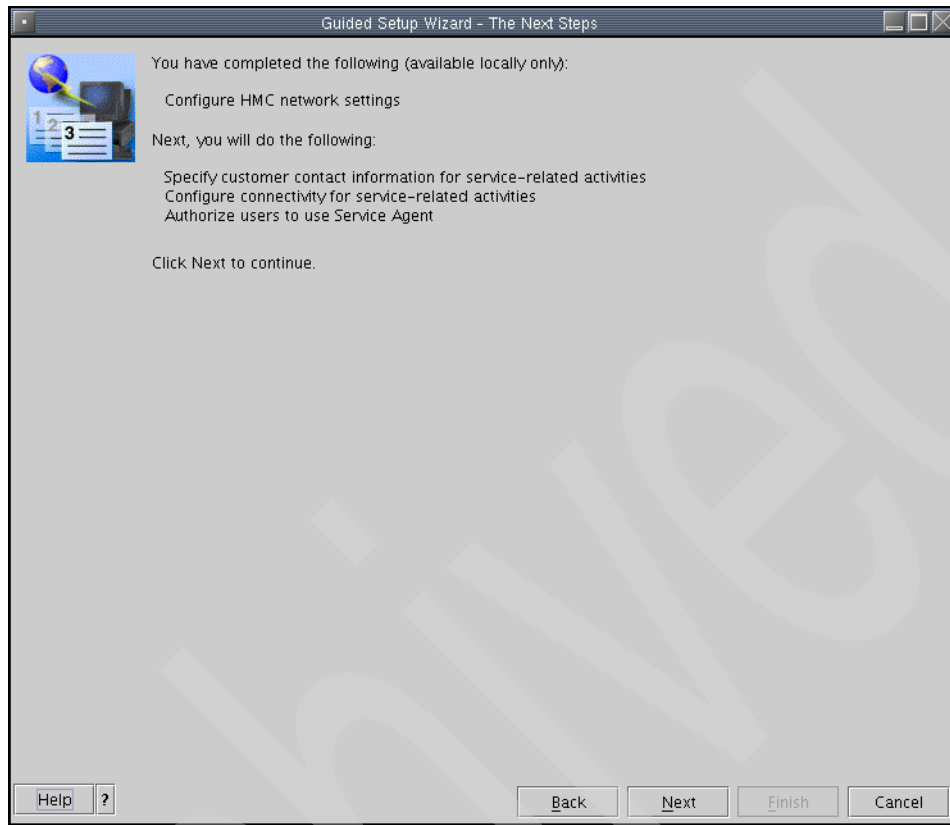


Figure 4-28 Guided Setup - End of network configuration

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

28. The Guide Setup Wizard - Specify Contact Information panel is presented (see Figure 4-29). This is the first of three panels which contain the contact details for your company (this information will probably be similar to the WRKCNTINF information stored in OS/400 if you have previous iSeries systems. The information entered here is used by IBM when dealing with problems electronically reported (calling home), as well as software updates. You should enter valid contact information for your own location. The fields marked with '*' are mandatory and must be completed.

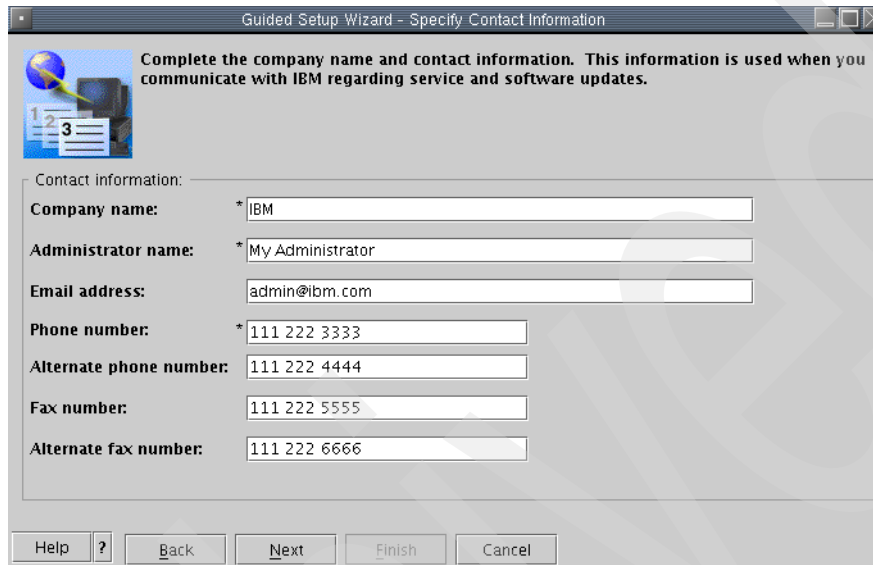


Figure 4-29 Guided Setup - Specify contact information

In our example we enter our location details and click **Next** to continue.

29. The second Contact Information panel is now displayed (see Figure 4-30). You should enter your contact address information on this panel. Again, you must complete the mandatory fields (*). Use the drop down menus to select your Country/Region and State/Province settings.

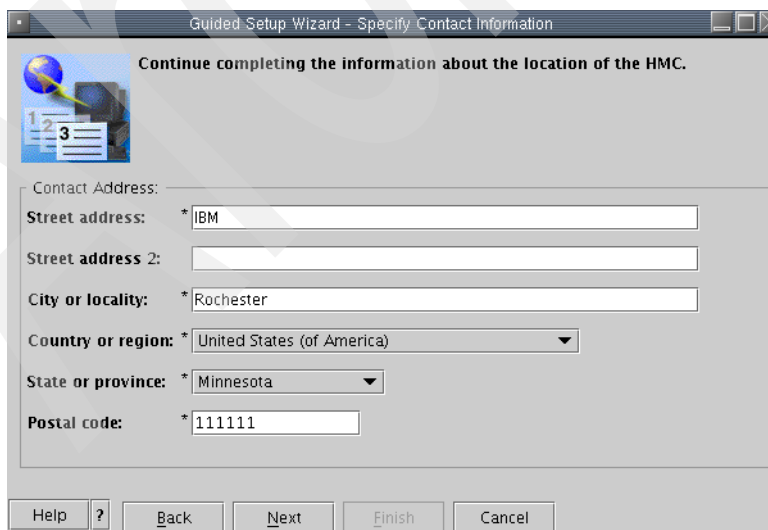


Figure 4-30 Guided Setup - Contact information part 2

For our example we enter our contact address details and click **Next** to continue.

30. The last panel for the Contact Information is now shown (Figure 4-31). You should enter the location details of this HMC here. If the location address is the same as the contact address used in the previous step, then click **Use the administrator mailing address**. Otherwise fill in the correct HMC location address details.

If this HMC will use the supplied modem to connect to IBM (call home) for service and support, then enter the modem telephone number in the Remote Support Information panel at the bottom of the panel.

Guided Setup Wizard - Specify Contact Information

Continue completing the information about the location of the HMC.

Location of the HMC:

Use the administrator mailing address

Street address: *

Street address 2:

City or locality: *

Country or region: * Afghanistan

State or province: *

Postal code: *

Remote Support Information

Modem number: 111 222 9999

Help ? Back Next Finish Cancel

Figure 4-31 Guided Setup - Modem Number

In our example, we used the same address for both contact and HMC, so we selected the Use the administrator mailing address. We also entered our Remote Support modem number in the Modem number field at the bottom of the panel.

This completes the contact information part of the HMC Guided Setup wizard service and support. Click **Next** to continue.

31. The Guided Setup Wizard - Configure Connectivity to Your Service Provider panel is now displayed (Figure 4-32). You can select by which communications method you wish to connect to IBM (call home) for service and support related functions. There are four service applications available on the HMC:

- **Electronic Service Agent™** - Monitors your managed systems for problems and if enabled, reports them electronically to IBM.
- **Inventory Scout Services** - Reports hardware and software information from your managed systems to IBM.
- **Remote Support Facility** - Enables the HMC to call out to IBM for problem reporting as well as enabling remote access to your managed systems (if enabled by customer).
- **Service Focal Point** - Collects system management issues in one central point for all your partitions.

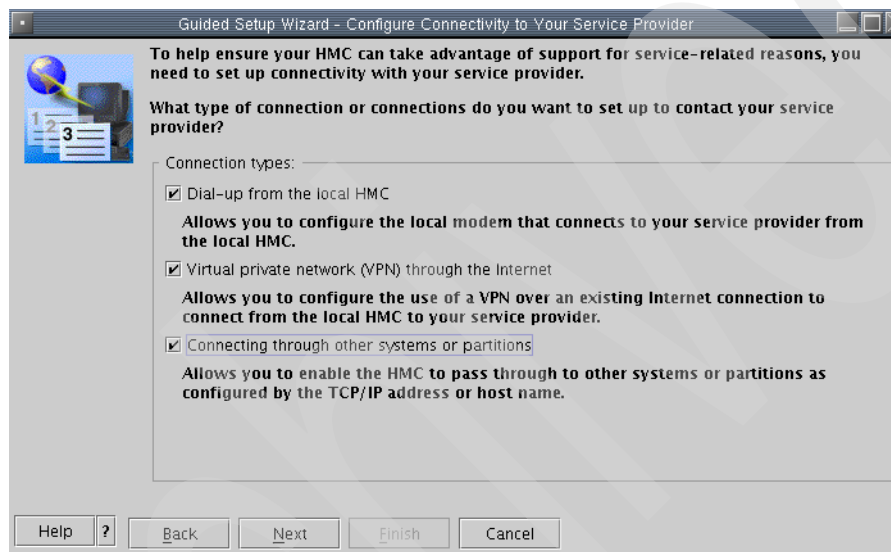


Figure 4-32 Guided Setup wizard - Configure Connectivity to Your Service Provider

You can select which connectivity method you wish to use when communicating electronically with IBM. There are three options available to you:

- **Dial-up from the local HMC** - This option will use the IBM supplied modem with the HMC to dial in to the IBM support network. You may choose this option if your HMC doesn't have a high speed Internet connection through an open network or has only been configured in a private network.
- **Virtual private network (VPN) through the Internet** - This option will use a high speed Internet connection to connect to the IBM support network. This is the fastest option available on the HMC; however, your company may restrict this type of connection.
- **Connecting through other systems or partitions** - This option sends information through another system in your network that can connect to IBM. The pass-through system could be another HMC or a partition running i5/OS V5R3.

In this example configuration we select all three connectivity options for demonstration purposes only. Normally you would only select the options valid for your environment.

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

32. The Agreement for Service Programs panel is now shown (see Figure 4-33). Read the agreement details carefully and click **Accept** or **Decline**.

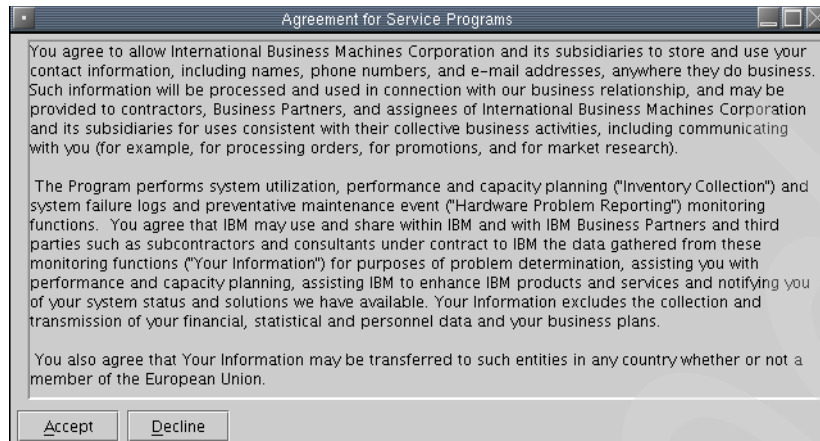


Figure 4-33 Guided Setup - Agreement for Service Programs

In our example configuration we click **Accept** to accept the terms and conditions of the IBM Agreement for Service Programs. We then return to the previous panel.

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

33. The Guided Setup Wizard - Configure Dial-up from the Local HMC panel is now displayed (see Figure 4-34).

In the modem panel, click the **Modem Configuration** button to set the modem parameters. You can set the Dial Type (Tone/Pulse), Wait for dial tone, Enable speaker, and the Dial prefix values (see the Modem Configuration window in Figure 4-34).

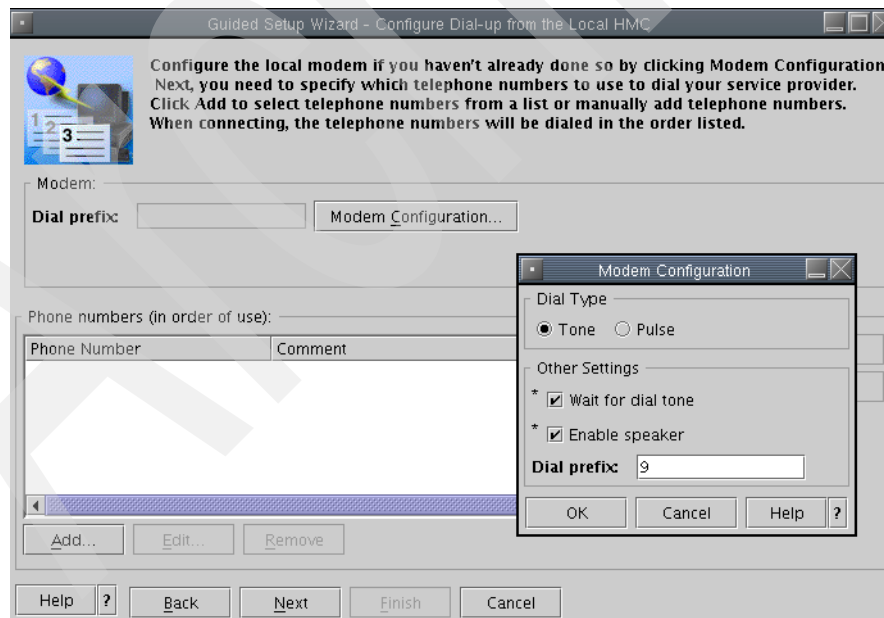


Figure 4-34 Guided Setup wizard - Dial-up connection configuration

In our example we enter **9** in the Dial prefix field and click **OK**. We then click the **Add** button in the Phone numbers (in order of use) panel to add the IBM support service phone number.

34. The Add Phone Number window is launched (see Figure 4-35). Use the drop down menus to select your Country/region and then your State/province.

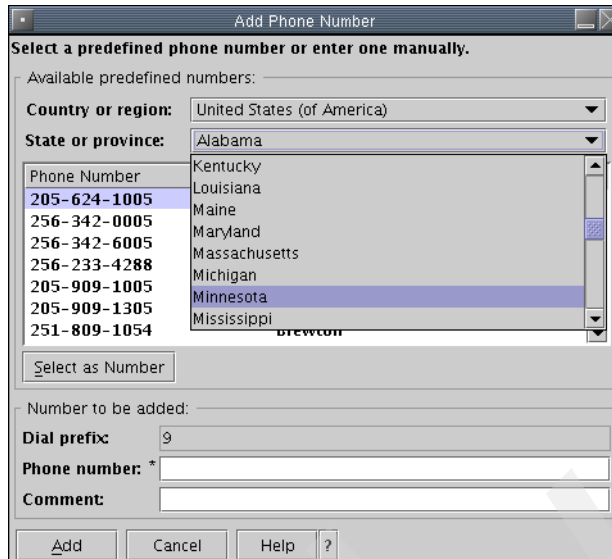


Figure 4-35 Guided Setup wizard - Add Phone Number

For our example we select **United States (of America)** for our Country/region and **Minnesota** for our State/province. You should select the relevant values for your location.

After you have selected your Country/region and State/province, a list of available IBM support service numbers are listed. You should select the phone number nearest to your location and click the **Select as Number** button. The phone number will then be populated in the Phone number field at the bottom of the panel (see Figure 4-36). You can also manually add phone numbers if you know your IBM support service number.

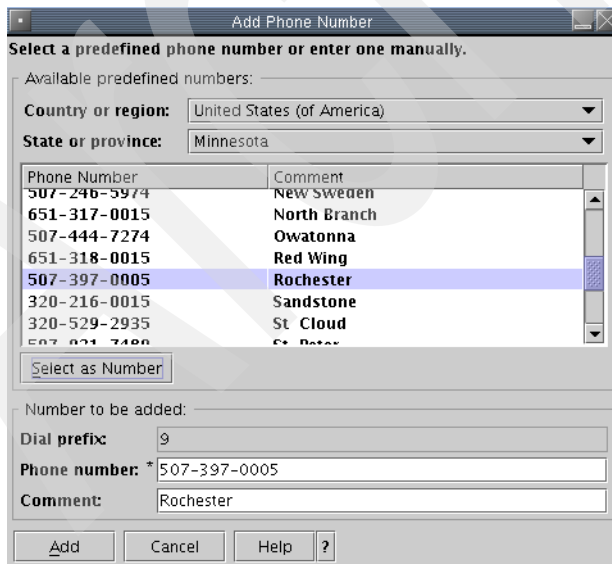


Figure 4-36 Guided Setup wizard - Select number

In our example we select Rochester and click the **Select as Number** button followed by the **Add** button.

35. We return to the Guided Setup Wizard - Configure Dial-up from the Local HMC panel (see Figure 4-37). You can add additional phone numbers by repeating the same procedure again and selecting a different number.

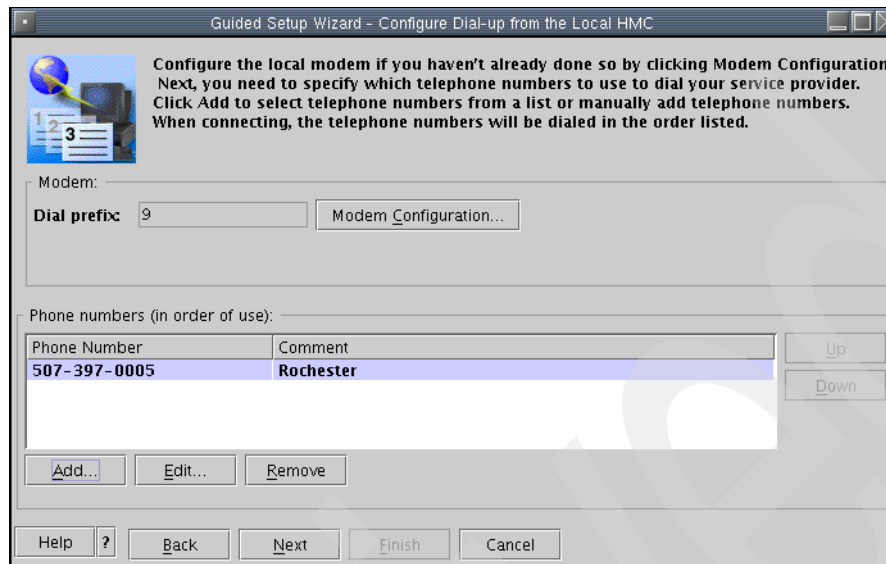


Figure 4-37 Guided Setup wizard - Dial-up configuration

This finishes our configuration for the Dial-up connection for the HMC. We click **Next** to continue.

36. The Guided Setup wizard - Use VPN using an Existing Internet Connection panel is displayed (see Figure 4-38).



Figure 4-38 Guided Setup wizard - Use VPN using an Existing Internet Connection

We click **Next** accept the VPN connection for our HMC support services and continue with the Guided Setup wizard.

37. The Guided Setup Wizard - Configure Connectivity using a Pass-Through System panel is shown (Figure 4-39). The HMC can use another system in your network which already has a VPN or dial-up connection to IBM service and support. This system could be an i5/OS V5R3 partition or another HMC.

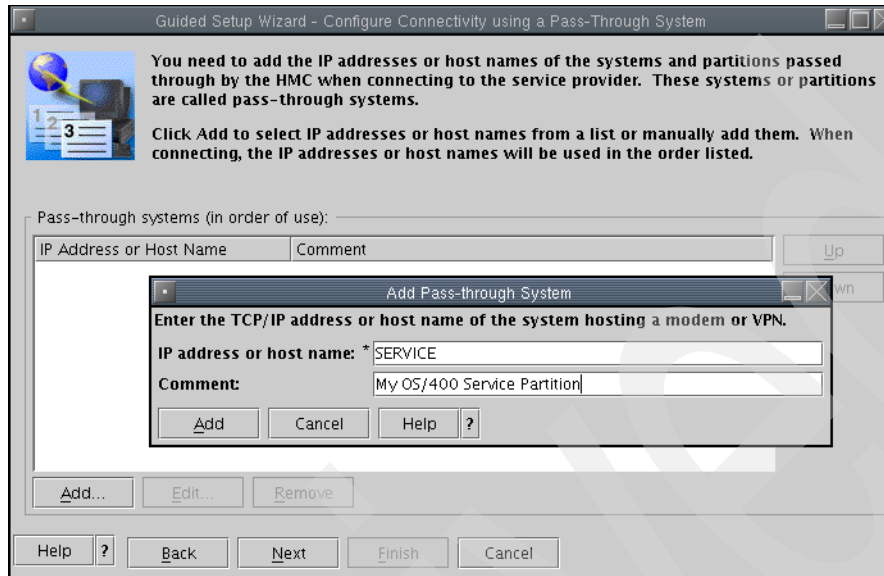


Figure 4-39 Guided Setup wizard - Pass-Through connectivity part 1

Click the **Add** button and enter the IP address or host name of your pass-through system. Type some meaningful comment text and click the **Add** button to accept the values entered. You can add multiple pass-through systems here. The order listed will be the order in which the pass-through systems will be used.

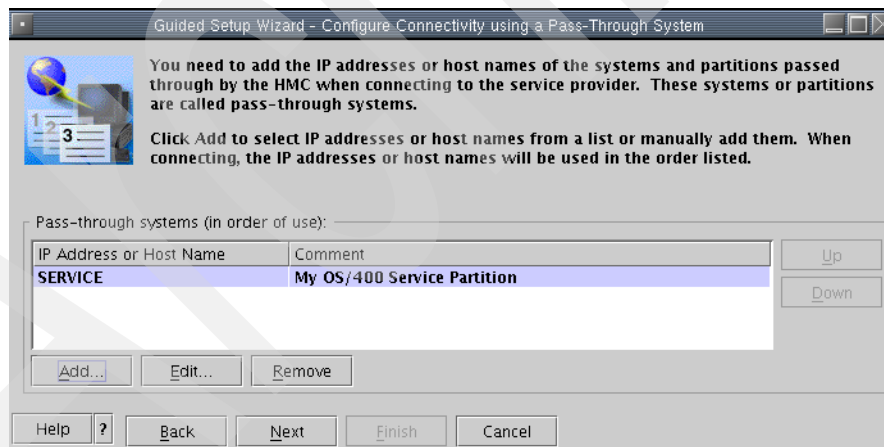


Figure 4-40 Guided Setup wizard - Pass-Through connectivity part 2

In our example we decided to use an i5/OS V5R3 partition call SERVICE which had a VPN connection to IBM service and support.

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

38. The Guided Setup Wizard - Authorize Users for Electronic Service Agent panel is now displayed (see Figure 4-41).

The information collected and sent to IBM by the HMC can be seen on the IBM Electronic Service Agent Web site:

<http://www.ibm.com/support/electronic>

To access this data on the Web, you must have a registered IBM ID and authorized that ID through the HMC. You can register IBM IDs via the Web site:

<https://www.ibm.com/registration/selfreg>

Enter a valid IBM ID and an optional second IBM ID if required, in the Web authorization panel. The Guided Setup will only allow you to authorize two user IDs to access the data sent by the HMC to IBM. However, you can submit as many registrations as you like by clicking **Service Applications** → **Service Agent** → **eService Registration** from the HMC desktop.

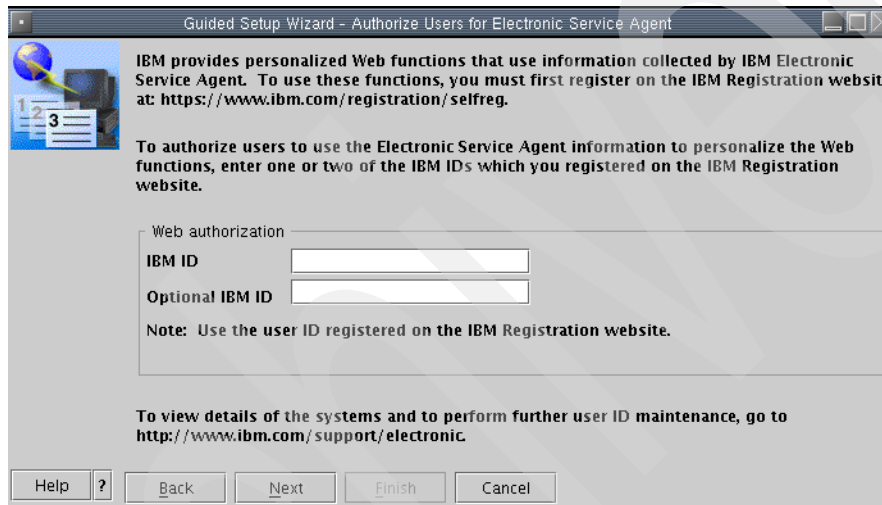


Figure 4-41 Guided Setup wizard - Authorized Users for ESA

If you do not have valid IBM ID, you can choose to leave this panel blank and manually complete this information later by clicking **Service Applications** → **Service Agent** → **eService Registration** from the HMC desktop.

In our example, we decide to register our IBM ID later. Click **Next** to continue.

39. The Guided Setup Wizard - Notification of Problem Events display is shown (see Figure 4-42). The HMC can alert your administrators of problems with the HMC or its managed systems via e-mail.

You can choose whether to notify your administrators of only problems reported to IBM (Only call-home problem events) or of all problem events generated.

Enter the IP address and port of your SMTP server. Then click the **Add** button and enter your administrator's e-mail address and the notification type required. Click **Add** to accept these values and return to the previous panel. You may enter multiple e-mail addresses by repeating this process.

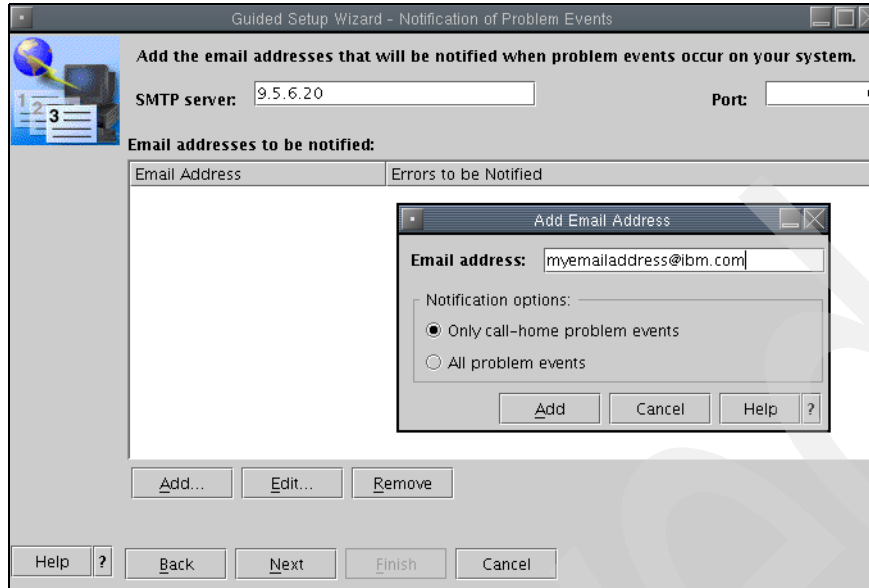


Figure 4-42 Guide Setup Wizard - Notification of Problem Events

In our example, we enter the SMTP server IP address/port and our administrator's e-mail address. We will only alert our administrator when a call-home problem event is generated.

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

40. The Guided Setup wizard - Summary panel is displayed (Figure 4-43). You can see all the changes that the Guided Setup wizard is **about** to make.

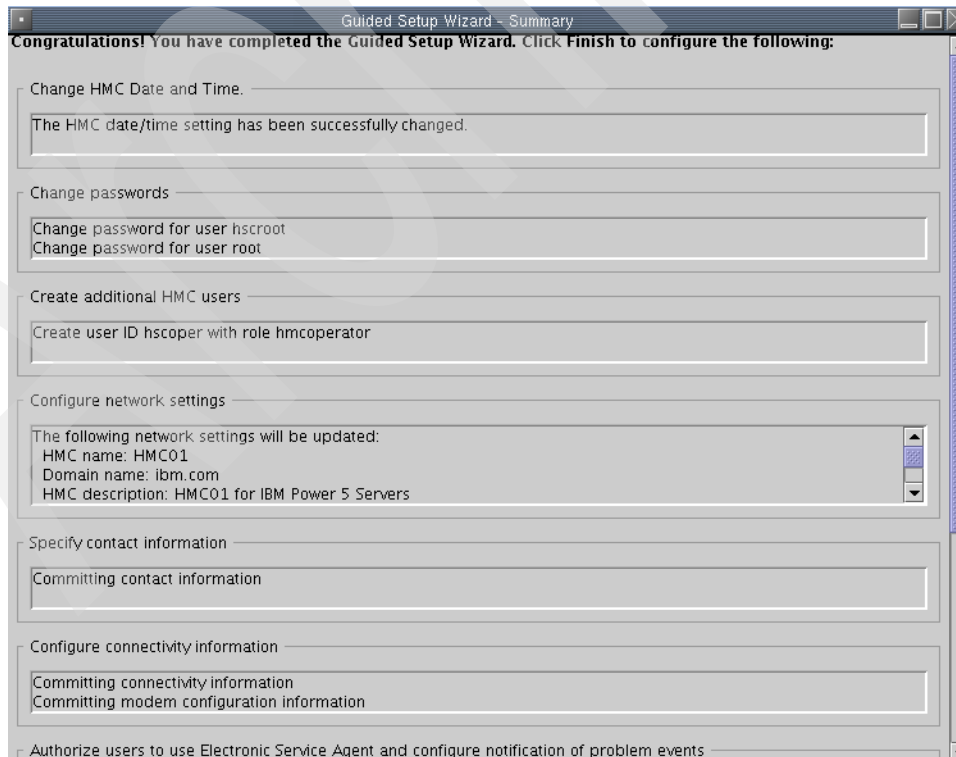


Figure 4-43 Guided Setup wizard - Summary panel - top

Important: At this stage nothing has actually been changed on the HMC. If you press the **Cancel** button, all changes made through the Guided Setup will be lost.

In our example, we click the **Finish** button to apply all our HMC changes.

41. The Guide Setup Wizard - Status panel is displayed (Figure 4-44). As each task completes, its status is automatically updated.

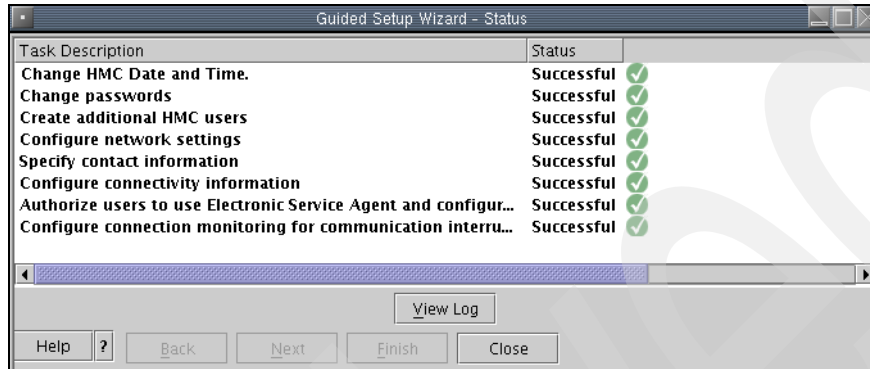


Figure 4-44 Guided Setup Wizard - Status panel

You can review the log by clicking the **View Log** button; this is especially useful if for any reason you have tasks that have a failed status. Figure 4-45 shows an example of our successful log file output.

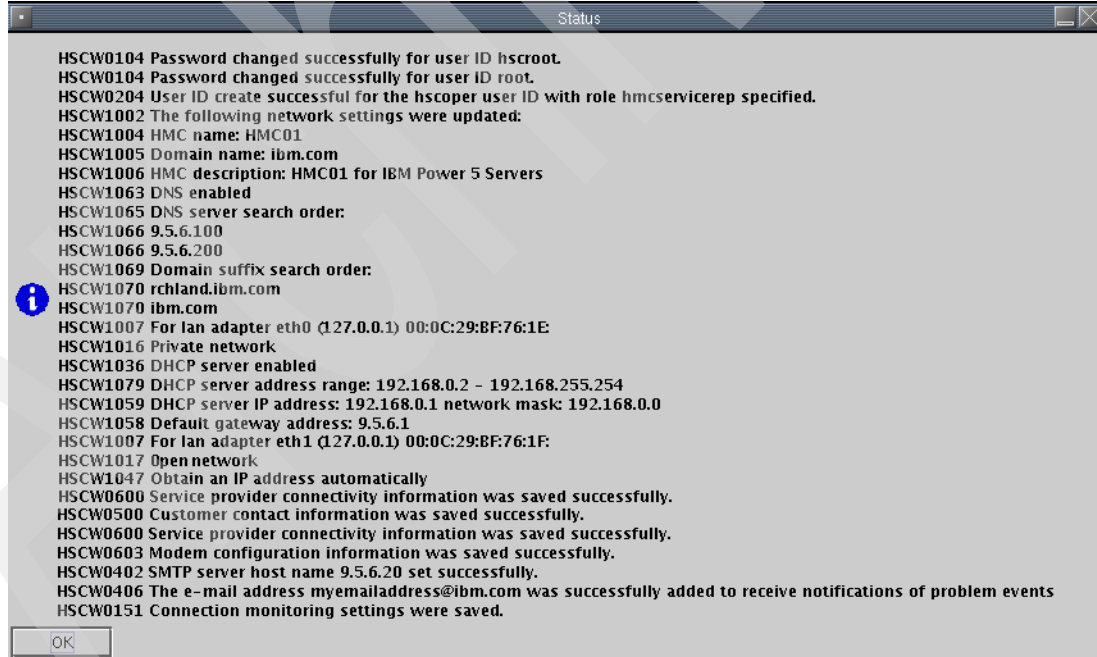


Figure 4-45 Guided Setup Wizard - Status log file

Click **OK** to return to the previous Guided Setup Status panel and click **Close**.

If you have configured any network settings during the Guided Setup Wizard, then you will probably receive a message asking you whether you wish to reboot the HMC (see Figure 4-46).

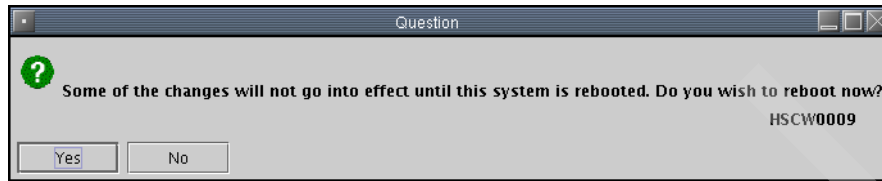


Figure 4-46 Guided Setup wizard - Reboot message

In our example, we click **Yes** to reboot the HMC and activate our new network settings.

This completes the HMC Guided 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 direct connected to the HMC, some tasks could be missed at first.

For example, you may not have all the user profiles set up, or you may not have both the LAN adapters set up.

Archived



Partition creation using the HMC

In this chapter we discuss the following topics:

- ▶ System and partition profiles
- ▶ Creating an i5/OS logical partition using the HMC
- ▶ Creating additional partition profiles for an existing logical partition
- ▶ Changing the default profile for a partition

5.1 System and partition profiles

This section discusses the concept of system and partition profiles and how they are used.

5.1.1 System profiles

A system profile is a collection of one or more partition profiles. System profiles can be used to specify which partition profiles are activated at the same time.

5.1.2 Partition profiles

A partition profile represents a particular configuration for a logical partition. A partition profile contains information about the resources assigned to the partition. These resources include memory, processor processing capacity, and physical I/O slots. Partition profiles also include information about which devices provide important partition functions (for example, load source, console, and alternate IPL).

Note: Unlike previous versions of logical partitioning, a partition no longer owns any resources. When a profile is activated, resources are allocated to the partition if they are available. Partition activation may fail if certain required or critical resources are not available or if processing resource minimums cannot be met.

Each logical partition has at least one partition profile associated with it. A logical partition can have more than one partition profile associated with it. However, only one profile can be active at any given time. Where there are two or more partition profiles for a logical partition, one of the profiles needs to be specified as the default partition profile. For a logical partition with only one profile, that profile is automatically the default partition profile. The default partition profile is used when the partition is activated if no other profile is specified.

5.1.3 Uses of partition profiles

Partition profiles provide a degree of flexibility that was not available with previous implementations of logical partitioning.

Partition profiles can be useful when it is not possible to dynamically move resources into a partition. For operating systems that currently do not have dynamic LPAR (DLPAR) support, partition profiles can be used to simulate the effect of DLPAR. Moving a processor between two partitions would involve the use of four profiles. The initial profile for partition A would be powered off and another profile with one less processor would be started. The initial profile for partition B would then be powered off and another profile started with one additional processor.

A partition could have a regular profile for normal day to day operations and a second profile for large month or quarter end processing. A development partition could also have another profile for situations where very large compiles or builds are needed. The provider of disaster recovery services can maintain several partition profiles for a given partition and activate one based on which customer is coming in. A standby partition may have different profiles representing different types of workloads that can be deployed when needed.

Note: Activating a profile for a partition requires an IPL of that partition.

5.2 Creating an i5/OS logical partition through the HMC

Creating a partition is a multiple step process. Partition creation can be accomplished through either the HMC GUI (graphical user interface) or CLI (command line interface). We will focus on using the HMC GUI.

Important: Typically all partition creation and management is performed through the Hardware Management Console (HMC). The CLI is an advanced option and still requires an HMC.

5.2.1 Partition creation overview

Table 5-1 is an overview of the steps involved in creating a new logical partition, and can be used as a checklist. These are further discussed in the following sections.

Table 5-1 Checklist for creating a new logical partition

Task No	Description	Comment
	Start the partition wizard	
	Partition name ID and type	
	Workload management participation	
	Partition profile name	
	Partition profile memory values	
	Partition profile processors value	
	Partition profile interactive values (if available)	
	Allocation of physical I/O resources	
	Allocation of virtual I/O resources	
	Tag partition critical resources (i5/OS only)	
	Virtual and HSL Opticonnect	
	Power control for guest partitions	
	SFP surveillance and auto start settings	

5.2.2 Starting the create partition wizard

If initial setup and configuration of the HMC is required, refer to Chapter 3, “HMC overview and planning” on page 45.

Attention: An HMC can manage multiple servers. Before creating a logical partition, make sure that the correct server is selected.

1. If you are not already connected to the HMC, sign in with an HMC user profile that has either System Administrator or Advanced Operator authority. The capabilities of various HMC roles are discussed in 7.3, “HMC User Management” on page 238.
2. From the Navigation Area on the left hand side of the HMC display, select the managed host where the logical partition will be created and expand as follows:

Server and Partition → Server Management

After selecting **Server Management**, the right hand side pane is populated with logical partition information. If any logical partitions were already created, they can be viewed by expanding the **Partitions** folder. See Figure 5-1.

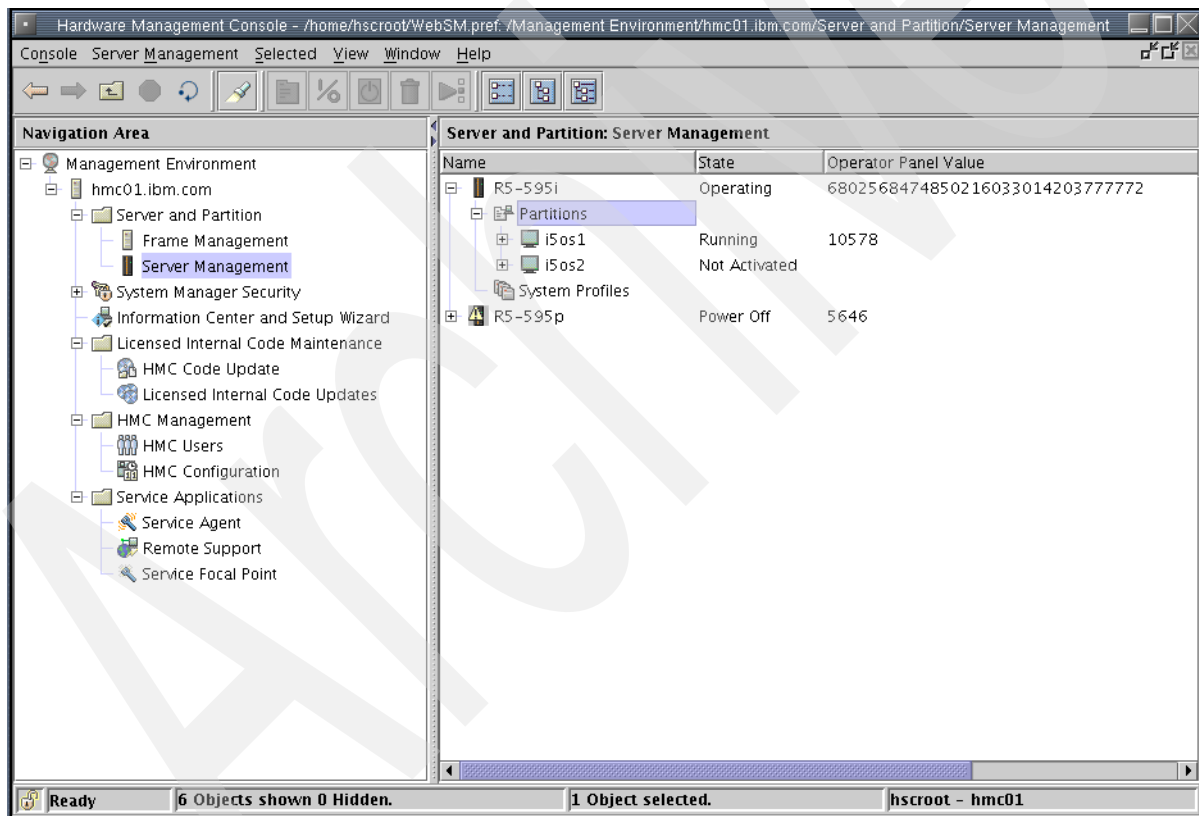


Figure 5-1 On the selected host, server management is selected

- To start the partition creation process, right-click **Partitions**, and from the pop up menu shown in Figure 5-2, select:

Create → Logical Partition

Alternatively, with server management selected in the Navigation Area, the logical partition creation wizard can be accessed from the menu bar as follows:

Selected → Create → Logical Partition

Either way, once selected, the Create Logical Partition wizard starts.

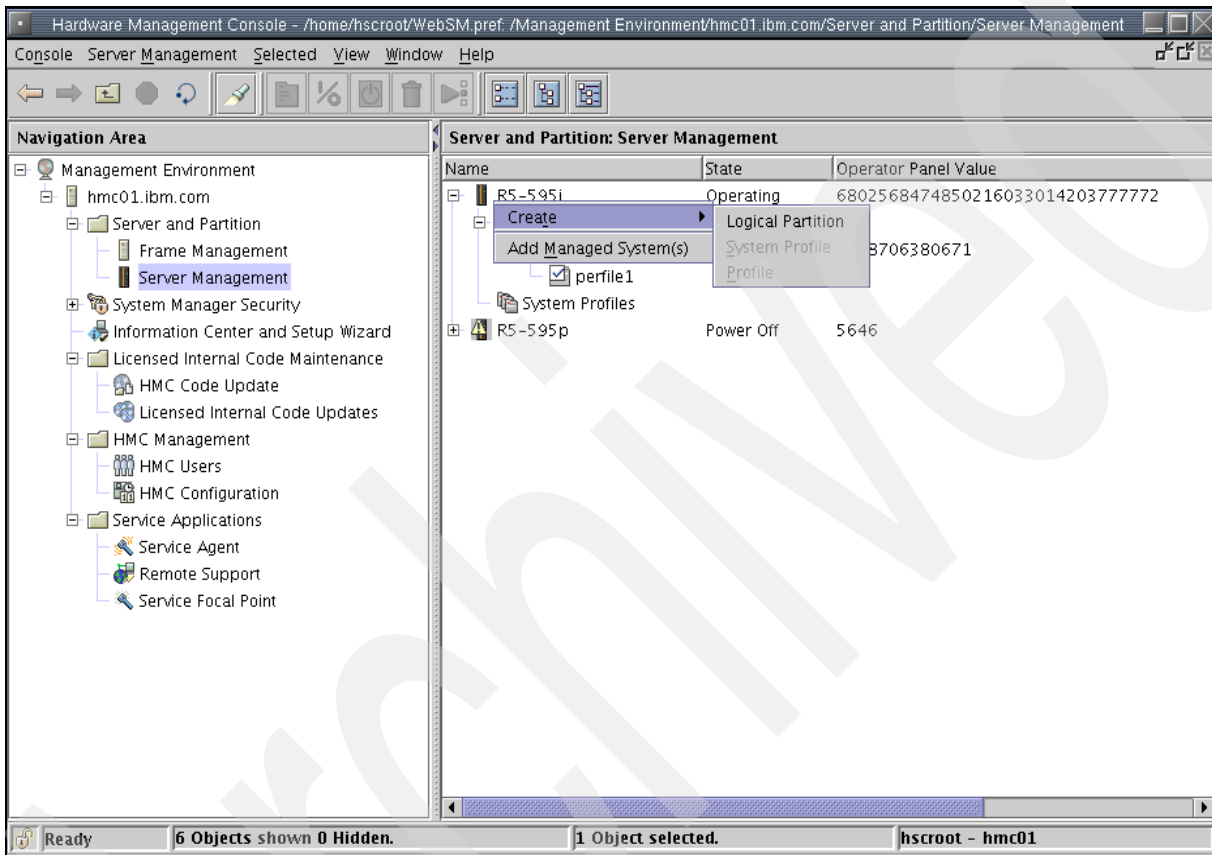


Figure 5-2 Invoking the create partition wizard

It may take a number of seconds before the wizard pane opens (Figure 5-3), since the HMC has a lot of information gathering to perform.



Figure 5-3 Partition wizard first pane

5.2.3 Specifying the partition name, partition ID, and partition type

Each partition needs to have the following pieces of information:

- ▶ Partition ID Number, a numeric value between 1 and 254
- ▶ Partition Name, a text string up to 31 characters
- ▶ Partition Type, indicating which operating system will be installed on the partition:
 - AIX or Linux
 - i5/OS
 - Virtual I/O Server

Figure 5-4 shows a sample partition name and partition ID for an i5/OS. Once the fields are filled in, click **Next >** to proceed to the next panel.

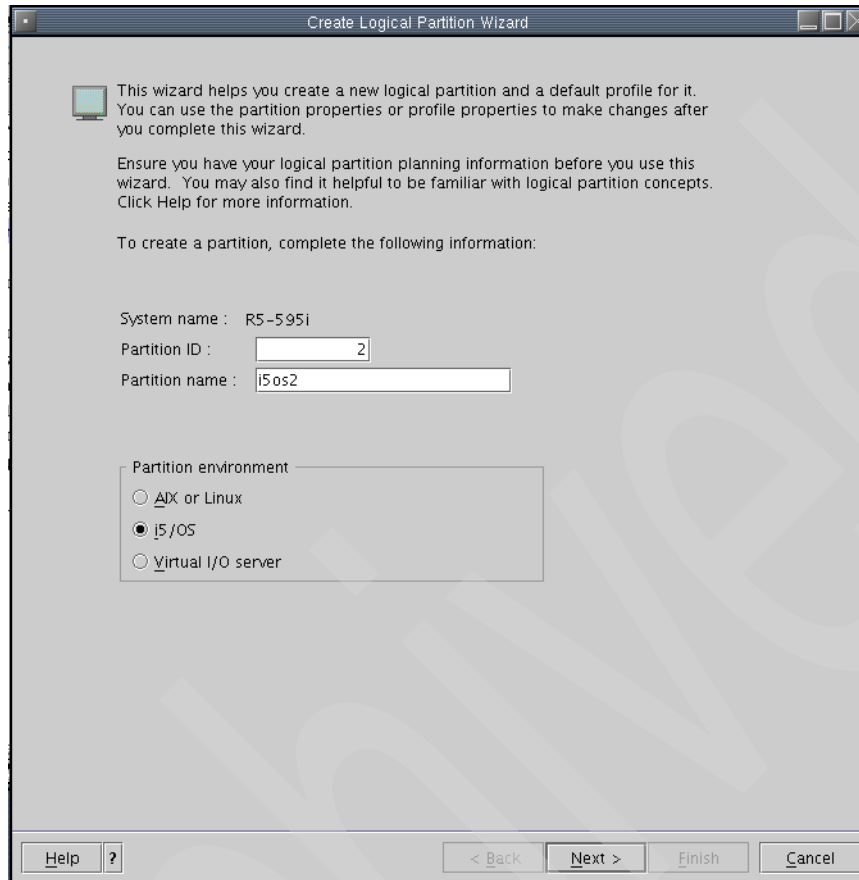


Figure 5-4 Specifying the partition name, partition ID, and partition type

5.2.4 Workload management group

If the partition will participate in a workload management group, select **Yes** and specify the **GroupID** (Figure 5-5).

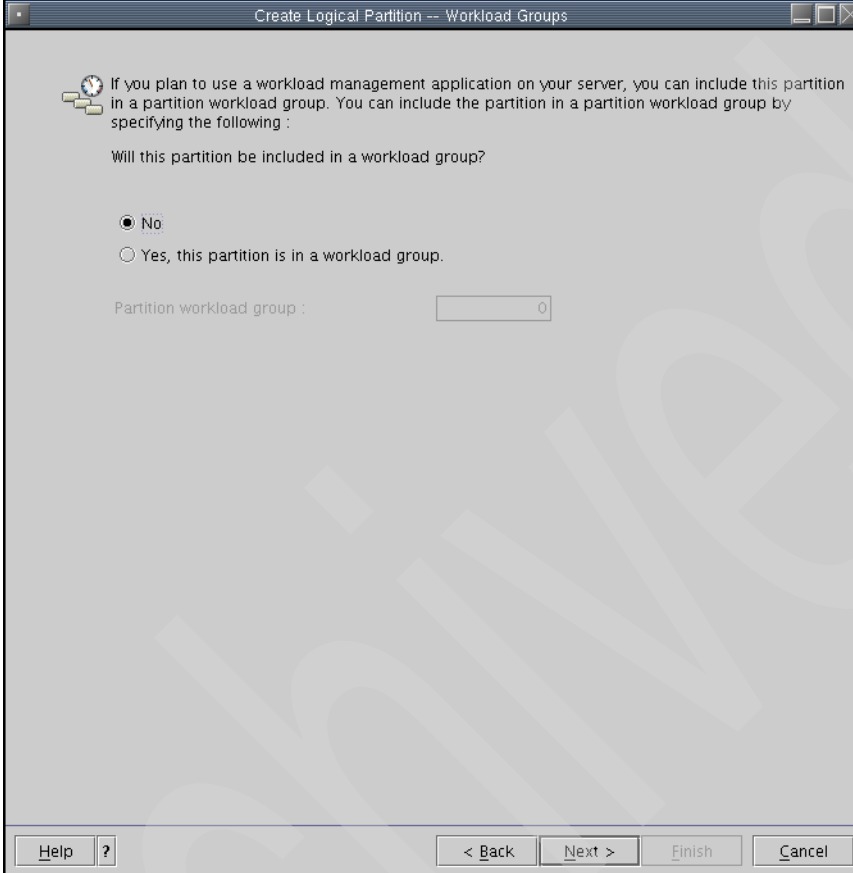


Figure 5-5 shows a dialog box titled "Create Logical Partition -- Workload Groups". The dialog contains the following text and controls:

- Text: "If you plan to use a workload management application on your server, you can include this partition in a partition workload group. You can include the partition in a partition workload group by specifying the following :
- Text: "Will this partition be included in a workload group?"
- Radio buttons:
 - No
 - Yes, this partition is in a workload group.
- Text: "Partition workload group :"
- Text input field: "0"
- Buttons: "Help", "< Back", "Next >", "Finish", "Cancel"

Figure 5-5 Workload management group

Once complete, click **Next >** to proceed to the next panel.

5.2.5 Partition profile name

Each logical partition requires at least one partition profile. The profile name can contain up to 31 characters. The profile name should be descriptive — *month end processing*, for example.

Once complete, click **Next >** to proceed to the next panel (Figure 5-6).

The screenshot shows a window titled "Create Logical Partition Profile". Inside, there is a folder icon and text explaining that a profile specifies processor, memory, and I/O allocations. It asks for information to create a default profile. The fields are: System name: R5-5951, Partition name: I5os2, Partition ID: 2, and Profile name: profile1. Below these is a checkbox labeled "Use all the resources in the system." At the bottom are buttons for Help, Back, Next >, Finish, and Cancel.

Figure 5-6 Specifying the partition profile name

5.2.6 Partition memory

Figure 5-7 shows the initial partition memory panel that is displayed. Memory can be allocated in a combination of megabytes (MB) or gigabytes (GB). Megabyte allocations are restricted to multiples of the logical memory block size. Currently, the value for the logical memory block size is 16 MB.

Notice that the default MB values do not automatically disappear when you add values into the GB window. These values must be removed manually. Also, the 128 MB minimum value is too small a minimum for an i5/OS partition.

Specify desired, minimum and maximum amounts of memory for this profile using a combination of the gigabyte and megabyte fields below.

Installed memory (MB): 11534336

Current memory available for partition usage (MB) : 10485504

Minimum memory	Desired memory	Maximum memory
10 GB	20 GB	30 GB
128 MB	128 MB	128 MB

Help ? < Back Next > Finish Cancel

Figure 5-7 Initial partition memory panel

There are three memory values for a partition profile:

- ▶ **Minimum memory:** This amount of memory is required for the partition. The profile will fail to activate if the minimum memory is not met.
- ▶ **Desired memory:** This is the requested amount of memory for the partition. On profile activation, the partition will receive an amount of memory between the minimum and desired amounts depending on what is available.
- ▶ **Maximum memory:** The maximum represents the upper limit for memory. Memory allocations cannot go above the maximum. In order to go above the maximum, the maximum value needs to be changed.

Once complete, click **Next >** to proceed to the next panel.

5.2.7 Partition processors

The next resource to configure for a partition profile is processing capacity. The choices are

- ▶ Dedicated processors
- ▶ Capped shared processors
- ▶ Uncapped shared processors

Figure 5-8 shows the processor selection panel. The two main choices are Dedicated and Shared. Dedicated processors are intended for use solely by the partition to which they are assigned. Shared processors allow for fractional processors to be assigned. This additional flexibility does require some overhead and may not be suitable for all workloads.

For additional information on the differences between dedicated, capped shared, and uncapped shared processors, refer to Chapter 1, “Introduction to LPAR on IBM System i5” on page 1.

Dedicated processors

To use dedicated processors for a partition profile, select the **Dedicated** radio button and click **Next >**.

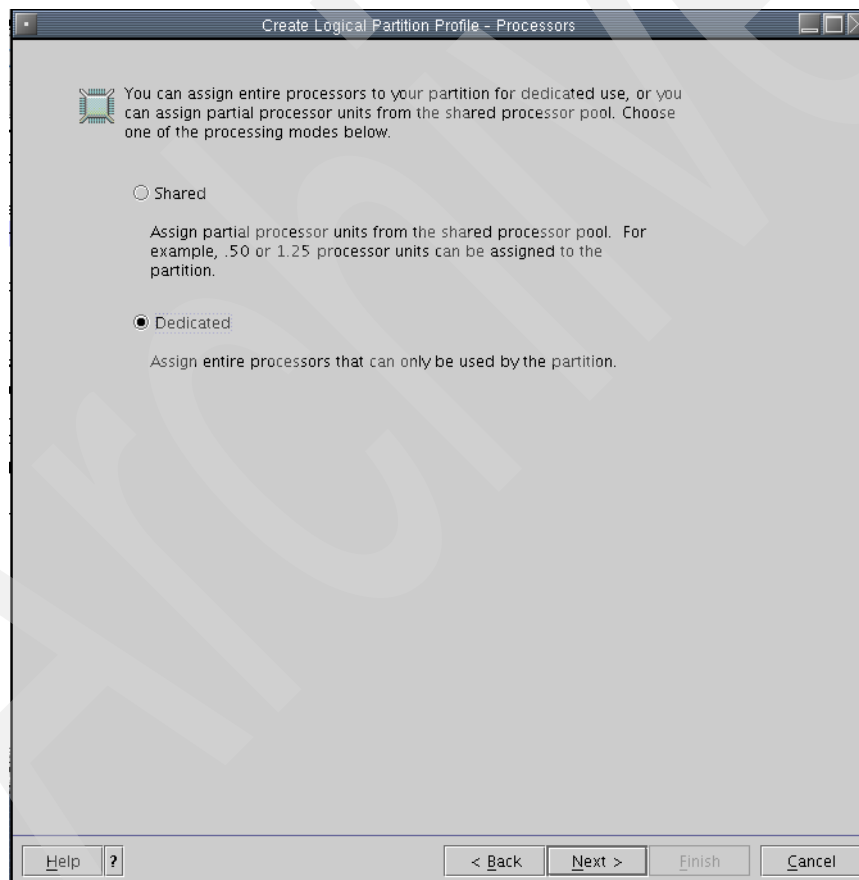


Figure 5-8 Choosing dedicated processors for a partition profile

For any dedicated processor partition profile, three values are required (Figure 5-9).

- ▶ **Desired processors:** This is the requested amount of processors for the partition. On profile activation, the partition will receive a number of processors between the minimum and desired amounts depending on what is available.
- ▶ **Minimum processors:** This number of processors is required for the partition. The profile will fail to activate if the minimum number of processors is not met.
- ▶ **Maximum processors:** The maximum represents the upper limit for processors. Processor allocations cannot go above the maximum. In order to go above the maximum, the maximum value needs to be changed.

Figure 5-9 shows a dialog box titled "Create Logical Partition Profile - Processing Settings". The dialog contains a processor icon and the instruction: "Specify the desired, minimum, and maximum processing settings in the fields below." Below this, there are four input fields: "Total number of processors : 32", "Minimum processors: 1", "Desired processors: 1", and "Maximum processors: 1".

Figure 5-9 Initial dedicated processor panel

Figure 5-10 shows a sample completed dedicated processor configuration. This partition profile would require at least 1 dedicated processor in order to start. Depending on whether or not processor resources are over committed, this partition profile will be allocated between 1 and 4 processors when activated. As configured, this profile will not allow for more than 7 processors to be in the partition.

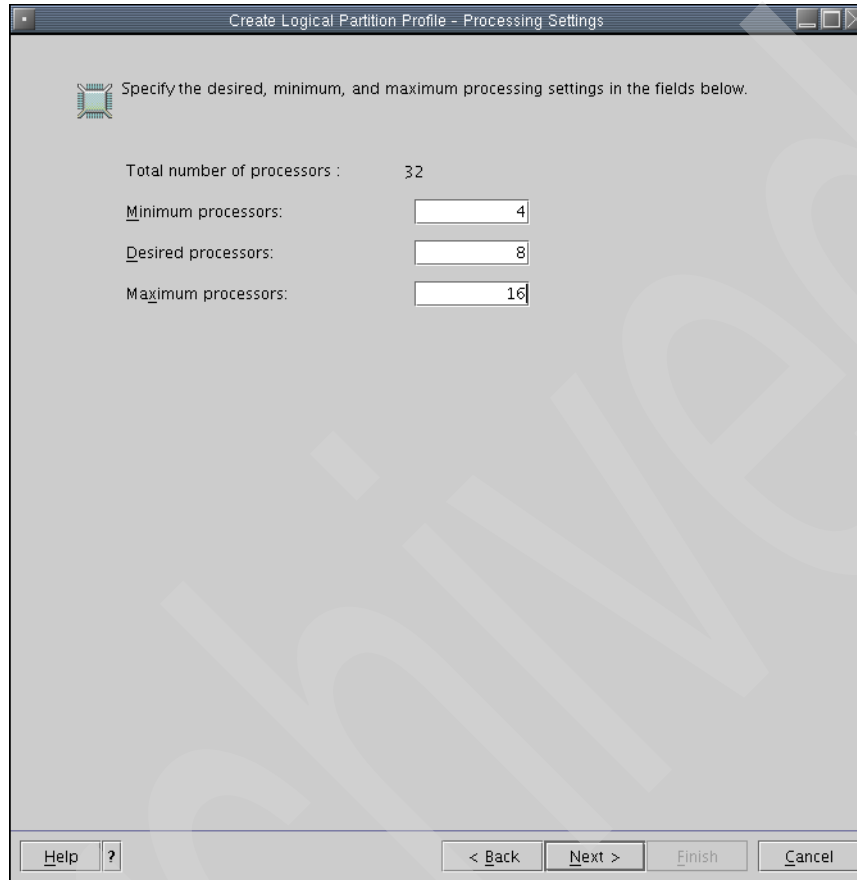


Figure 5-10 Sample completed processor configuration

Once complete, click **Next >** to proceed to the next panel.

Capped shared processor partition

To use shared processors for a partition profile, select the **Shared** radio button as shown in Figure 5-11 and click **Next >**.

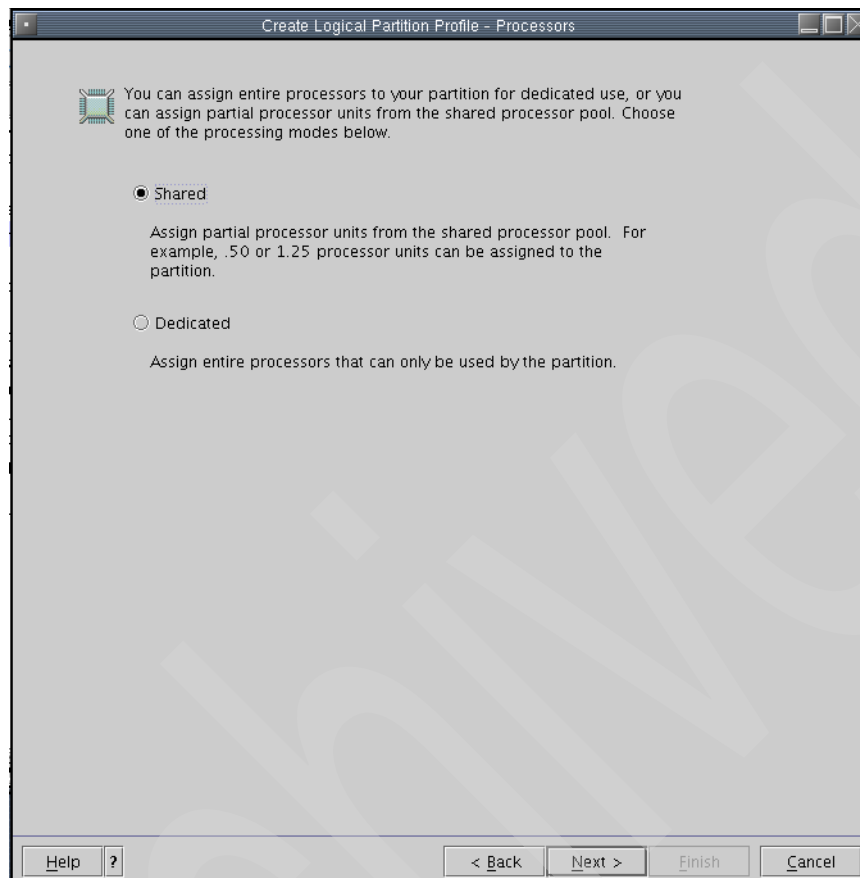


Figure 5-11 Choosing shared processors for a partition profile

There are several pieces of information required for a shared processor partition. The first three of these are:

- ▶ **Desired processing units:** This is the requested amount of processing units for the partition. On profile activation, the partition will receive a number of processing units between the minimum and desired amounts depending on what is available.
- ▶ **Minimum processing units:** This number of processing unit is required for the partition. The profile will fail to activate if the minimum number of processing units is not met.
- ▶ **Maximum processing units:** The maximum represents the upper limit for processing units. Processing unit allocations cannot go above the maximum. In order to go above the maximum, the maximum value needs to be changed.

Once these are filled in, click the **Advanced** button to bring up the sharing mode dialog (Figure 5-12).

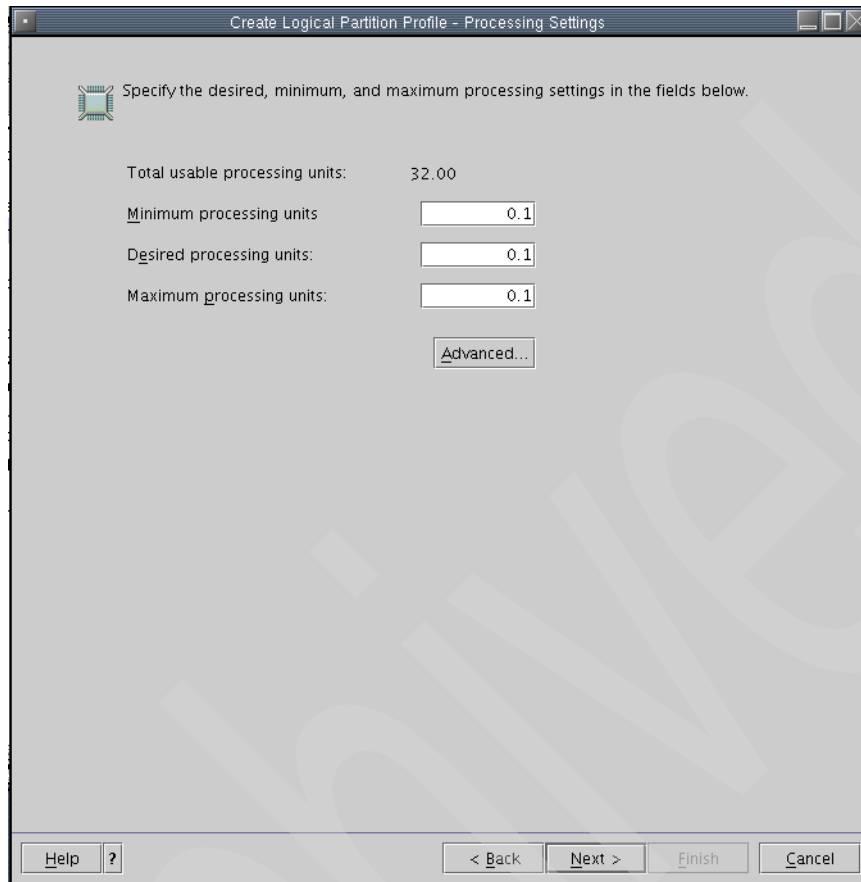


Figure 5-12 Initial shared processor panel

In the sharing mode properties, click the **Capped** radio button (Figure 5-13). The desired, minimum, and maximum number for virtual processors need to be specified. At a minimum, use the values for desired, minimum, and maximum processor units rounded up to the next whole number. For example, for 1.25 processor units, use at least 2 for the number of virtual processors.

Once complete, click **OK** to close the Sharing Mode Properties dialog, and then click **Next >** to proceed to the next panel.

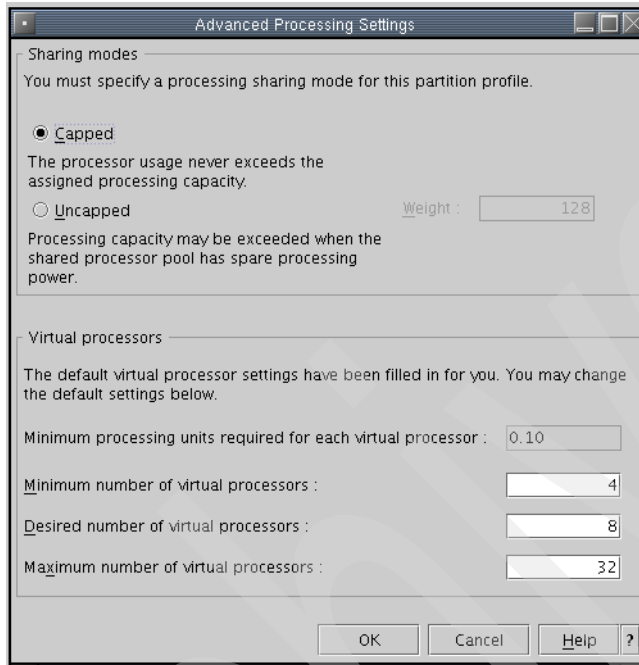


Figure 5-13 Sharing mode properties for a capped shared processor partition

Uncapped shared processor partition

For the most part, the creation of an uncapped shared processor partition profile is identical to creating a capped shared processor partition profile. In order to create an uncapped shared processor partition profile, click the **Advanced** button to bring up the sharing mode dialog. As shown in Figure 5-14, make sure that the **Uncapped** radio button is selected. The **Weight** of the partition can be changed if the default of 128 is not desired.

Once complete, click **OK** to close the Sharing Mode Properties dialog, and then click **Next >** to proceed to the next panel.

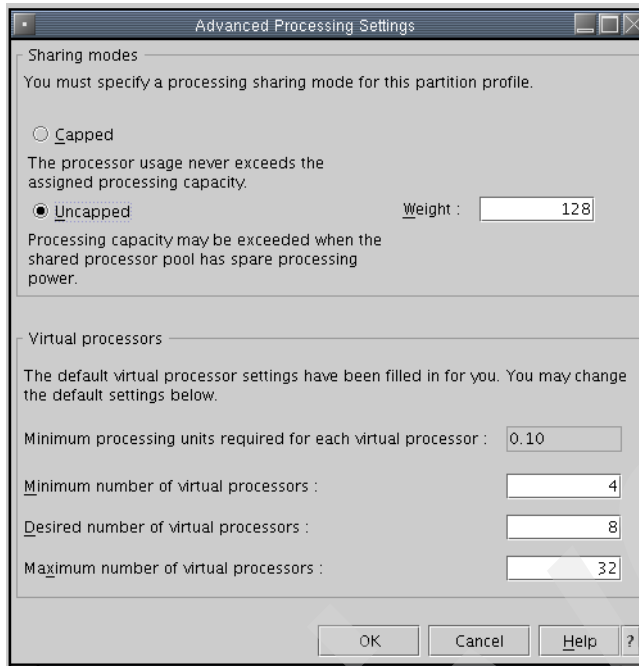


Figure 5-14 Sharing mode properties for an uncapped shared processor partition

5.2.8 Interactive (5250 OLTP) capacity

Depending on the particular system model, assigning interactive capacity for the partition profile may or may not be required. Zero interactive and enterprise edition systems do not require the interactive feature to be assigned to the profile.

The 5250 OLTP screen will not display on models other than the i520.

5.2.9 Allocate physical I/O resources

The next step in the creation of a partition profile is the allocation of physical IO resources to the partition profile (Figure 5-15).

Information about a particular resource can be obtained by selecting the resource and clicking the **Property** button.

Resources can be added to the partition profile as either required or desired. A required resource is needed for the partition to start when the profile is activated. The load source disk unit is an example of a required resource. Desired resources are assigned to the partition if they are available when the profile is activated. The absence of a desired resource does not prevent the partition profile from activating. A tape unit or optical device might be considered a desired resource.

Multiple items can be selected while holding down the control key and left clicking the items to be selected. Each additional item may need to be selected twice for it to become highlighted. A range of resources can be selected while holding down the shift key and then selecting the first and last item in the range.

Once selected, add the resources to the profile by clicking either **Add as Required** or **Add as Desired**.

Once complete, click **Next >** to proceed to the next panel.

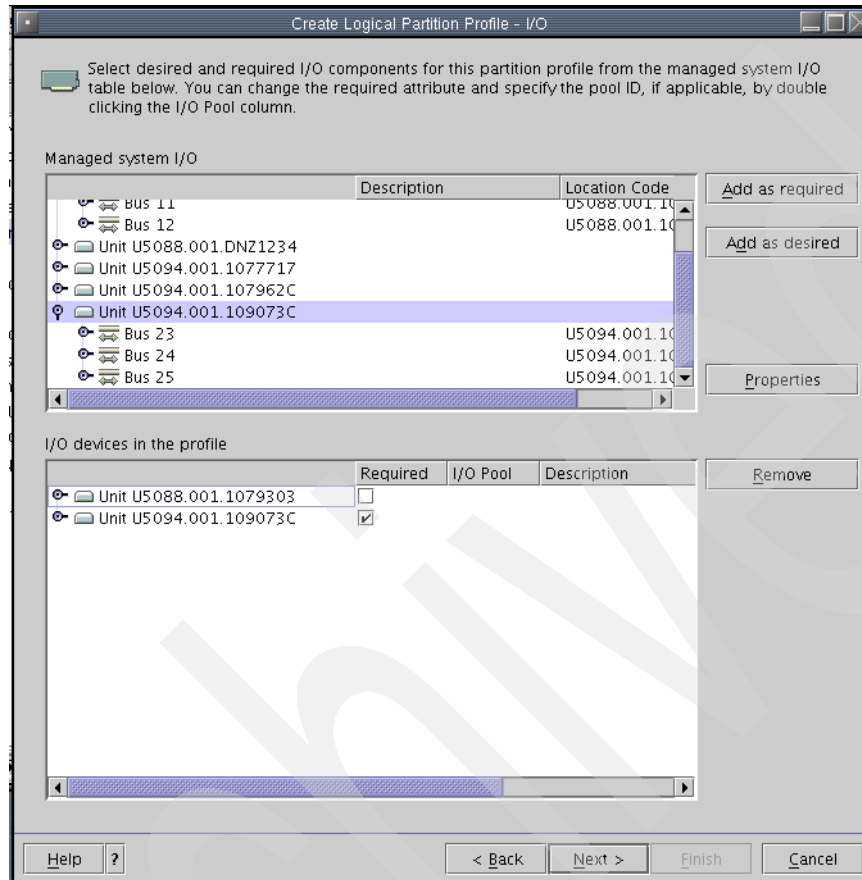


Figure 5-15 Allocating physical I/O to a partition profile

Location codes

Hardware may or may not be identified in the HMC by resource type and model. This information is provided by the converged Hypervisor only if this information has been provided to it by the operating system.

In the absence of resource type and models, location codes are used.

Location codes can identify either the physical location of a resource or the logical path that the system uses to communicate with a particular resource.

Physical location codes

An example of a physical location code is as follows:

U970305010ABCDE-P3-C31

In this location code, the resource is contained in a unit of type 9703, model 050, and has serial number 10-ABCDE. This resource connects to backplane P3 and is in the connector labeled C31. The serial number may be omitted if it is not known.

Logical path location codes

An example of a logical path location code is as follows:

U970305010ABCDE-P3-C31-T2-L23

The first portion (through the *T* prefix) of the logical path logical code is the physical location code for the resource that communicates with the desired resource. The string that follows after the T prefix identifies the particular resource.

Note: It is possible for a device to have more than one logical path location code. An external tape device that is connected to two IO adapters would have two logical path location codes.

5.2.10 Virtual IO adapters

Virtual IO adapters give a partition the flexibility to use certain types of resources with requiring the physical hardware to be present. If the partition will be using virtual IO, make sure that the **Yes** radio button is selected. Otherwise, select the **No** radio button. Once complete, click **Next >** to proceed to the next panel (Figure 5-16).

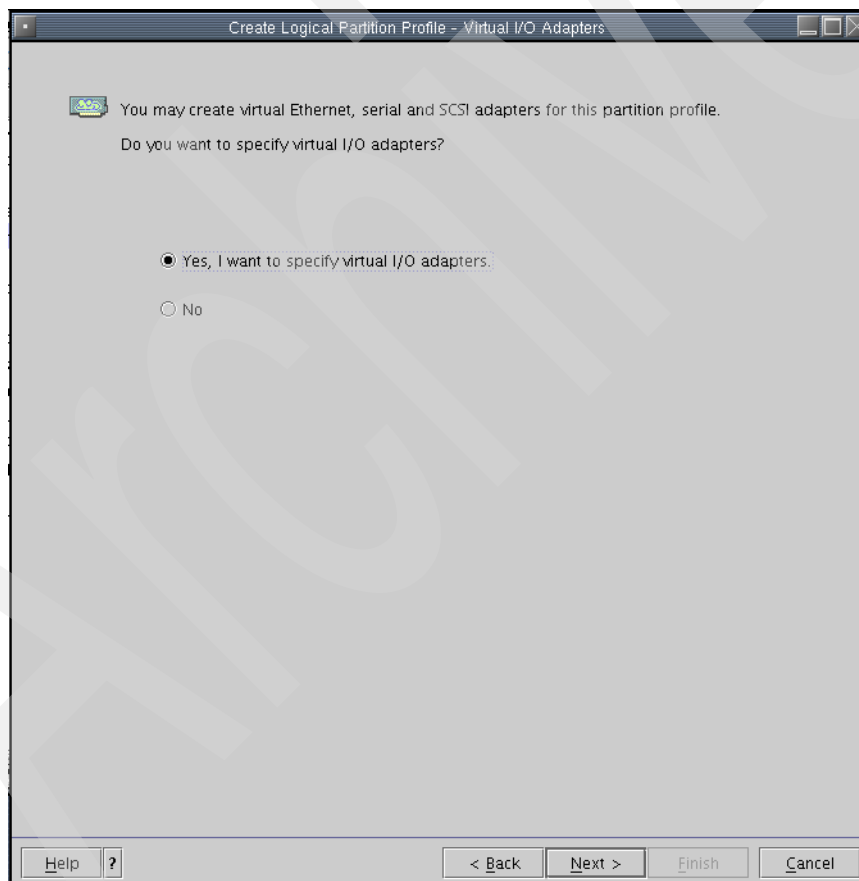


Figure 5-16 Specifying virtual IO

Creating Virtual IO Adapters

The following virtual IO adapters can be created:

- ▶ Virtual Ethernet
- ▶ Virtual Serial
- ▶ Virtual SCSI

All virtual IO adapters reside on a single virtual system bus (Figure 5-17). The **maximum number of virtual adapters** is a user editable field that specifies how many virtual IO adapters can connect to the virtual system bus. Clicking **Next >** will advance to the next panel. In order to create a virtual IO adapter, select the adapter type radio button and click the **Create** button.

Slot numbers 0 and 1 are reserved for system use.

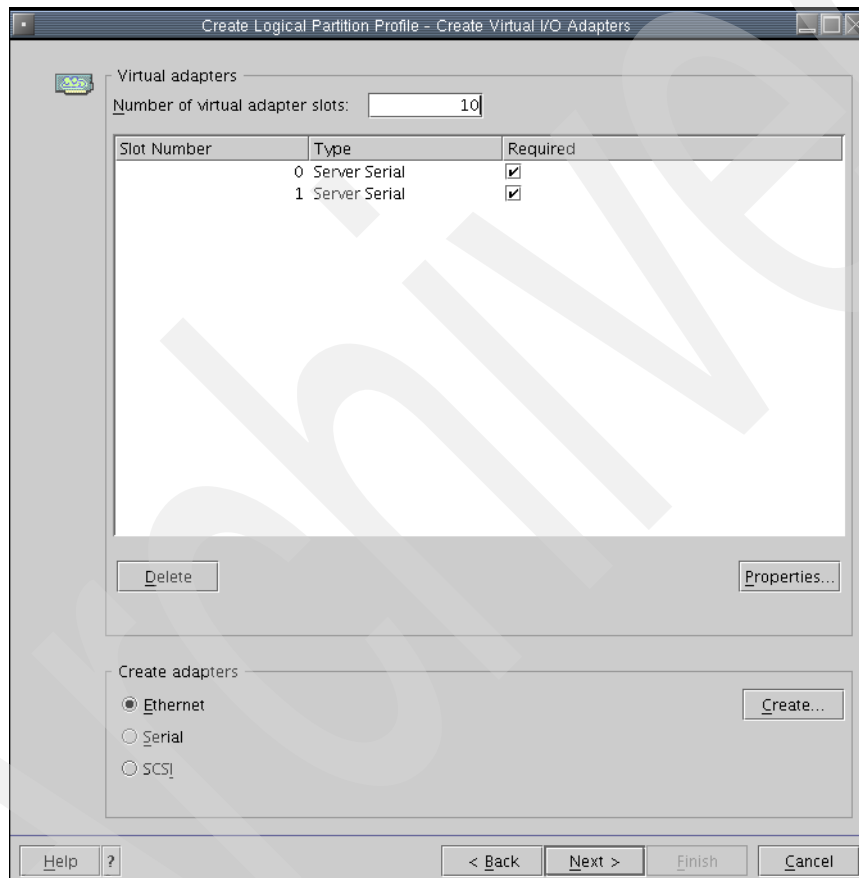


Figure 5-17 Creating virtual IO adapters

Virtual Ethernet

Virtual ethernet allows for high speed (1 gigabit) interpartition communications. The slot number refers to the position of the virtual ethernet adapter on the virtual system bus. **Port virtual LAN ID** is analogous to the current virtual LAN/Ethernet ID. For two partitions to be able to talk, they will need to have a virtual ethernet adapter that has the same port virtual LAN ID. Having the virtual ethernet adapter occupy the same slot is not a requirement. See Figure 5-18.

Migrating an existing VLAN or virtual ethernet configuration requires special consider. Refer to the info center article, “Convert your preexisting virtual Ethernet configuration”, at: <http://publib.boulder.ibm.com/eserver/>

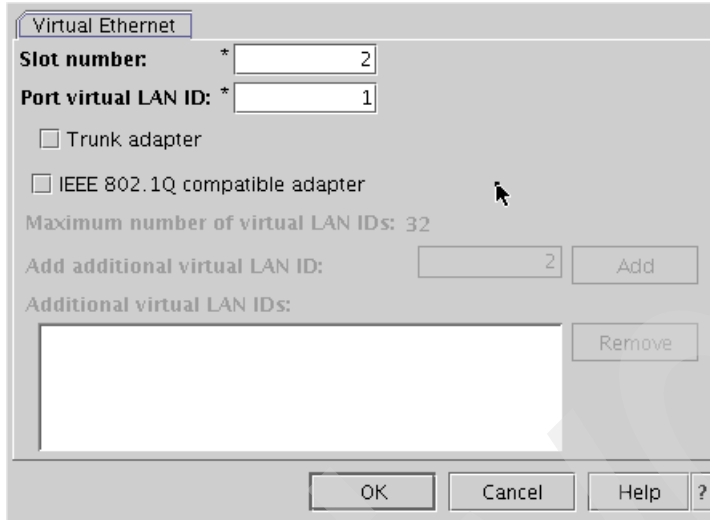


Figure 5-18 Creating a virtual ethernet resource

Virtual Serial

Virtual serial (see Figure 5-19) allows for the creation of an internal point-to-point connection. This connection is between the partition and either the HMC or another partition. An adapter is either client adapter or a server adapter. A client adapter is for a partition that wishes to use a resource that is provided by a server adapter.

In creating a virtual serial client adapter, the remote partition needs to be specified along with the slot number containing the virtual serial server adapter on the remote partition. By default, the HMC gets the console for a hosted guest partition. A pair of virtual serial adapters can be used to allow another partition to provide console function for the hosted guest partition as well.

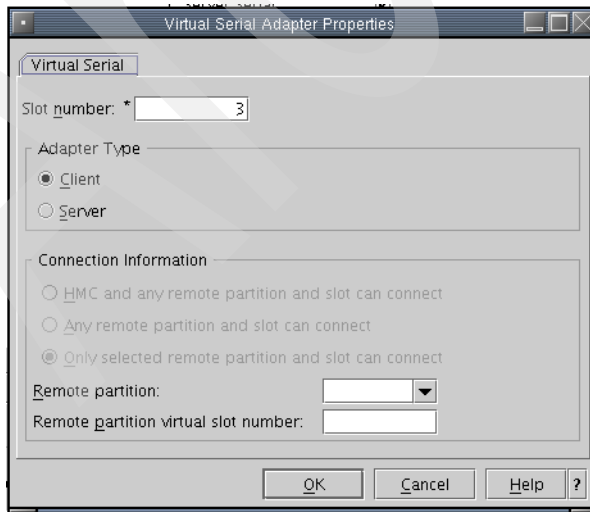


Figure 5-19 Creating a virtual serial client resource

To create a virtual serial server adapter, select **Server** for the adapter type (Figure 5-20). The connection information also needs to be specified. Connection information determines who can connect to the resource.

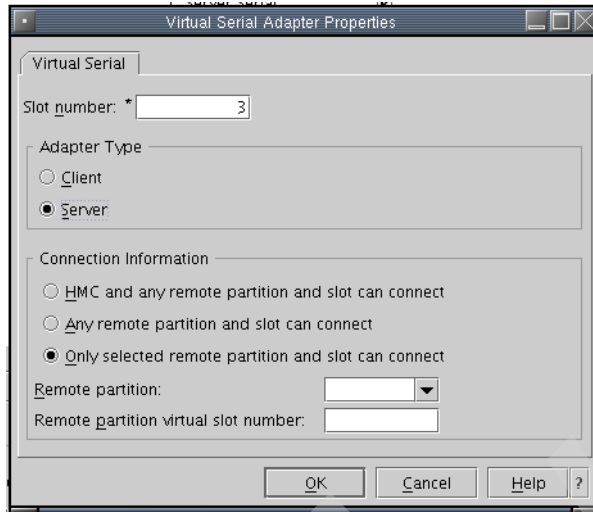


Figure 5-20 Creating a virtual serial server resource

Virtual SCSI

Virtual SCSI (see Figure 5-21) allows for a partition to use storage resources that physically reside in another partition. Storage resources include disk, tape, and optical. As with virtual serial, the Server adapter type is for a partition that is providing the resource, and the Client adapter type is for a partition that will be using the resource. A pair of virtual SCSI adapters can be created to provide disk access for a hosted guest partition.

Restriction: i5/OS does not support the use of virtual SCSI resource. An i5/OS partition can only act a server of virtual SCSI resources to AIX and Linux partitions.

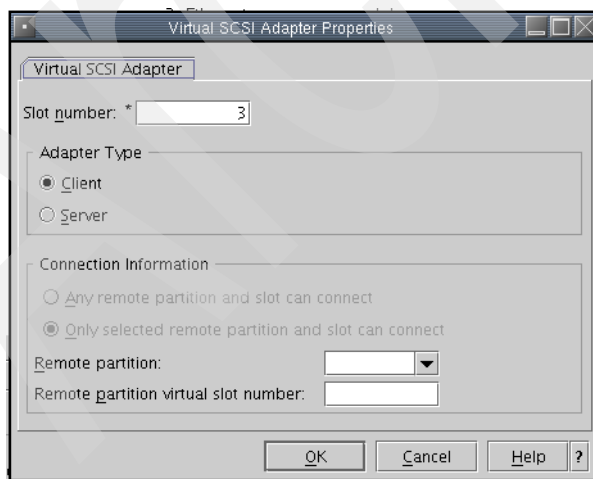


Figure 5-21 Creating a virtual SCSI client resource

To create a virtual SCSI server adapter, select **Server** for the adapter type (see Figure 5-22). The connection information also needs to be specified. Connection information determines who can connect to the resource. For a virtual disk that contains data that only needs to be accessed in a read only fashion, allowing **Any remote partition and slot can connect** would be okay. For write access, specifying **Only selected remote partition and slot can connect** may be a better choice.

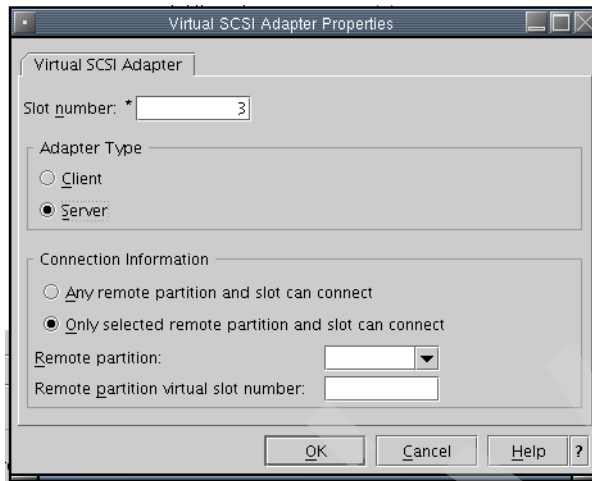


Figure 5-22 Creating a virtual SCSI server resource

5.2.11 Tagging partition resources

An i5/OS partition requires certain resources in order to function. These resources include a load source and console. For install, an alternate restart (D IPL) device is also required. The wizard will now ask us to tag these resources. Some of the resource tags are optional. Highlight the desired resource and click the **Select** button. After the resource is selected, click **Next>** to advance to the next panel.

Load source

The load source is used for IPLs from either the A or B side of the Licensed Internal Code. Selecting the load source IOP resource specifies what is to be used for regular IPLs and where to place Licensed Internal Code during an install. See Figure 5-23.

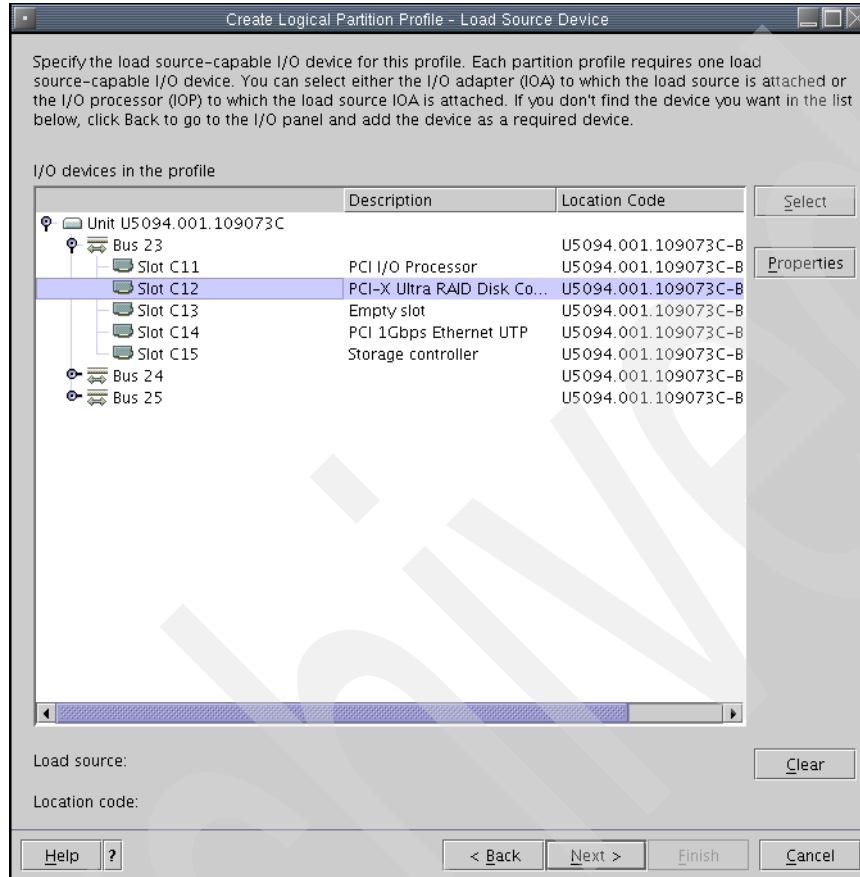


Figure 5-23 Selecting load source IOP resource

Alternate IPL (restart) device

The Alternate IPL device (Figure 5-24) is used for D mode IPLs when Licensed Internal Code needs to be installed or restored.

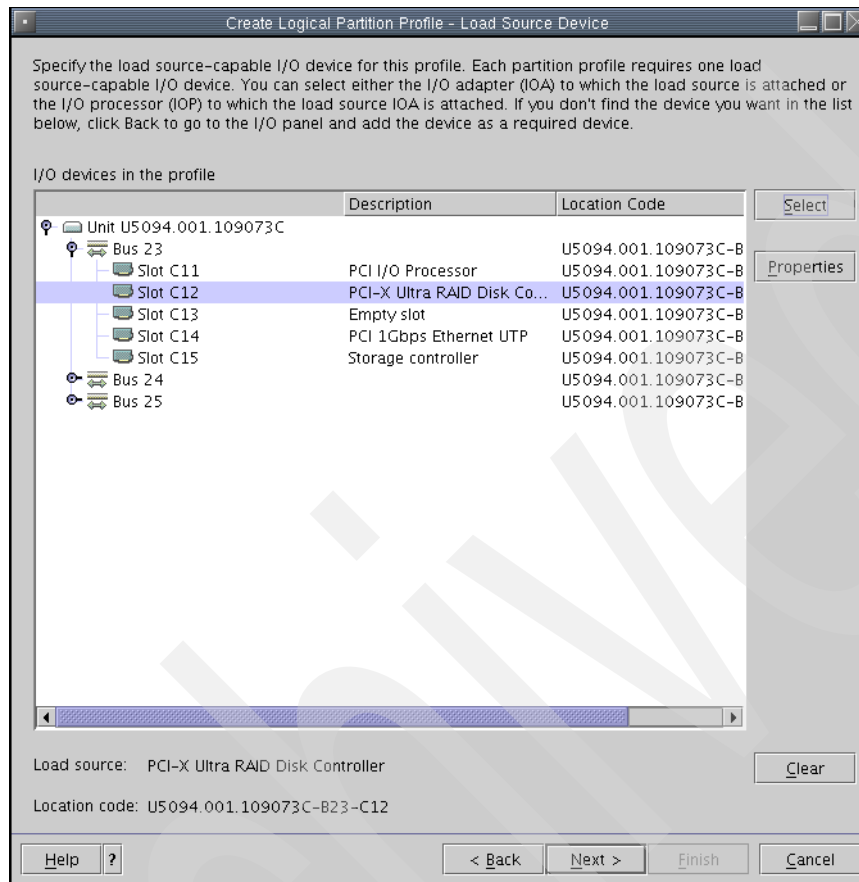


Figure 5-24 Selecting the Alternate IPL resource

Operation console device

Selecting the operations console device resource is optional, this was previously called Electronic Customer Support (ESC). See Figure 5-25.

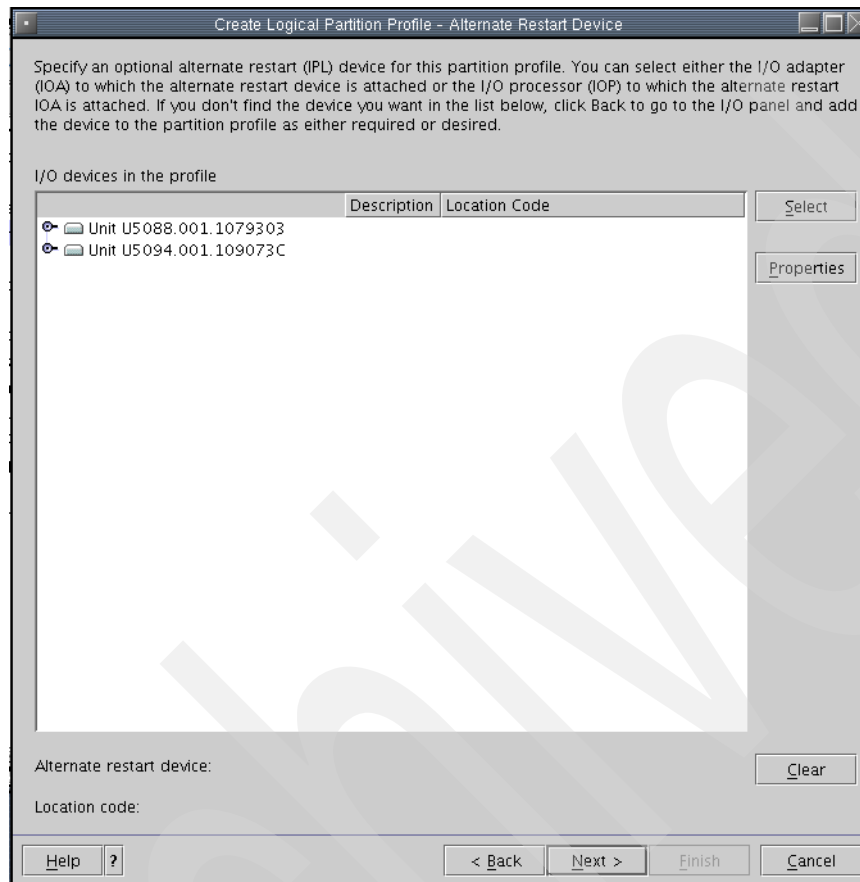


Figure 5-25 Selecting the Operations console device

Some support functions, such as RSSF (Remote Service and Support Facility), require that an ECS resource is selected. Additional information regarding RSSF can be found in the registered software knowledge base at:

<https://techsupport.services.ibm.com/as400.world/requestir?action=login>

Once inside the registered knowledge base, follow this path:

Software Knowledge Base Documents → Remote Support → RSSF → General

Console

The console provides a display (Figure 5-26) to interact with the partition. Certain functions, such as full system saves and dedicated service tools (DSTs), need to be initiated at or from the console. Chapter 2, “i5/OS consoles under System i5” on page 31 has more information regarding the console options that are available on System i5 hardware.

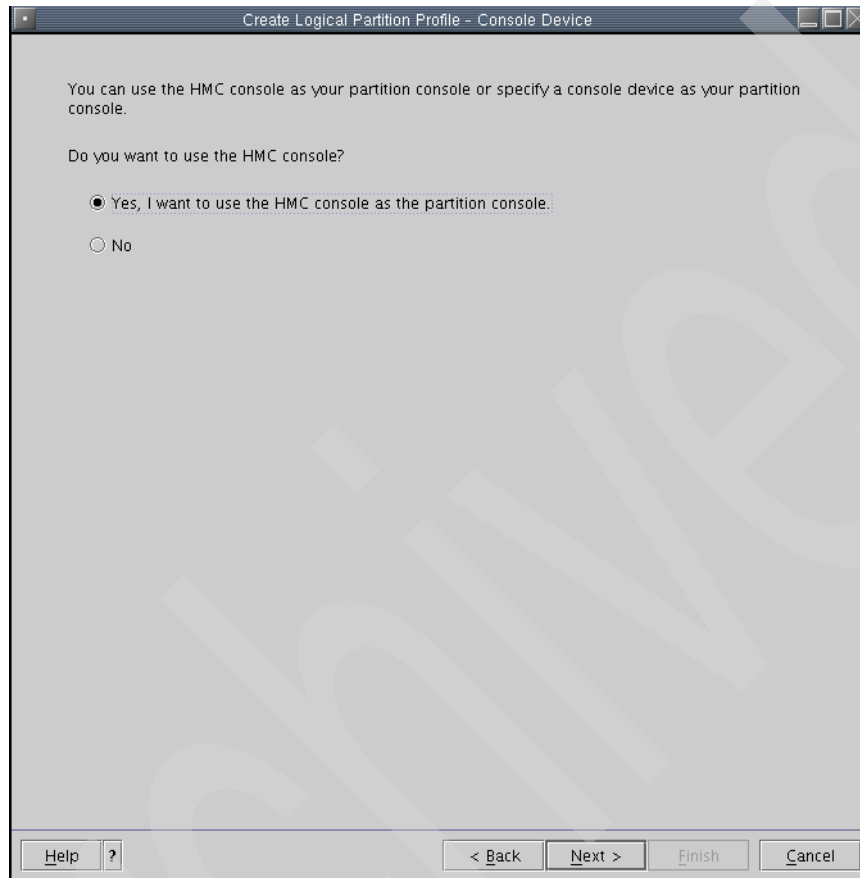


Figure 5-26 Selecting an HMC or different console

Using the HMC as the console

The default option is to use the HMC to provide console function for the partition. If this is desired, click **Next >** to advance to the next panel.

Specifying console resource

If some device other than the HMC is to provide console function, select the radio button labeled **No, I want to specify a console device** and click **Next >** as shown in Figure 5-27.

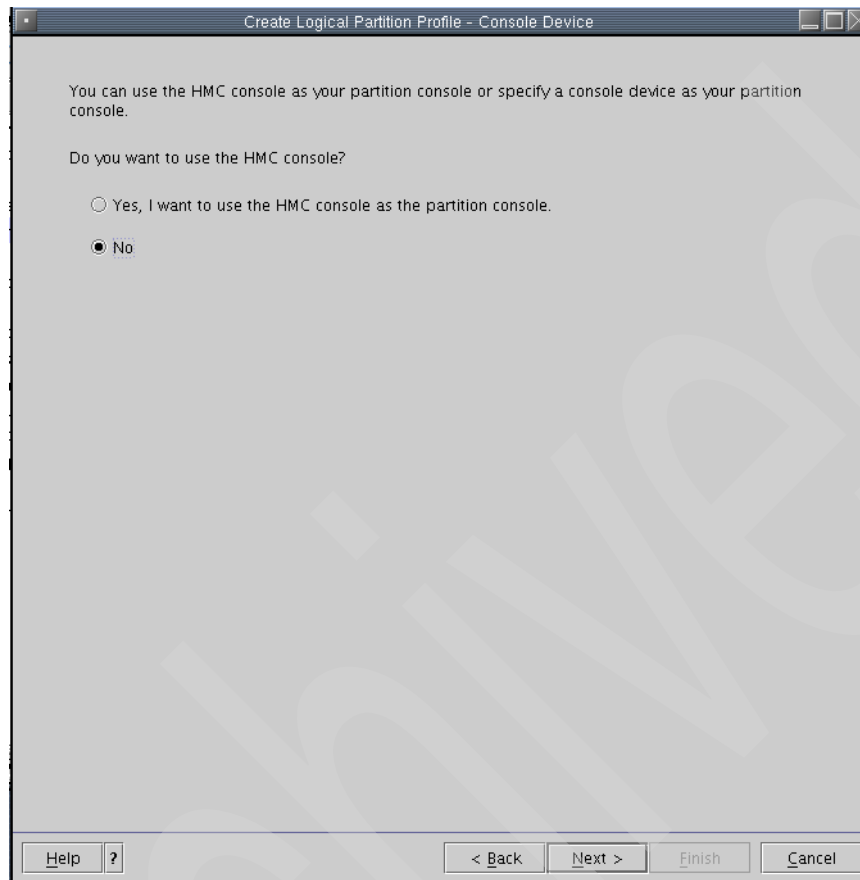


Figure 5-27 Selecting a console other than the HMC

As when selecting other partition resources, you are now presented with a dialog similar to the one in Figure 5-28.

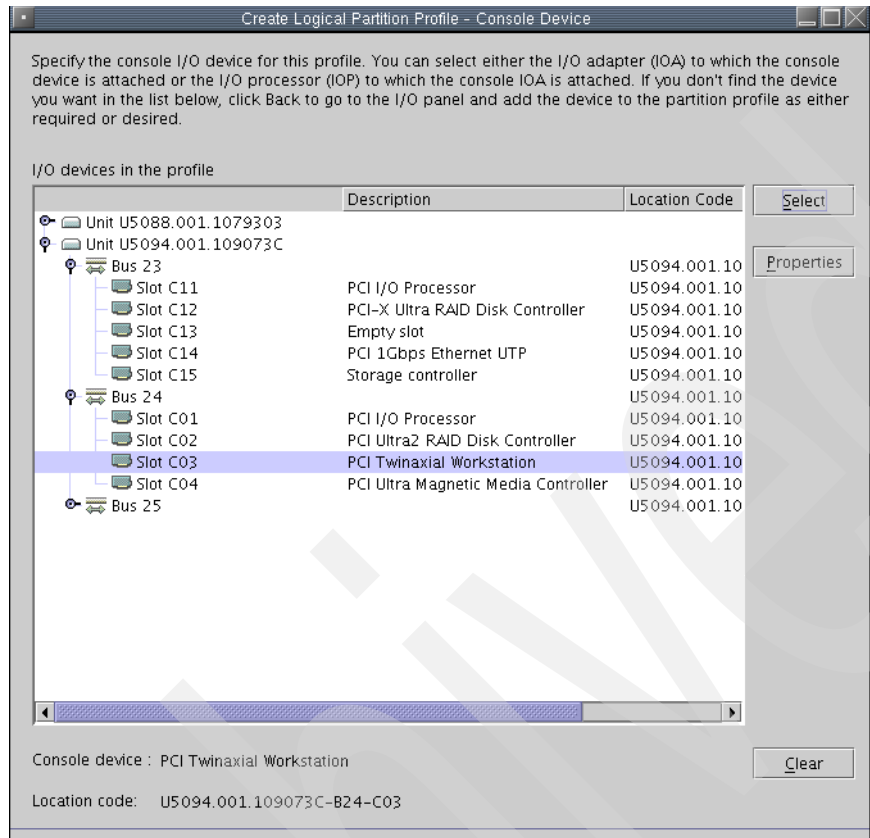


Figure 5-28 Selecting the console resource

Alternate console

An alternate console can provide console functions if the primary console is not functioning or not available. Some functions such as operating system install cannot be performed from an alternate console. Selecting an alternate console for a partition profile is optional. See Figure 5-29.

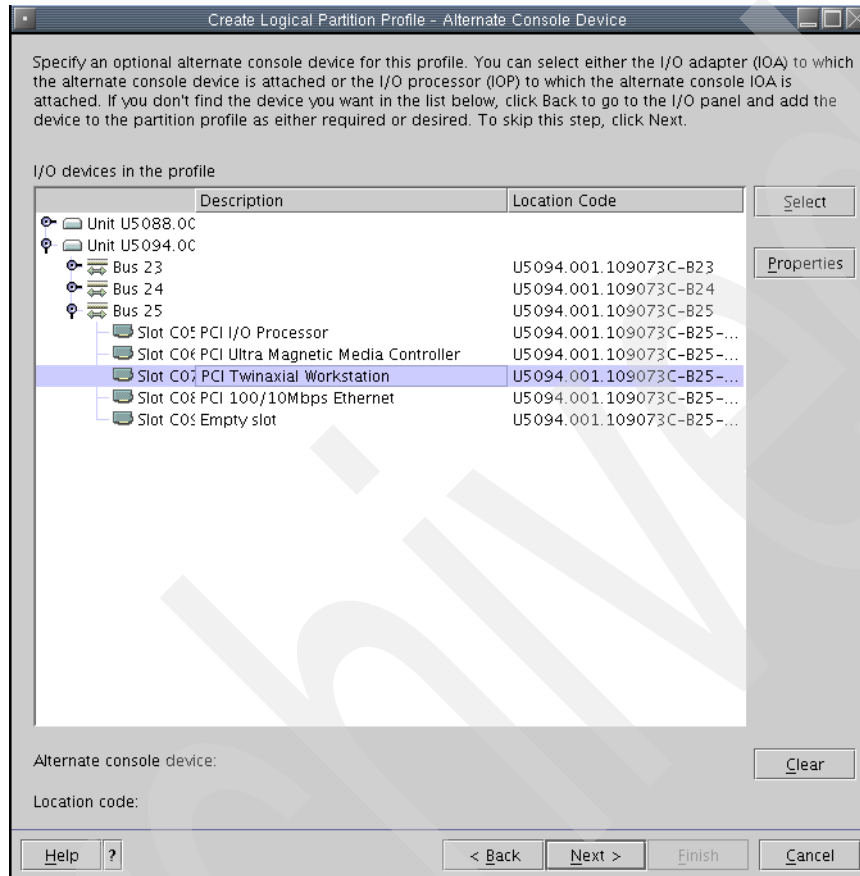


Figure 5-29 Selecting the alternate console resource

5.2.12 Opticonnect

If the partition profile will be using either **Opticonnect** or **HSL Opticonnect**, this can be specified by selecting the appropriate check box as shown in Figure 5-30.

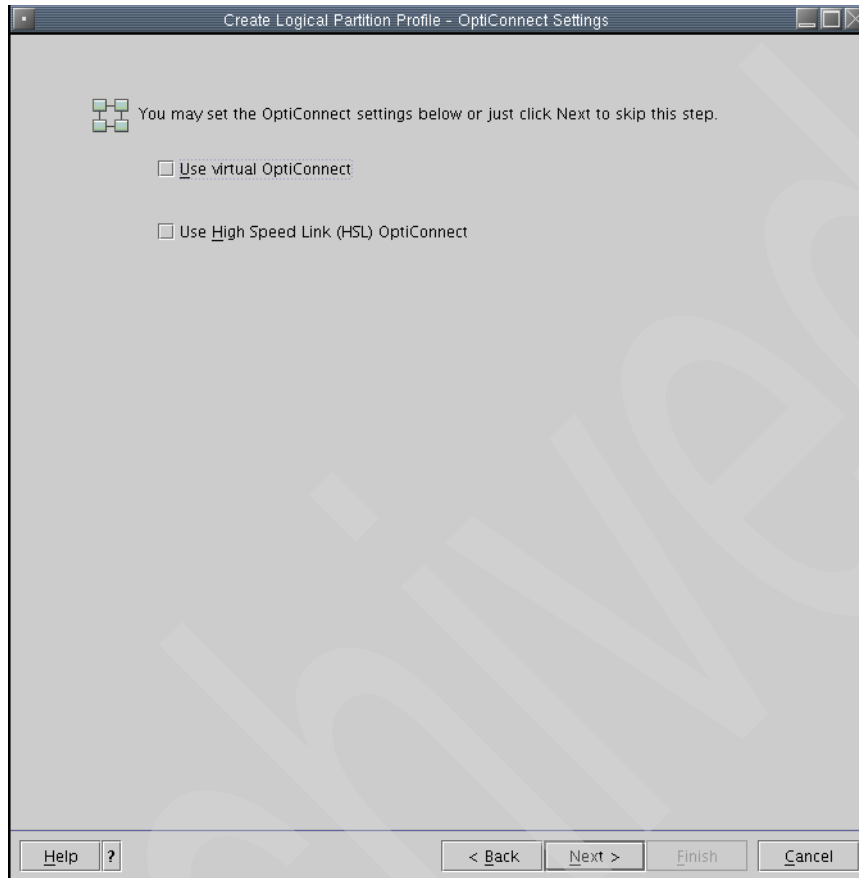


Figure 5-30 Specifying the opticonnect participation

5.2.13 Specifying power control partitions

For a hosted guest partition, the HMC by default gets power control, the ability to power the partition on and off. Specifying a power control partition allows another partition to have the same capability. Click **Add** to add another partition to the power control list for the partition, and click **Next >** to advance to the next panel (Figure 5-31).

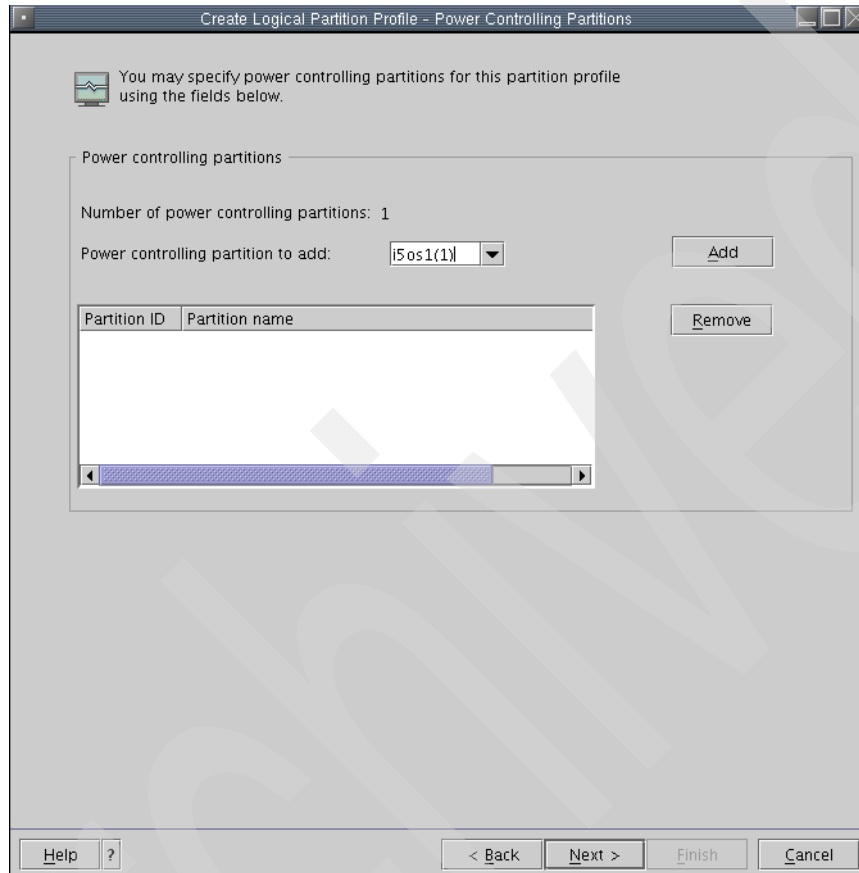


Figure 5-31 Specifying power control partitions

5.2.14 Miscellaneous profile settings

Automatically boot when managed system is powered on: This option allows for a partition profile to automatically IPL when the managed system is powered on from a powered off state or re-IPL'd. It is similar to the System IPL Action in preSystem i5 LPAR (Figure 5-32).

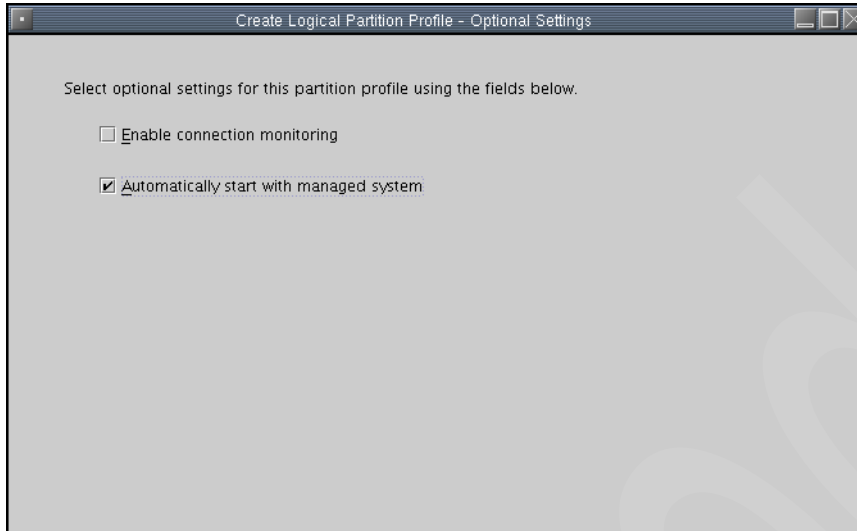


Figure 5-32 Miscellaneous partition profile settings

5.2.15 Review profile summary

Before the partition is created, the profile is displayed for final review as shown in Figure 5-33. If no changes are required, click **Finish** to have the partition profile created. Otherwise, use the **< Back** button to find the desired panel and make the required changes.

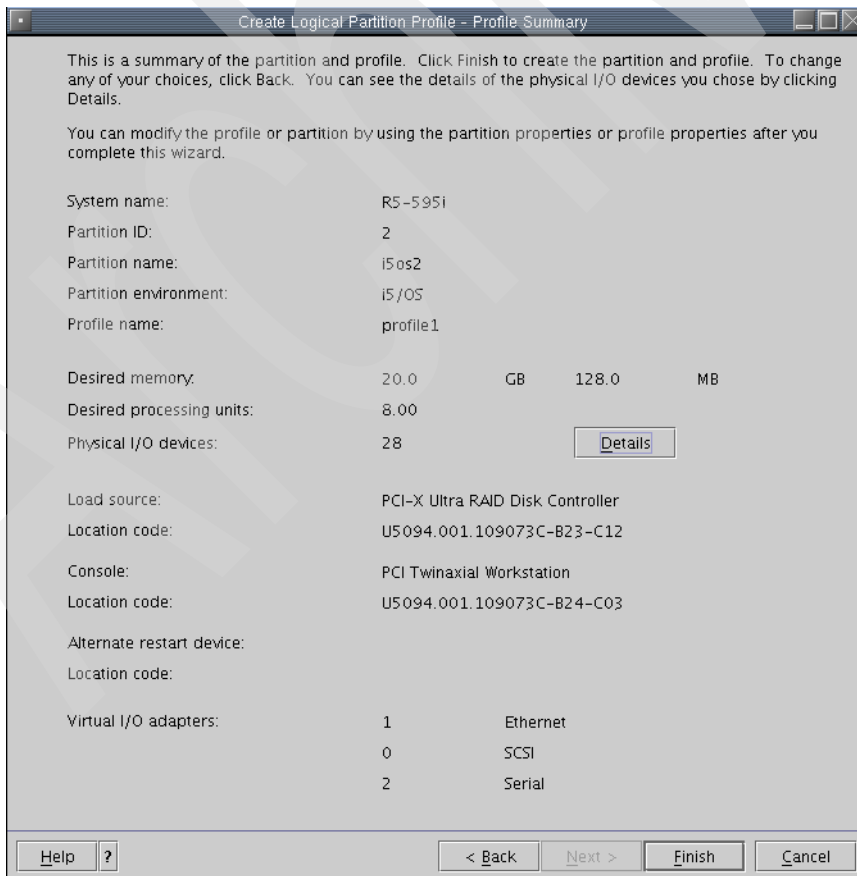


Figure 5-33 Partition profile is displayed for review

5.2.16 New partition profile has been created

Figure 5-34 shows the partition and partition profile. This partition is brand new and requires that an operating system be installed before the partition can be functional.

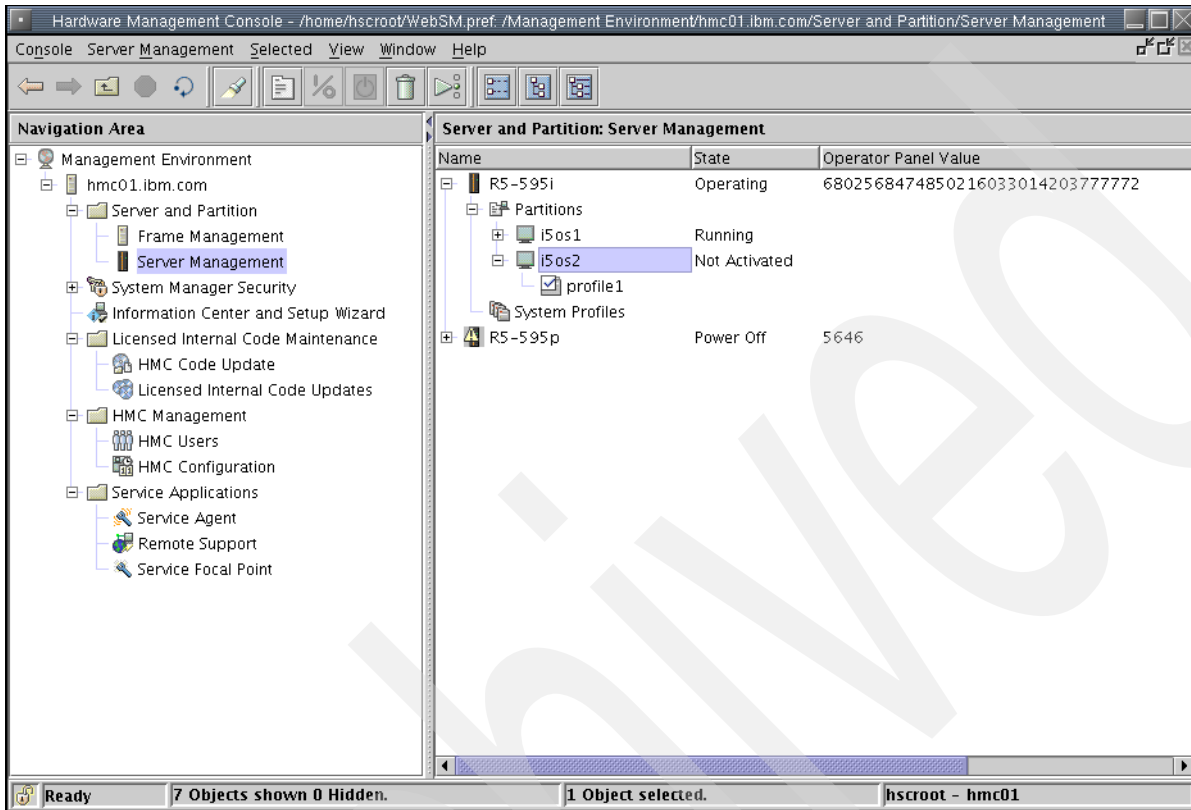


Figure 5-34 The new partition profile has been created

5.3 Creating another profile for an existing partition

Figure 5-35 shows the flows for creating a partition profile for an existing logical partition. This is similar to creating a brand new logical partition. Select the partition and right-click. From the pop up menu, select:

Create → Profile

The Create Partition Profile wizard loads and guides you through the remainder of the process.

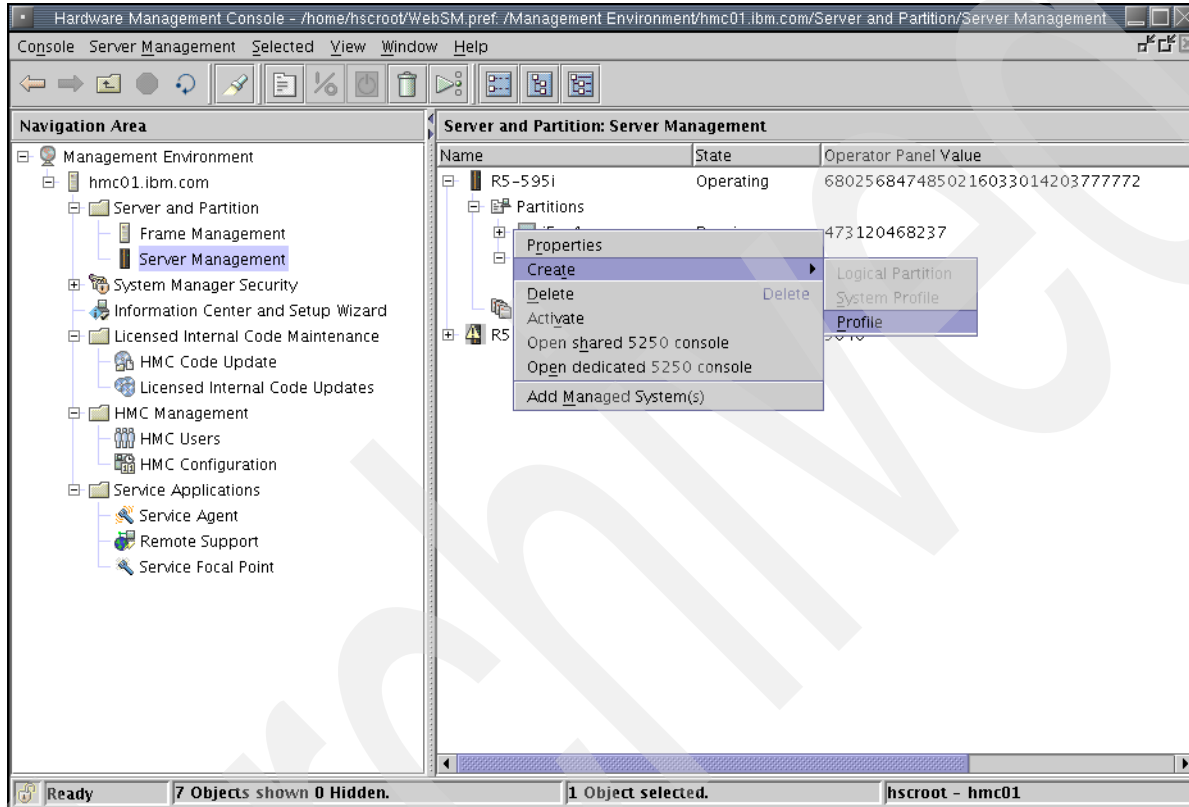


Figure 5-35 Starting the create profile process for an existing logical partition

Each partition profile needs to have a name. After that, the creation of another partition profile follows the same flow as presented earlier (Figure 5-36).

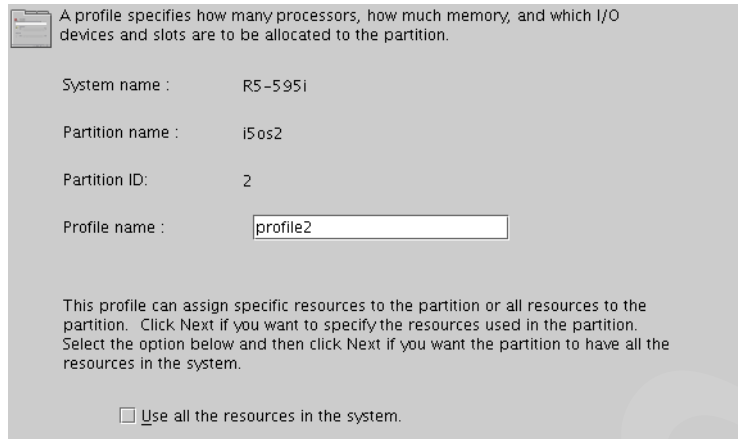



Figure 5-36 Specifying a name for the partition profile

Figure 5-37 shows a logical partition that has more than one partition profile. The  icon indicates which profile for the logical partition is the default profile. The default profile specifies which profile is automatically started when a partition is activated.

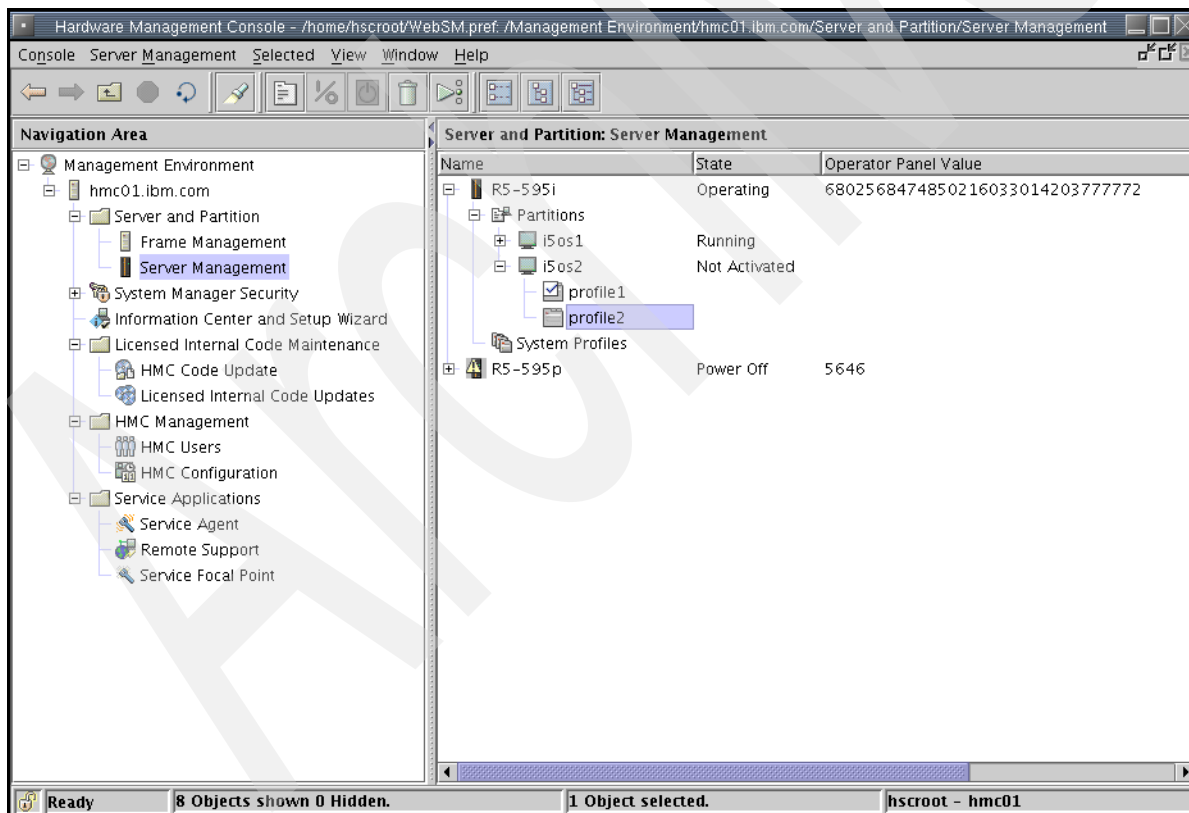


Figure 5-37 The second partition profile has been created

5.4 Changing the default profile for a partition

Changing the default profile is a relatively straightforward operation. From the **Servers and Partitions: Server Management** pane, right-click the desired partition. From the pop up menu, select **Change Default Profile** (Figure 5-38).

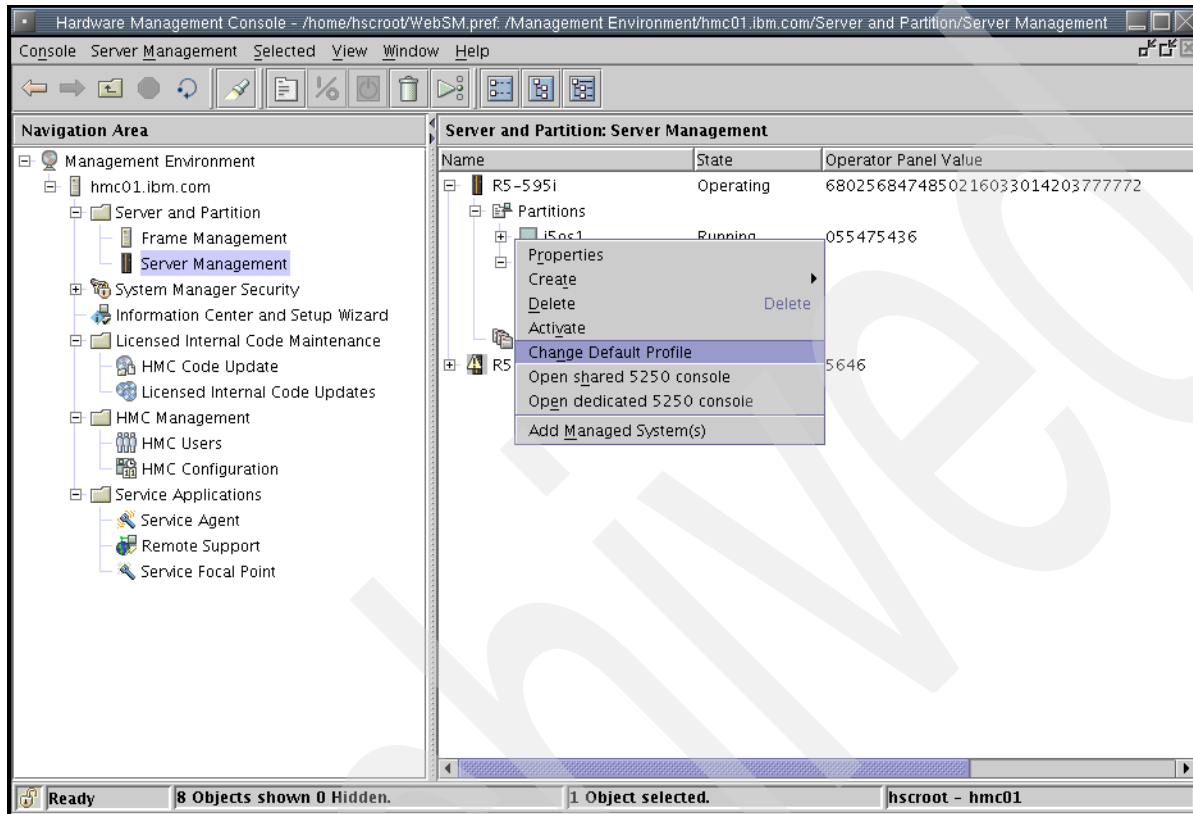


Figure 5-38 Changing the default profile for a partition

Figure 5-39 shows a dialog box that displays the list of profiles that are associated with the selected partition. From the drop down menu, select the desired partition profile and click **OK**.

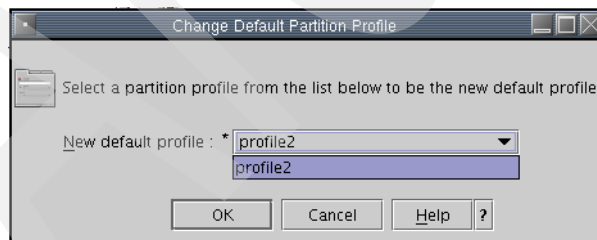


Figure 5-39 Selecting a new default profile

Archived

Working with the HMC

In addition to providing an interface for creating logical partitions, the Hardware Management Console (HMC) is used to manage logical partitions once they are created.

In this chapter, we discuss the following topics:

- ▶ Accessing LPAR functions
- ▶ Viewing partition properties
- ▶ Starting and stopping partitions
- ▶ Performing dynamic LPAR (DLPAR) functions
- ▶ Performing service functions on partitions
- ▶ Remote management:
 - HMC CLI
 - Web SM
- ▶ Cross reference of iSeries control panel functions to HMC options

6.1 Accessing LPAR functions

Attention: An HMC can manage multiple servers. Before performing an operation on a logical partition, make sure that the correct server is selected.

On the HMC, LPAR functions for a logical partition can be accessed in two main ways. Throughout this chapter, the menu will be referred to as the LPAR functions menu.

Pop up menu

In the navigation area, select the managed system and drill down as follows:

Server and Partition → Server Management

The **Server and Partition: Server Management** pane is displayed. Expand the **Partitions** folder, select the desired partition, and right-click. The pop up menu that is displayed is shown in Figure 6-1.

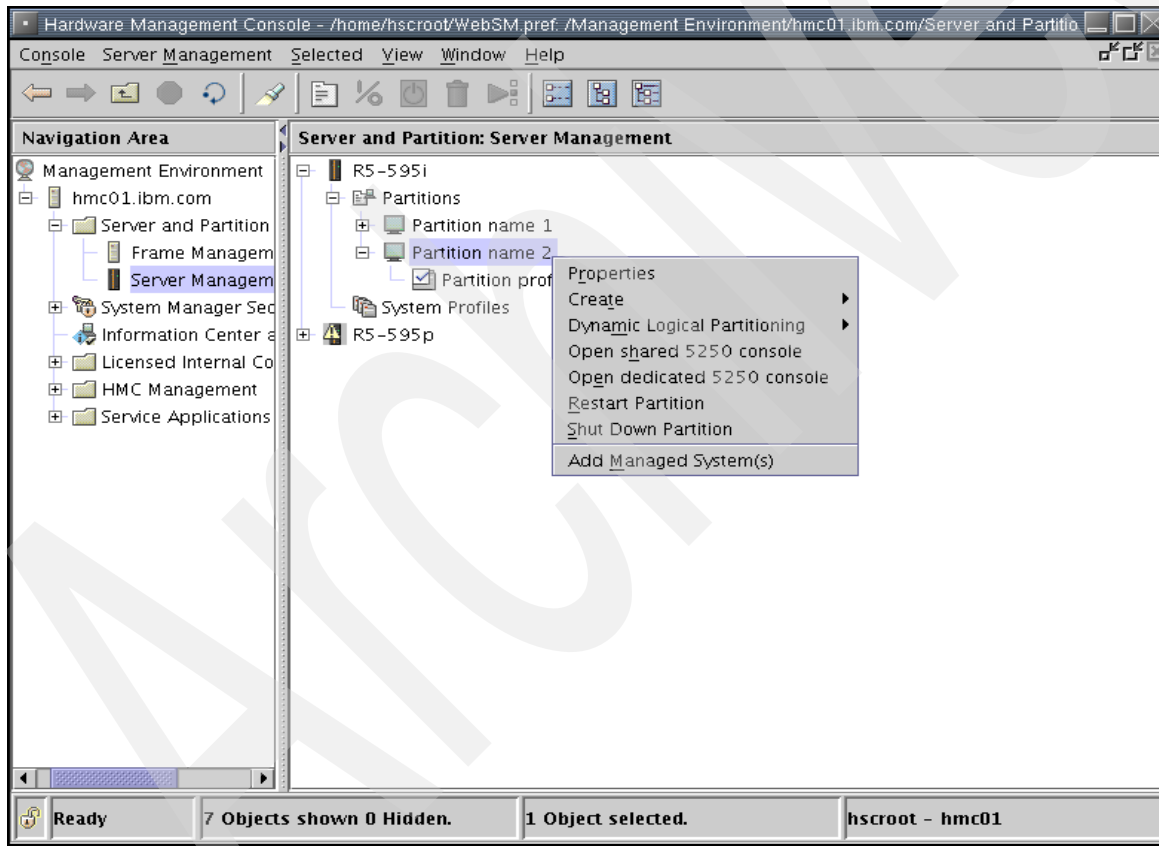


Figure 6-1 Accessing LPAR functions via the pop up (context) menu

Selected menu bar item

Select the desired partition using the same method as in "Pop up menu" on page 178. Instead of right-clicking the partition, from the menu bar, use the **Selected** menu item. The resulting menu items are shown in Figure 6-2.

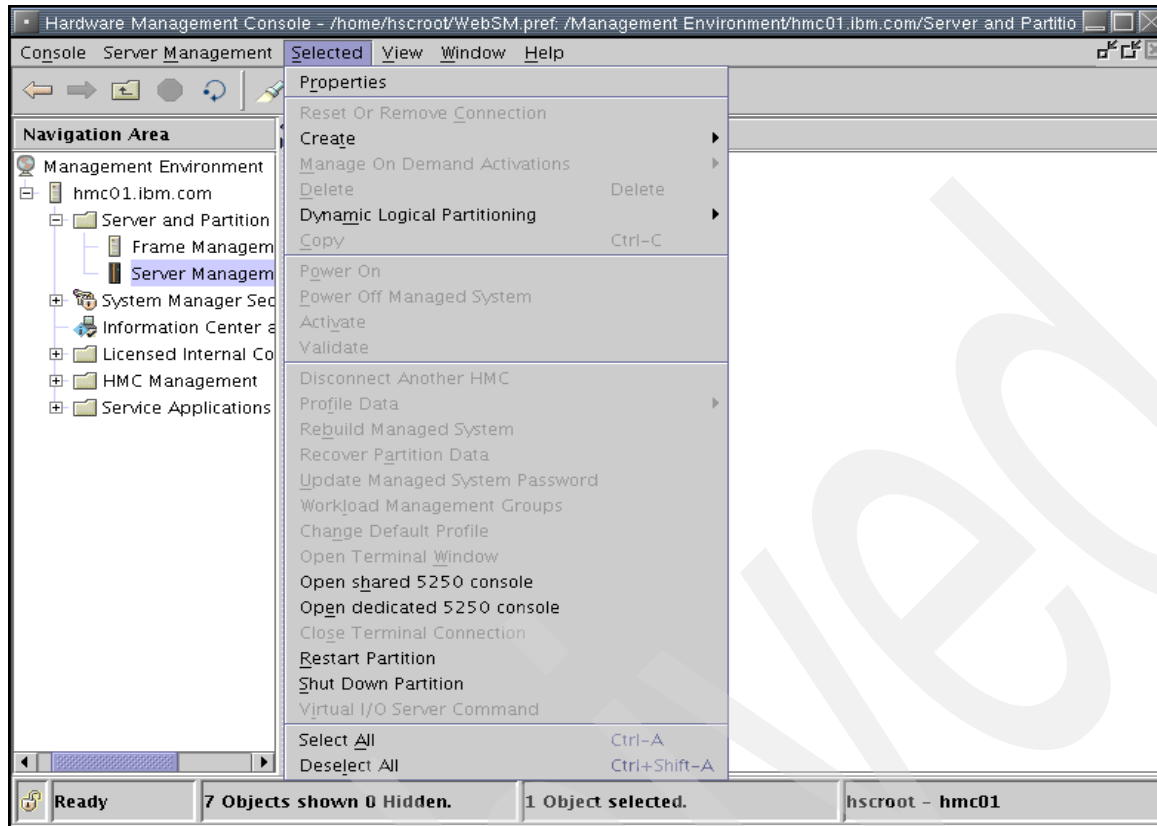


Figure 6-2 Accessing LPAR functions via the Selected menu bar item

6.2 Viewing partition properties

Some of the information concerning a partition is common to the properties of both the partition and the partition profile. In other cases, specific information can only be found in one of the two places.

6.2.1 Partition properties

To access the properties for a partition, access the LPAR functions menu and select **Properties**. Ceasing the LPAR functions menu is discussed in 6.1, “Accessing LPAR functions” on page 178.

General

As shown in Figure 6-3, the **General** tab displays basic information about a logical partition. The partition **Name** and **ID** are values that were assigned during partition creation. The partition **Environment** identifies what type of operating system is installed on the partition, and the **OS version** identifies which release of that operating system.

If the partition is active, IPL'd, the **Current profile** reflects which partition profile was specified when the partition was activated. The **System** value reflects the machine type/model and serial number of the managed system that contains this partition. Partition status is reflected in the **State** of the partition. Some common states are shown in Table 6-1.

Table 6-1 Partition states and their meaning

State	Description
Off	The partition is powered off.
Power On	The partition is in the process of powering on.
On	The partition's operating system is running.
Power Off	The partition is in the process of powering off.
Failed	The partition has encountered an error in the early IPL path.
Unit Attention	The partition encountered a run time failure. Check the reference code for the partition and take the appropriate action.

The **Resource configuration** value indicates if the partition profile is bootable. Possible values are:

- ▶ Configured (bootable)
- ▶ Not configured (not-bootable)

A profile is considered not bootable when a critical resource is not present. Examples would include insufficient processors and memory to meet the minimum requirements. Also, if a partition is powered off and not part of a workload management group, the Hypervisor can steal resources that are needed by other partitions. So, if a critical IO resource is no longer present, due to a hardware failure or having been removed, the profile is not bootable as well (Figure 6-4).

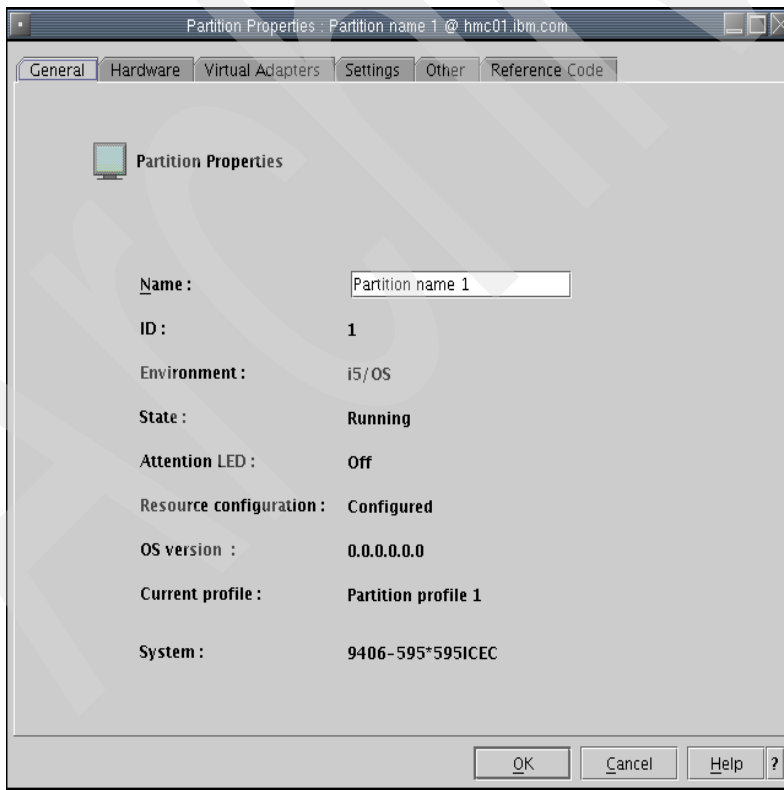


Figure 6-3 The General tab of the partition properties dialog

Hardware

The **Hardware** tab shows which hardware resources are currently in use by the partition. There are two sub-tabs on this dialog. One lists the I/O resources, and the other lists both processors and memory.

I/O

The **I/O** sub-tab (Figure 6-4) shows which I/O resources are currently assigned to the partition. The hierarchical view can be expanded to display which buses within a given unit are available to a partition. An individual system bus can be further expanded to display which individual IO slots are allocated to the partition.

Information about the type and model of a particular card residing in a given slot is only displayed if the converged Hypervisor has informed the HMC. In order to obtain this information, the converged Hypervisor depends on being directed by the operating system that owns and uses that particular slot. The converged Hypervisor does not query the slot itself for this information (Figure 6-5).

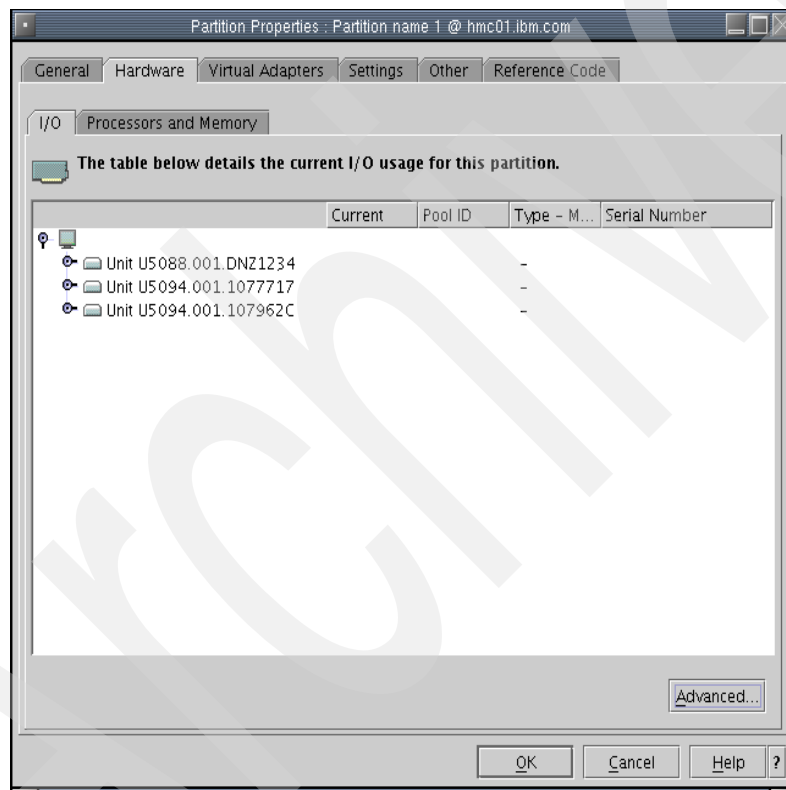


Figure 6-4 Hardware - I/O tab on the partition properties

Processors and Memory

The **Processors and Memory** sub-tab (Figure 6-5) displays the current processing configuration for the partition. The minimum, maximum, and current values for processing units and virtual processors is displayed. The processor **Sharing mode** identifies what kind of processors are being used: dedicated, capped shared, and uncapped shared. For an uncapped partition, the partition weight is also displayed.

Allow shared processor utilization authority indicates if the partition has the authority to view utilization information of the entire shared processor pool. Without this authority, the partition can only obtain shared processor pool information about itself.

Under **Memory**, the minimum, maximum, and current values are displayed.

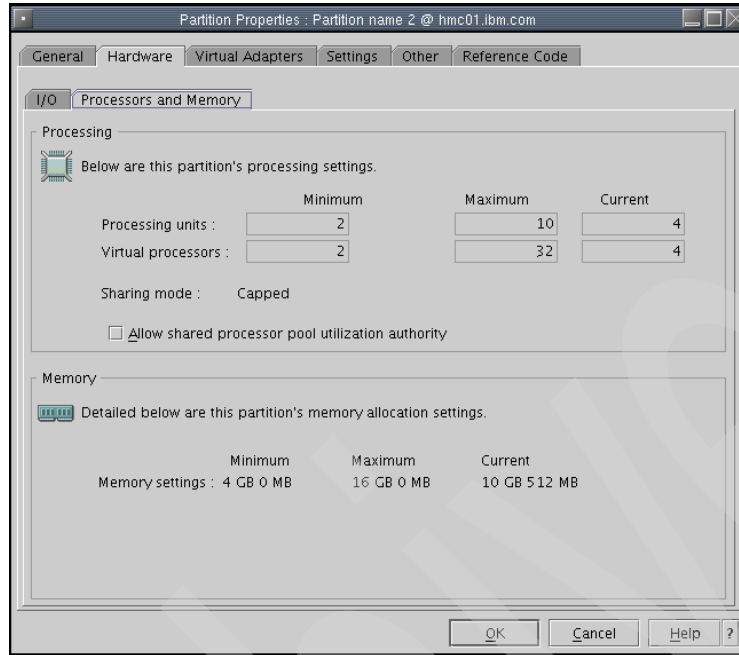


Figure 6-5 Hardware - Processors and Memory tab on the partition properties

Virtual Adapters

If the partition has any configured virtual I/O adapters, they are shown on the **Virtual Adapters** tab (Figure 6-6) of the partition properties dialog. There are separate sub-tabs for virtual ethernet, serial, and SCSI adapters.

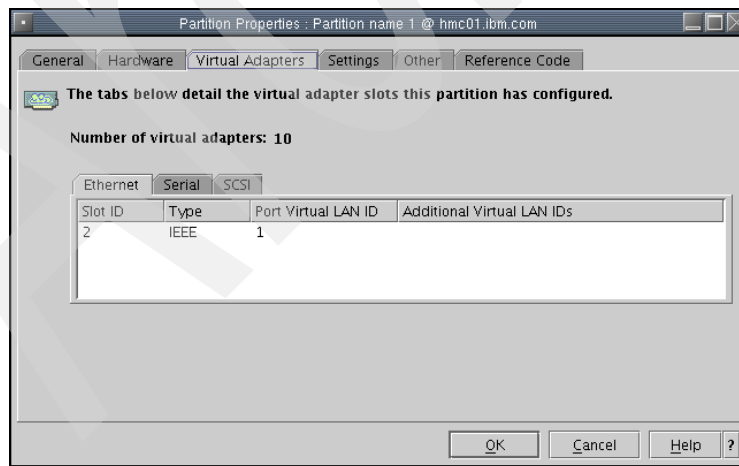


Figure 6-6 Virtual devices tab on partition properties dialog

Settings

The **Settings** tab (Figure 6-18 on page 192) shows information about partition boot, service support and tagged I/O.

For partition IPL source and mode settings, refer to 6.3.1, “Changing IPL side and mode” on page 191. **Automatically start with managed system** specifies if a partition to perform an IPL when the entire managed system is IPL'd. The default partition profile is started in this case. If unchecked, the partition would need to be manually started after the managed system is IPL'd.

Service partition setting indicates if this partition is identified as the service partition. In the absence of the HMC, the service partition can report software and hardware errors to IBM.

The **tagged I/O** section identifies which resources are marked for special purpose functions like the load source and console.

Other

The **Other** tab (Figure 6-7) shows information about this partition's participation in an eWLM group under the **Workload Management** section. If this partition provides power control for a hosted guest partition, that is displayed under the **Power Controlling Partitions** section. The **Communication** section shows information regarding HSL Opticonnect and Virtual Opticonnect.

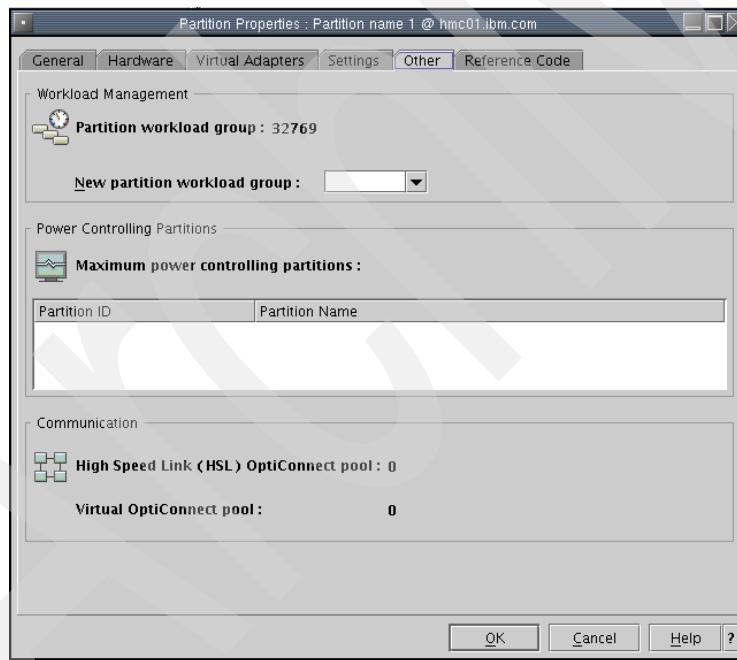


Figure 6-7 Other partition properties

Reference Code

For an explanation of the **Reference Code** tab, see 6.5.1, “Displaying reference code information” on page 199.

6.2.2 Partition profile properties

This section covers the properties available for partitions through the tabs at the top of each properties pane.

General

As with the General tab for partition properties in Figure 6-3 on page 180, the partition profiles **General** tab (Figure 6-8) displays some basic information about a partition profile. The **System Name** is the managed system that contains this partition and partition profile. **Partition ID**, **Partition Name**, and **Partition environment**, are identical to what is displayed on the General tab for partition properties (6.2.1, “Partition properties” on page 179). **Profile name** reflects the name of the selected profile.

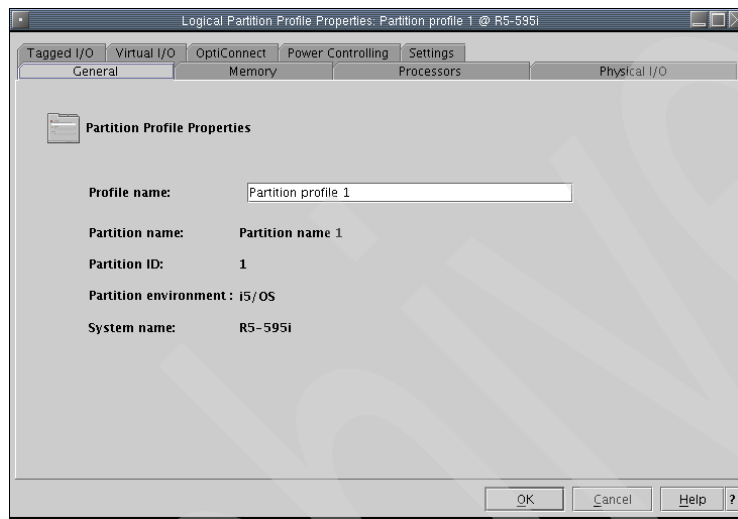


Figure 6-8 Partition profile General tab

Memory

The **Memory** tab (Figure 6-9) displays information about the profile’s memory configuration. If desired, changes to the profile’s memory configuration can be made here. These are not dynamic LPAR (DLPAR) changes and will take affect when this profile is next restarted. For DLPAR memory changes, refer to the memory portion of 6.4, “Performing dynamic LPAR (DLPAR) functions” on page 195.

The concepts of minimum, desired, and maximum memory are discussed in 5.2.6, “Partition memory” on page 148.

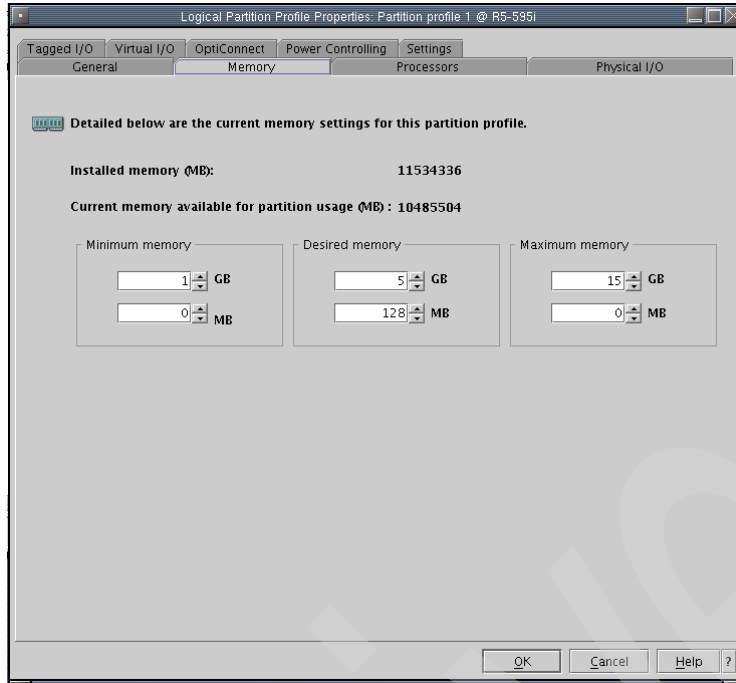


Figure 6-9 Partition profile Memory tab

Processors

The **Processors** tab displays information about the profile's processor configuration. There are two distinct views, depending on processing mode: Dedicated (Figure 6-10) and Shared (Figure 6-11).

Depending on the processing mode, either **Total managed system processors** or **Total managed system processing units** reflects the total processing capacity that the physical system can provide.

If desired, changes to the profile's processing configuration can be made here. These are not dynamic LPAR (DLPAR) changes and will take effect when this profile is next restarted. For DLPAR processor changes, refer to the processor portion of 6.4, "Performing dynamic LPAR (DLPAR) functions" on page 195.

The concepts of minimum, desired, and maximum processors are discussed in 5.2.7, "Partition processors" on page 149. For an uncapped shared processor partition, the partition weight can also be adjusted. The distinctions between capped shared and uncapped shared are discussed in Chapter 1, "Introduction to LPAR on IBM System i5" on page 1.

Allow idle processors to be shared specifies if a partition with dedicated processors should have those processors available for use by an uncapped partition. The sharing of the dedicated processors could only occur if the partition with the dedicated processors was powered off or otherwise unavailable. In the majority of cases, this box should be checked.

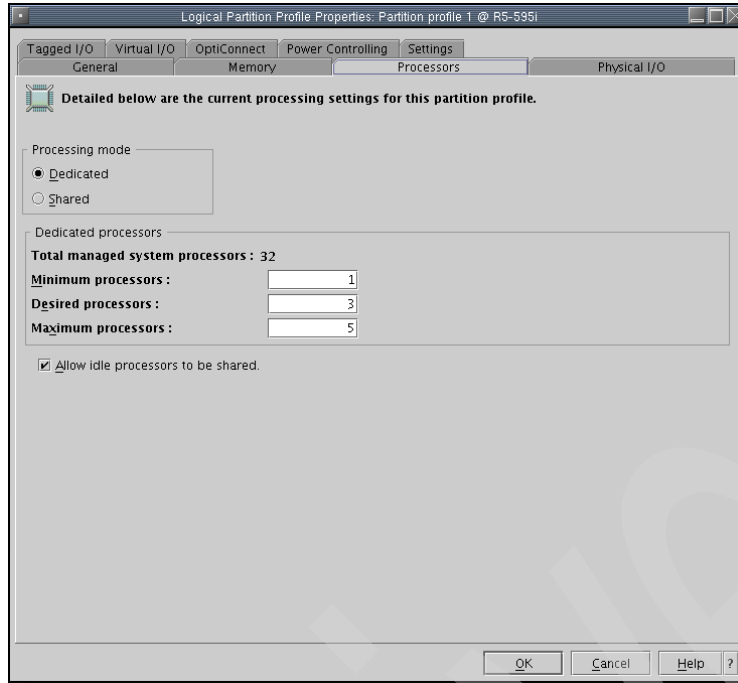


Figure 6-10 Partition profile Processors tab for dedicated processors

The processor properties of a shared processor partition profile is somewhat different. As with dedicated processors, changes to the profile's processing configuration can be made here. These are not dynamic LPAR (DLPAR) changes and will take affect when this profile is next restarted. For DLPAR processor changes, refer to the processor portion of 6.4, "Performing dynamic LPAR (DLPAR) functions" on page 195.

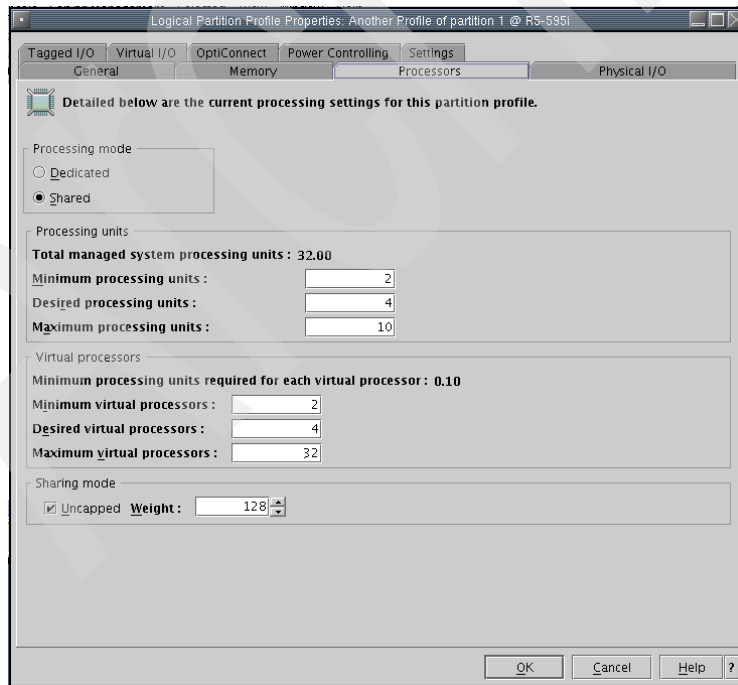


Figure 6-11 Partition profile Processors tab for shared processors

Physical I/O

The **Physical I/O** tab (Figure 6-12) identifies what physical I/O is available on the entire system and what is assigned to profile. Changes to the profile's allocation of physical I/O adapters can be performed here.

These are not dynamic LPAR (DLPAR) changes and will take affect only when this profile is next restarted. For DLPAR physical adapter changes, refer to the physical adapter portion of 6.4, "Performing dynamic LPAR (DLPAR) functions" on page 195.

Adding a resource as required or desired is discussed in 5.2.9, "Allocate physical I/O resources" on page 155.

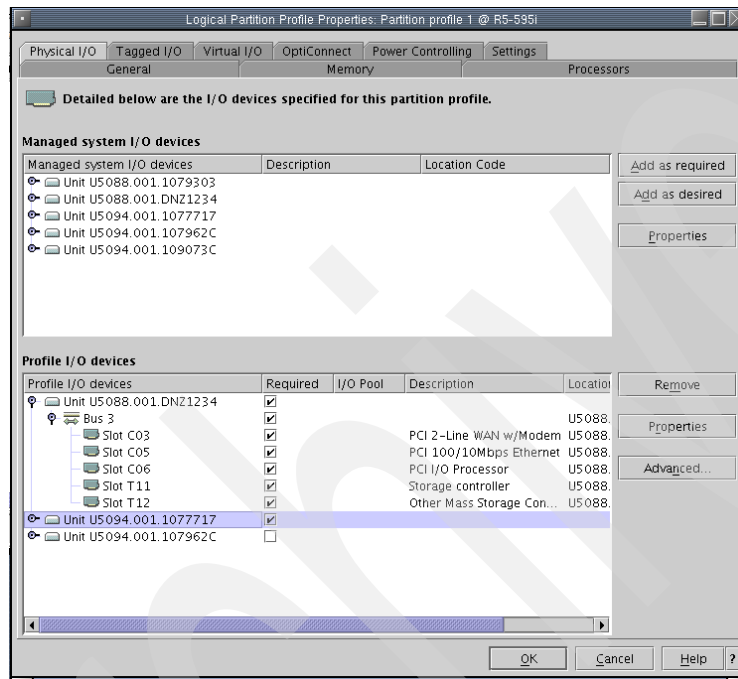


Figure 6-12 Partition profile Physical I/O tab

Tagged I/O resources

The **Tagged I/O** tab (Figure 6-13) identifies which resources are selected to perform partition critical functions. Some of these, like load source and console, are required for the profile to start. Others, such as alternate restart device, alternate console, and operations console, are only required in particular circumstances.

Any of these resources can be changed by clicking **Select** and choosing a new resource. These changes take effect the next time the profile is restarted.

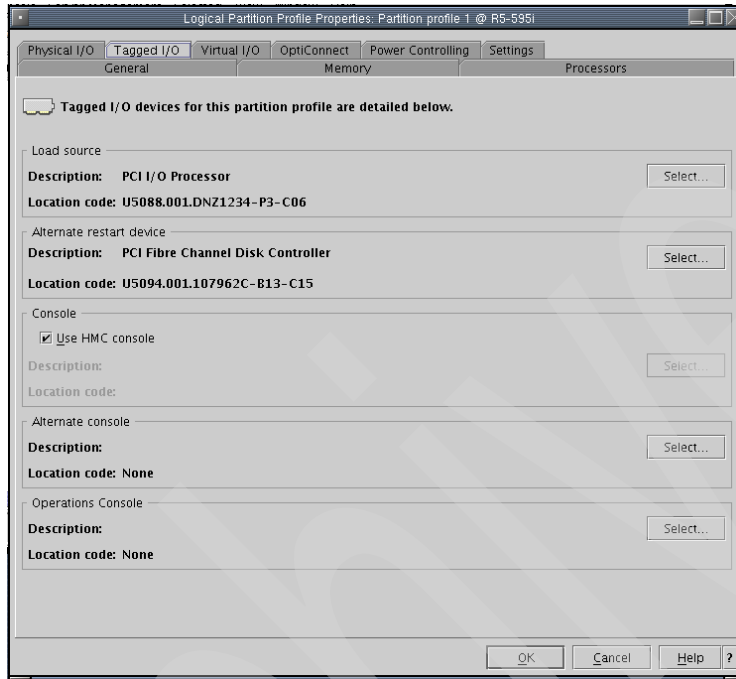


Figure 6-13 Tagged I/O devices defined in a partition profile

Virtual I/O resources

For details on the **Virtual I/O** tab, see 6.4.4, “Virtual IO adapters” on page 197.

Opticonnect

The **Opticonnect** tab displays the profile’s opticonnect settings (Figure 6-14). Changing the virtual opticonnect setting takes effect on the next restart of the profile.

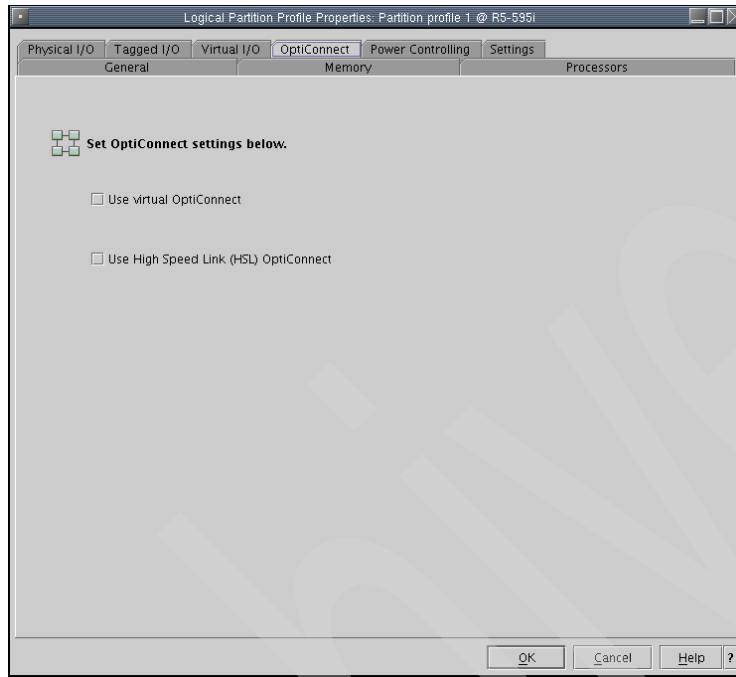


Figure 6-14 Profile opticonnect settings

Power Controlling

The **Power Controlling** tab (Figure 6-15) shows if this partition profile has power control for a hosted guest partition. Guest partitions can be removed from the list and added up to the maximum.

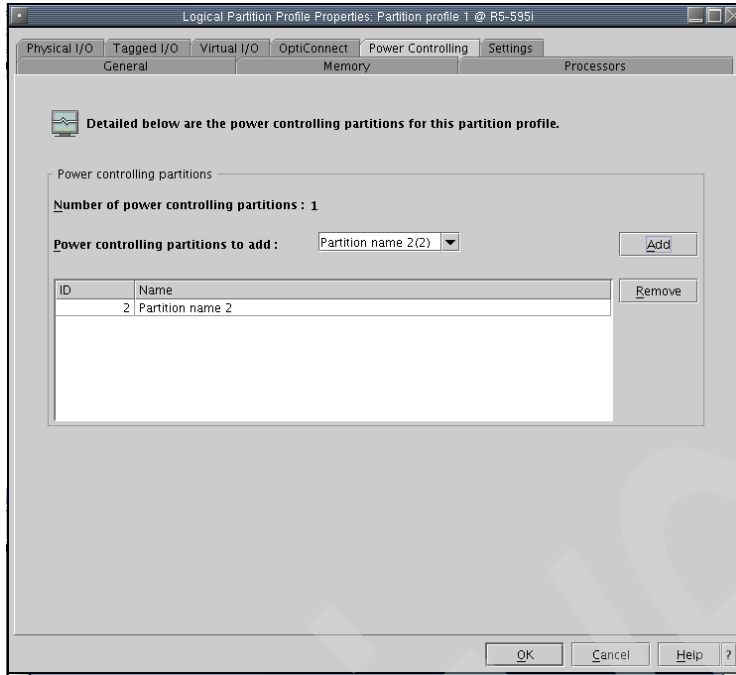


Figure 6-15 Partition profile Power Controlling tab

Settings

The **Settings** tab (Figure 6-16) of the partition profile dialog has the following options for a partition profile:

- ▶ Enable Service connection monitoring
- ▶ Automatically start when the managed system is powered on




Figure 6-16 Partition profile Settings tab

6.3 Starting and stopping partitions

One of the more common partition tasks involves starting and powering down partitions. In this section we discuss the following tasks:

- ▶ Changing the IPL side and mode for a logical partition
- ▶ Manually starting a powered off partition
- ▶ Restarting an operational partition
- ▶ Powering down a partition

6.3.1 Changing IPL side and mode

In order to change the IPL side (A,B, or D) and mode (manual or normal) for a logical partition, perform the following steps:

1. Access the LPAR functions menu using one of the methods in 6.1, “Accessing LPAR functions” on page 178.
2. From the LPAR function menu, select **Properties**. You will see a dialog similar to the one in Figure 6-17.

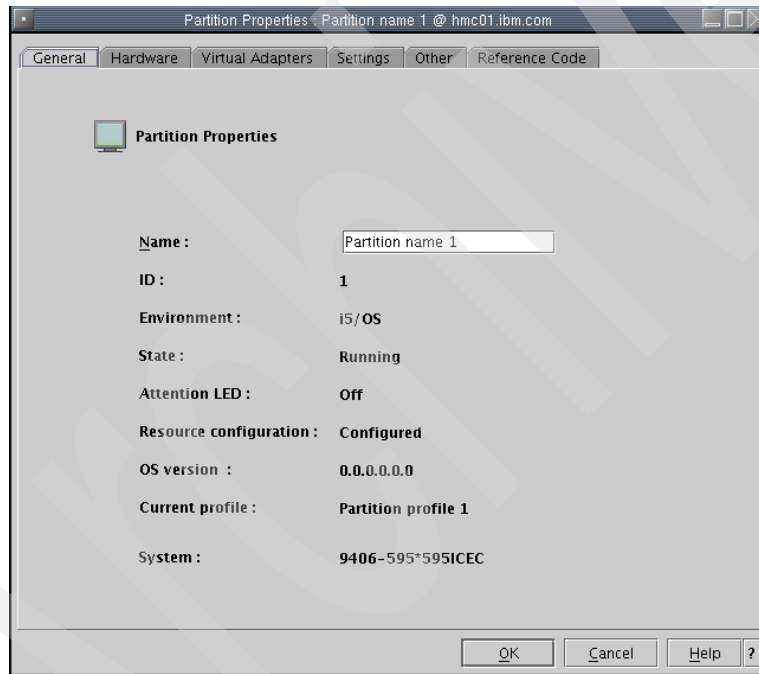


Figure 6-17 Partition properties dialog

3. Select the **Settings** tab. This dialog is shown in Figure 6-18.

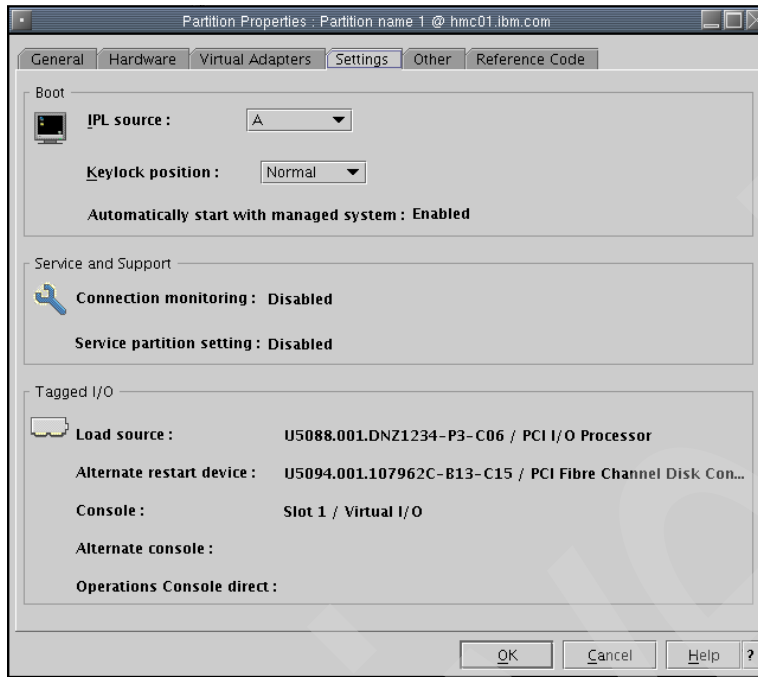


Figure 6-18 Partition settings

4. In the **Boot** section of the dialog, the **IPL Source** and **Keylock position** can be changed from their respective pull down menus.
5. Click **OK** once the desired changes have been made.

6.3.2 Starting a powered off partition

To IPL a partition involves activating that partition profile:

1. If needed, set the desired IPL side and mode as referenced in 6.3.1, “Changing IPL side and mode” on page 191.
2. From the LPAR functions menu, select **Activate**. A dialog like the one in Figure 6-19 is shown.
3. From the list of partition profiles shown, select the one to activate.
4. After the desired profile is selected, press **OK** to start the activation process. See Figure 6-19.

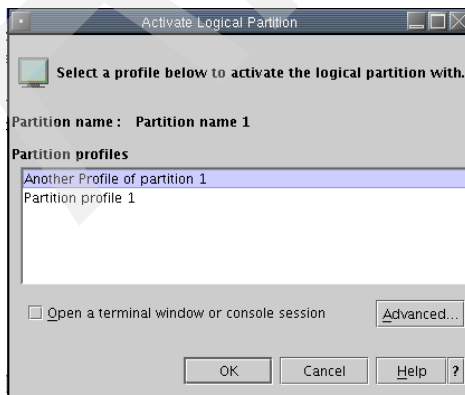


Figure 6-19 Selecting the profile to activate

Note: Activating a second profile for an active partition is not allowed and results in an error. To switch profiles for an active partition, the first profile needs to be deactivated (powered off) before the second profile can be activated (powered on).

6.3.3 Restarting a running partition

The preferred method of restarting an active partition is to issue the following command:

```
PWRDWSYS OPTION(*CNTRLD) DELAY(user specified delay) RESTART(*YES)
```

The default delay time for a controlled power down is 3600 seconds or one hour. Depending on the particular circumstances, this value may need to be changed. A logical partition can also be restarted using the HMC GUI:

1. Access the LPAR functions menu using one of the methods in 6.1, “Accessing LPAR functions” on page 178.
2. Select **Restart Partition**. A dialog like the one in Figure 6-20 on page 193 is shown.
3. Several options are listed for reboot type. Immediate should rarely be used, if at all. **Dump** and **Dump Retry** are discussed in 6.5.3, “Performing main storage dumps” on page 200.

Attention: Except under the direction of your next level of support, the use of the Immediate reboot option is discouraged. This is an abnormal end of the system. Improper use can result in an extremely long IPL and potentially in damaged objects.

4. Select the restart type and press **OK** to start. See Figure 6-20.

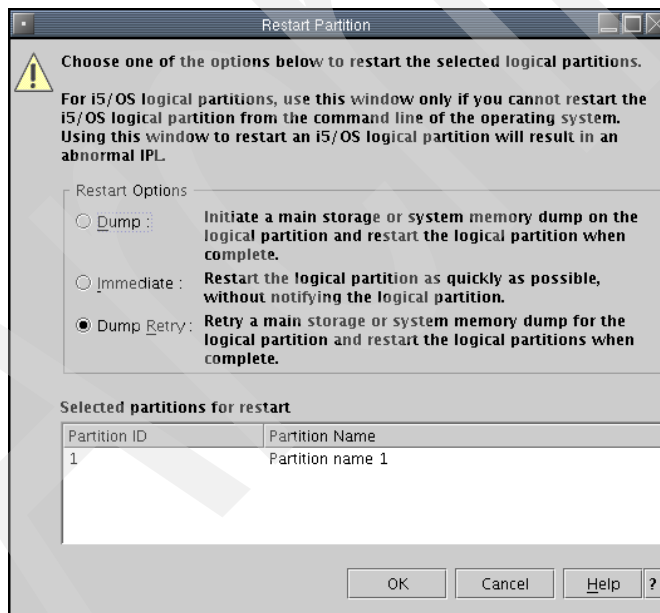


Figure 6-20 Partition restart options

6.3.4 Stopping (powering down) a running partition

The preferred method of powering down an active partition is to issue the following command.

```
PWRDWN SYS OPTION(*CNTRLD) DELAY(user specified delay) RESTART(*NO)
```

The default delay time for a controlled power down is 3600 seconds or one hour. Depending on the particular circumstances, this value may need to be changed.

In cases where a command line is not available due to a partition loop or wait state, the partition can be brought down as follows.

Tip: In the case of either a partition loop or wait, it is highly recommended to capture a main storage dump (MSD) of the partition. It is not required, but without a MSD, problem diagnosis and resolution is often severely restricted.

Refer to 6.5.3, “Performing main storage dumps” on page 200 for additional information.

1. Access the LPAR functions menu using one of the methods in 6.1, “Accessing LPAR functions” on page 178.
2. Select **Shut Down Partitions**. A dialog like the one in Figure 6-21 is shown.
3. Choose the desired shutdown type:
 - **Delayed:** This option is delayed power off and is the same as pressing the power button twice on a stand-alone machine or taking option 7 from the virtual panel. The operating system does not have an opportunity to shut down, but changed pages are written out to disk.
 - **Immediate:** This option is an immediate power down and is the same as panel function 8.

Note: When choosing between shutdown types, use **Delayed** first, and then select **Immediate** only if the delayed shutdown has not been progressing for several minutes.

- After selecting the desired shutdown type, press **OK** to start the shutdown process. See Figure 6-21.

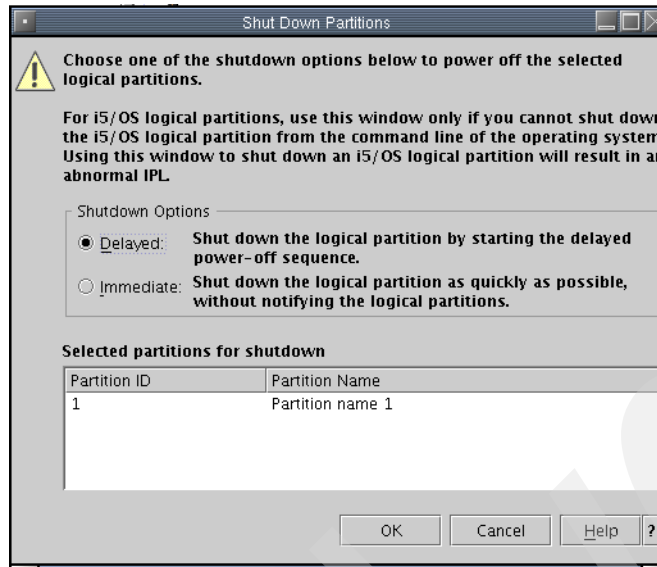


Figure 6-21 Partition shutdown options

Attention: Both delayed and immediate shutdown types are considered abnormal system ends and longer IPL times may result. Damaged objects are also possible. An immediate shutdown is more likely to result in an abnormal IPL.

6.4 Performing dynamic LPAR (DLPAR) functions

Dynamic LPAR (DLPAR) can be performed against the following types of resources:

- ▶ Physical Adapters
- ▶ Processors
- ▶ Memory
- ▶ Virtual IO Adapters

6.4.1 Physical adapters

In this section we discuss dynamic LPAR (DLPAR) operations on physical adapters.

Figure 6-22 shows how to access the physical adapter DLPAR options.

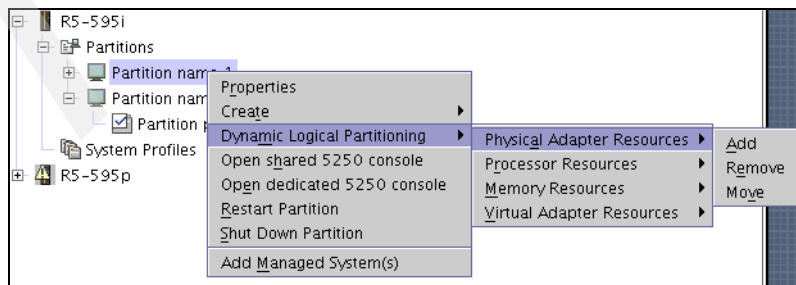


Figure 6-22 DLPAR functions for Physical Adapter Resources

In order to add, move, and remove physical adapters using DLPAR, perform the following steps:

Add

Dynamic Logical Partitioning → Physical Adapter Resources → Add

Remove

Dynamic Logical Partitioning → Physical Adapter Resources → Remove

Move

Dynamic Logical Partitioning → Physical Adapter Resources → Move

6.4.2 Processors

In this section we discuss dynamic LPAR (DLPAR) operations on processors.

Figure 6-23 shows how to access the processor DLPAR options.

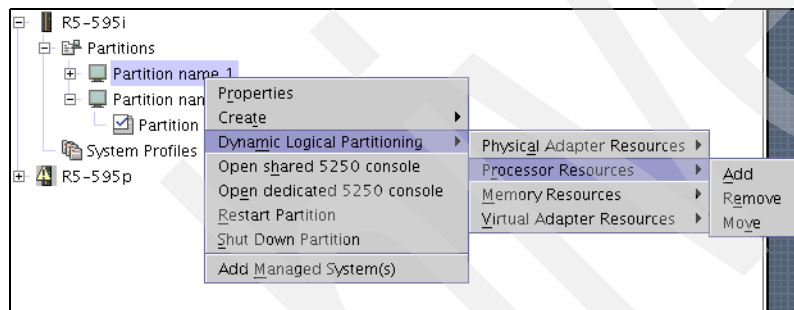


Figure 6-23 DLPAR functions for Processor Resources

In order to add, move, and remove processor resources using DLPAR, perform the following steps:

Add

Dynamic Logical Partitioning → Processor Resources → Add

Remove

Dynamic Logical Partitioning → Processor Resources → Remove

Move

Dynamic Logical Partitioning → Processor Resources → Move

6.4.3 Memory

In this section we discuss dynamic LPAR (DLPAR) operations on memory.

Figure 6-24 shows how to access the memory DLPAR options.

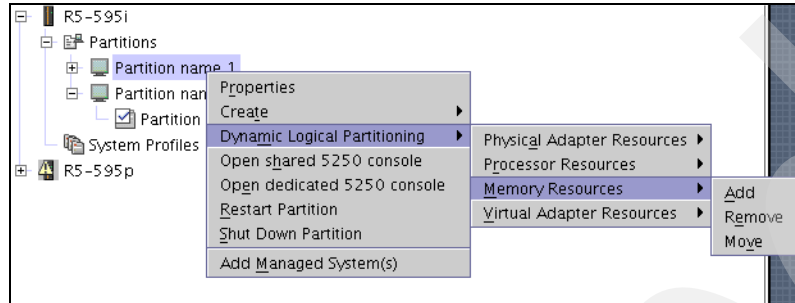


Figure 6-24 DLPAR functions for Memory Resources

In order to add, move, and remove memory resources using DLPAR, perform the following steps:

Add

Dynamic Logical Partitioning → Memory Resources → Add

Remove

Dynamic Logical Partitioning → Memory Resources → Remove

Move

Dynamic Logical Partitioning → Memory Resources → Move

6.4.4 Virtual IO adapters

While other dynamic LPAR (DLPAR) functions are performed at the partition level, creating additional virtual IO adapters or deleting ones that are no longer required is performed against a particular partition profile.

As shown in Figure 6-25, right-clicking a profile will display its context menu. From the context menu, select **Properties**.




Figure 6-25 Displaying properties for partition profile

From the properties dialog, select the **Virtual I/O** tab. Figure 6-26 shows a sample virtual IO adapter listing for a partition profile.

As shown in Figure 6-26, the properties on an existing virtual IO adapter can be displayed by selecting that IO adapter and clicking **Properties**. From the properties dialog, the settings for a virtual IO adapter can also be changed.

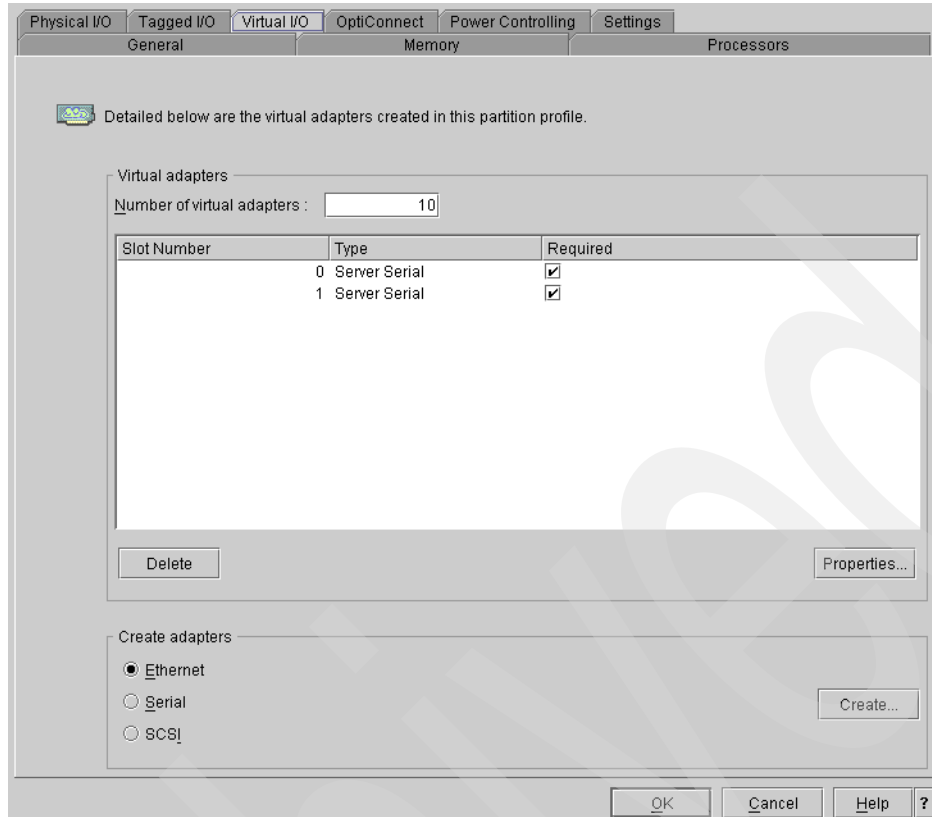


Figure 6-26 Virtual I/O properties for a partition profile

Create

Before creating a virtual IO adapter, increase the **Number of virtual adapters**, if required, to accommodate the additional adapters. To create a virtual IO adapter, select the desired adapter type and click **Create**.

For additional discussion on creating virtual IO adapters, refer to the following headings under 5.2.10, “Virtual IO adapters” on page 157:

- ▶ Virtual Ethernet - 5.2.10, “Virtual IO adapters” on page 157
- ▶ Virtual Serial - 5.2.10, “Virtual IO adapters” on page 157
- ▶ Virtual SCSI - 5.2.10, “Virtual IO adapters” on page 157

Delete

If a particular virtual IO adapter is no longer required, it can be deleted. Select the adapter and select **Delete**.

Properties

The **Properties** button will display properties of selected virtual I/O adapter.

6.5 Performing service functions on partitions

For information about accessing the LPAR functions menu, see 6.1, “Accessing LPAR functions” on page 178.

6.5.1 Displaying reference code information

In order to display reference code information for a partition, select **Properties** from the LPAR functions menu. Select the **Reference Code** tab on the properties dialog, as shown in Figure 6-27.

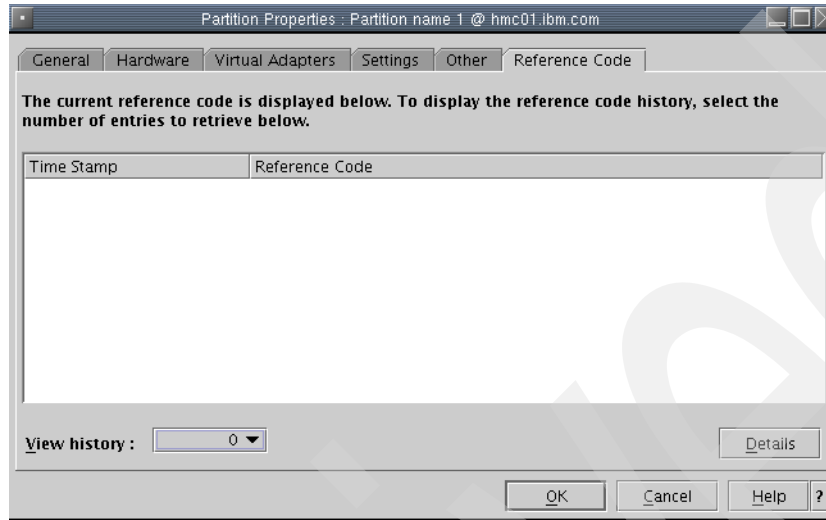


Figure 6-27 Displaying reference code information for a partition

6.5.2 Posting DST to the console

Dedicated Service Tools (DST) can be posted to the partition console during normal runtime. This allows for the majority of DST functions and options to be accessed while the partition is running without the need for a manual IPL. In order to post DST for a partition, the partition needs to be in manual mode.

Note: Posting DST disconnects any running console job. Before posting DST, sign off the console.

In order to post DST to the partition console, perform the following steps:

1. Expand **Service Applications** from the Navigation Area on the left side of HMC window, click **Service Focal Point**.
2. Click **Service Utilities**.
3. Select desired system unit, then click the **Selected** pull down menu to select **Operator Panel Service Functions...**
4. Select a partition to which DST will be posted, click **Partition Functions** pull down menu.
5. Select **Activate Dedicated Service Tools (21) - i5/OS** to post DST to the partition console.

You will need to signon to DST with the appropriate service tools user ID and password, as shown in Figure 6-28.

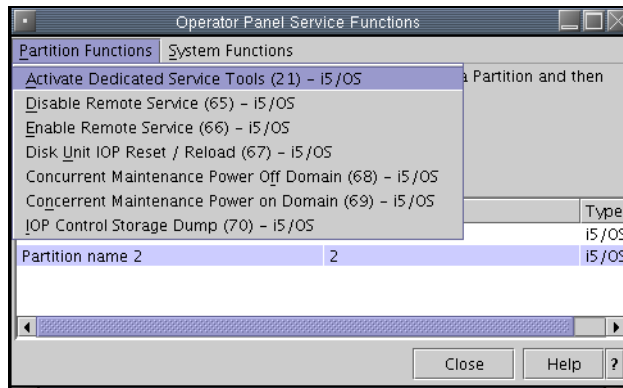


Figure 6-28 Posting DST to partition console

6.5.3 Performing main storage dumps

Attention: Perform a main storage dump under the direction of your next level of support. Incorrect use of this service tool can result in unnecessary down time and loss of debug information.

A main storage dump (MSD) is the contents of main storage, that is, system memory or RAM, from a single moment in time. Main storage dumps can either be initiated automatically by the partition in response to a severe Licensed Internal Code error, or manually by the user. In order to manually initiate a main storage dump, the partition needs to be in manual mode.

Manually, a main storage dump is initiated as follows:

1. If the partition is not in manual mode, change the mode to manual. See 6.3.1, “Changing IPL side and mode” on page 191 for additional information.
2. Bring up the reboot dialog for the partition (6.3.3, “Restarting a running partition” on page 193) and select **Dump** and then click **OK**.

Retry MSD IPL

During MSD processing there is a special short IPL before the MSD can be viewed or copied to either disk or media. This IPL is called the MSD IPL. In the event of an error, this IPL can be retried without losing the contents of main storage that have not yet been saved for diagnostic purposes.

The retry MSD IPL is initiated as follows:

1. If the partition is not in manual mode, change the mode to manual. See 6.3.1, “Changing IPL side and mode” on page 191 for additional information.
2. Bring up the reboot dialog for the partition (6.3.3, “Restarting a running partition” on page 193) and select **Dump Retry** and then click **OK**.

6.5.4 Working with remote service

The remote service support communications line is used by Operations Console (see Figure 6-29). Support tools like RSSF (Remote Service and Support Facility) can also use this line. From the HMC GUI, this line can be activated or deactivated.

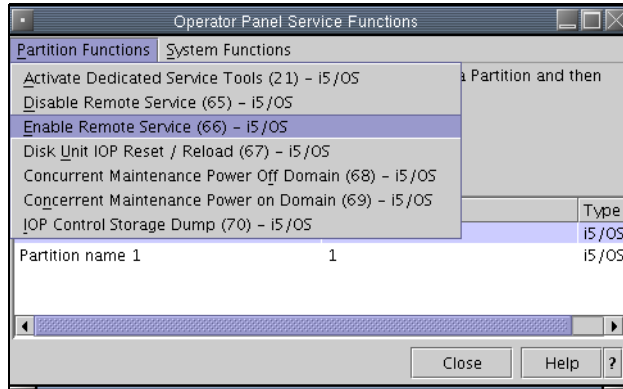


Figure 6-29 Activate and Deactivate remote service functions

In order to activate the remote service communication, perform the following steps:

1. Expand **Service Applications** from the Navigation Area on the left side of HMC window, click **Service Focal Point**.
2. Click **Service Utilities**.
3. Select desired system unit, then click the **Selected** pull down menu to select **Operator Panel Service Functions....**
4. Select a partition, and click the **Partition Functions** pull down menu.
5. Select **Enable remote Service (66) - i5/OS** to enable remote service communication to selected partition.

This function will attempt to activate the remote support communications line. Reference code D1008066 is posted if the line activation was successful. The reference code D1008065 indicates that the line failed to activate.

In order to deactivate the remote service communication, perform the following steps:

1. Expand **Service Applications** from the Navigation Area on the left side of HMC window, click **Service Focal Point**.
2. Click **Service Utilities**.
3. Select desired system unit, then click the **Selected** pull down menu to select **Operator Panel Service Functions....**
4. Select a partition, and click the **Partition Functions** pull down menu.
5. Select **Disable remote Service (65) - i5/OS** to disable remote service communication.

This function will attempt to deactivate the remote support communications line. Reference code D1008065 indicates that the line is no longer active.

6.5.5 Deleting a partition

If a logical partition is no longer needed, you can delete the partition and all of the associated profiles.

Important: While the partition may no longer be required, save any required user data before deleting the partition.

To delete a partition, access the LPAR functions menu and select **Delete**. A confirmation dialog is displayed. Click **OK** to confirm the partition delete or **Cancel** to back out the request, as shown in Figure 6-30.

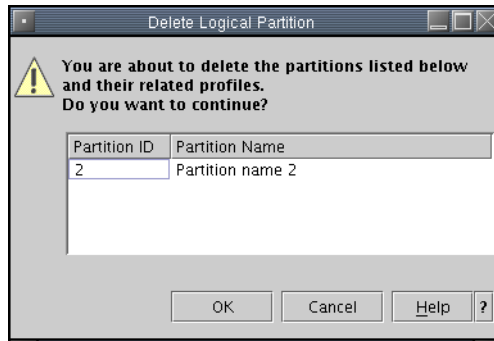


Figure 6-30 Confirming the deletion of a logical partition

Deleting a partition profile

If a partition has more than one partition profile, rather than deleting the entire partition, only the profile that is no longer needed can be deleted.

Restriction: An active or default profile cannot be deleted. The profile would need to be made inactive or another profile made default before the original profile could be deleted.

In order to delete a partition profile, select the profile and right click. From the pop up menu, select **Delete** as shown in Figure 6-31.




Figure 6-31 Deleting a partition profile

A confirmation dialog is displayed in Figure 6-32. Click **OK** to confirm the partition profile delete or **Cancel** to back out the request.

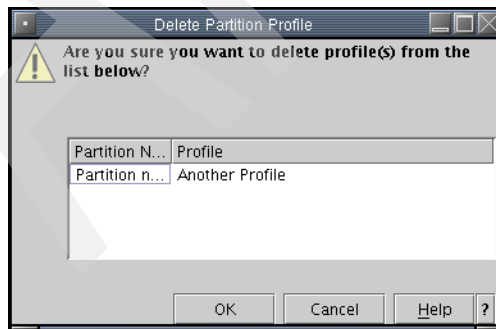


Figure 6-32 Confirming the deletion of a partition profile

6.5.6 Working with IOP functions

In this section we discuss operational considerations and actions that can be performed on IOPs while the system is running.

IOP reset (disk unit IOP reset/reload)

An IOP reset is only valid when certain disk unit subsystem error codes are posted.

Attention: Perform a disk unit IOP reset under the direction of your next level of support. Incorrect use of this service tool can result in unnecessary down time and loss of debug information.

This function performs a reset of the disk unit IOP and instructs the IOP to reload its microcode. See Figure 6-33.

Manually, you can initiate disk unit IOP reset/reload by performing the following steps:

1. Expand **Service Applications** from the Navigation Area on the left of HMC window, click **Service Focal Point**.
2. Click **Service Utilities**.
3. Select desired system unit, then click the **Selected** pull down menu to select **Operator Panel Service Functions...**
4. Select a partition, and click the **Partition Functions** pull down menu.
5. Select **Disk Unit IOP Reset / Reload (67) - i5/OS**.

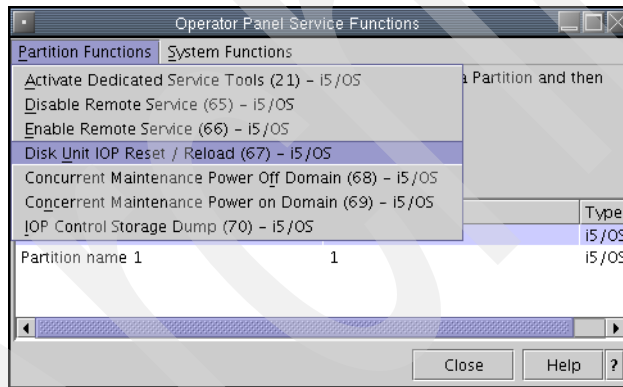


Figure 6-33 Resetting a disk unit IOP

IOP control storage dump

Attention: Perform an IOP control storage dump under the direction of your next level of support. Incorrect use of this service tool can result in unnecessary down time and loss of debug information.

This function dumps the contents of service processor memory to a predefined space on the load source disk unit and initiates a reset and reload of the service processor. See Figure 6-34.

Manually, you can initiate IOP control storage dump by performing the following steps:

1. Expand **Service Applications** from the Navigation Area on the left side of HMC window, click **Service Focal Point**.
2. Click **Service Utilities**.
3. Select desired system unit, then click the **Selected** pull down menu to select **Operator Panel Service Functions....**
4. Select a partition, and click the **Partition Functions** pull down menu.
5. Select **IOP control storage dump (70) - i5/OS**.

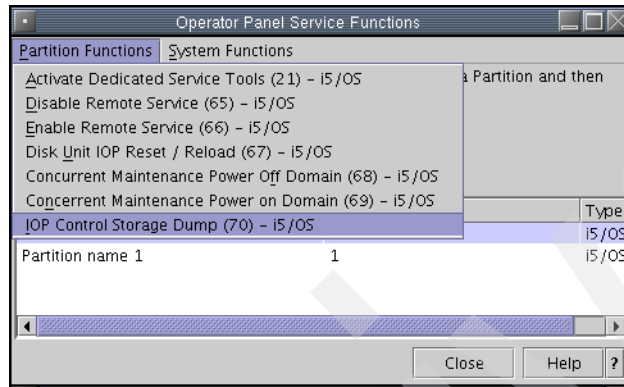


Figure 6-34 IOP control storage dump

6.5.7 Concurrent Maintenance Power Domain

Attention: Perform power domain functions under direction of your next level of support. Incorrect use of this service tool can result in unnecessary down time.

Concurrent maintenance allows for towers, individual slots or drive bays to be powered off for removing and installing hardware. This way, the entire platform does not need to be down for most hardware actions, as shown in Figure 6-25.

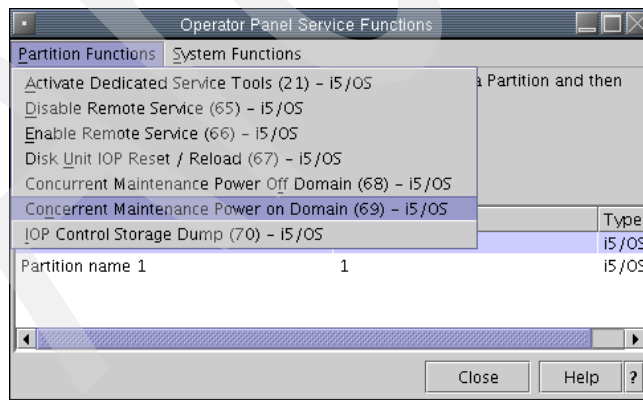


Figure 6-35 Concurrent Maintenance Power Domain functions

Power On I/O Domain

To perform power on I/O domain, do the following steps:

1. Expand **Service Applications** from the Navigation Area on the left side of HMC window, click **Service Focal Point**.
2. Click **Service Utilities**.
3. Select desired system unit, then click the **Selected** pull down menu to select **Operator Panel Service Functions...**
4. Select a partition, and click the **Partition Functions** pull down menu.
5. Select **Concurrent Maintenance Power on Domain (69) - i5/OS**.

Power Off I/O Domain

To perform power off I/O domain, do the following steps:

1. Expand **Service Applications** from the Navigation Area on the left side of HMC window, click **Service Focal Point**.
2. Click **Service Utilities**.
3. Select desired system unit, then click the **Selected** pull down menu to select **Operator Panel Service Functions...**
4. Select a partition, and click the **Partition Functions** pull down menu.
5. Select **Concurrent Maintenance Power Off Domain (68) - i5/OS**.

6.6 Remote management

Attention: Perform an IOP control storage dump under the direction of your next level of support. Incorrect use of this service tool can result in unnecessary down time and loss of debug information.

Remote management is enabled through a scripted secure shell environment. In this section we describe some of the remote function.

Enabling SSH

By default, SSH (secure shell) is disabled on a brand new HMC install. In order to use the HMC CLI to remotely manage a partitioned system, SSH needs to first be enabled. Enabling SSH takes affect immediately.

To enable SSH, perform the following tasks, as shown in Figure 6-37.

1. In the Navigation Area, expand the desired host (if there is more than one) and drill down as follows:

HMC Management → **HMC Configuration**

2. In HMC Configuration pane, select **Enable or Disable Remote Command Execution**.
3. Make sure the check box **Enable remote command execution using the SSH facility** is selected.
4. Click **OK** to activate the change.

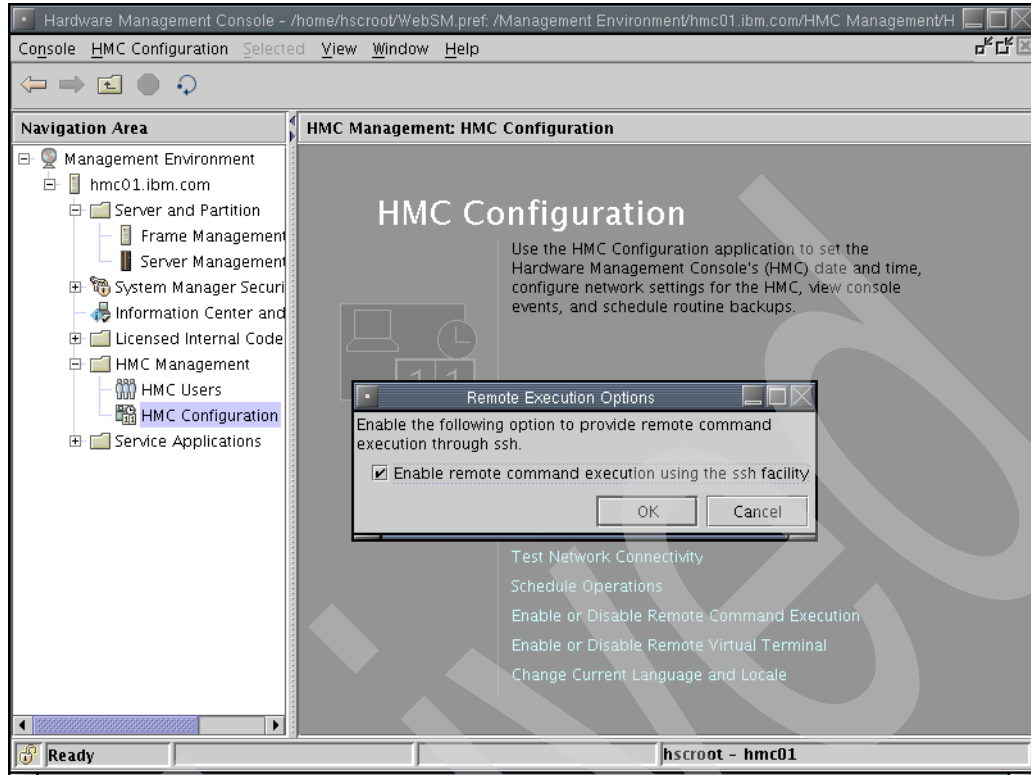


Figure 6-36 Enabling SSH (Secure Shell)

SSH client

In order to connect to the HMC via SSH, an SSH client needs to be installed on the client PC. One such client is PuTTY. Setup and install of the SSH client is outside of the scope of this document.

Connecting SSH client to the HMC

Whichever SSH client is used, you will need to connect to port 22 on the HMC. Figure 6-37 shows connecting with the PuTTY SSH client. When connected, you will be presented with a UNIX-like signon screen as shown in Figure 6-38. Signon with the appropriate HMC user ID and password.

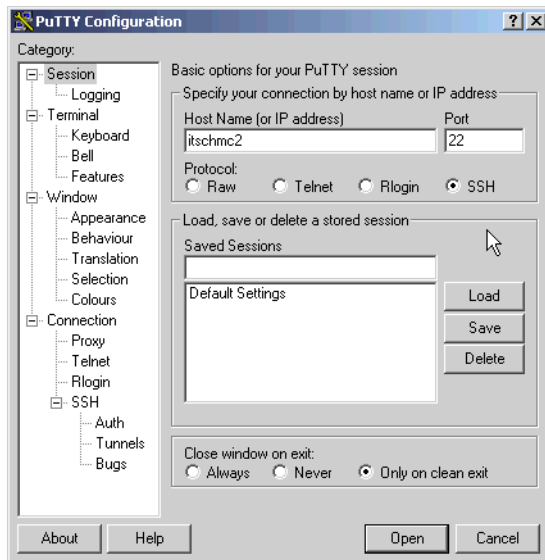


Figure 6-37 PuTTY SSH config




Figure 6-38 SSH logon to HMC

Example SSH command

Below is a relatively simple example that lists the partition ids, names, and states of the partitions on a managed system called Default2, as shown in Example 6-1.

Example 6-1 Sample lssyscfg command

```
hscroot@itschmc2:~> lssyscfg -r lpar -m Default2 -F lpar_id,name,state --header
lpar_id,name,state
2,lp 2 RPA,Off
1,lp 1 OS/400,Off
3,lp 3 MM,Off
123,a really long partition name,Off
```

Let us take a look at the command string parameters:

- lssyscfg — The name of the command: LiSt SYStem ConFiGuration
- r lpar — The resources to list: list the logical partition resources
- m Default2 — Says for which managed system to list the resources
- F lpar_id,name,state — Specifies the fields or attributes to be output
- header — Tells the command to display a header row with the outputted fields.

Tip: The help text for a command can be obtained by specifying the parameter `--help`. For example, `lssyscfg --help` will list the help text for the list system configuration command.

HMC commands

The HMC command line interface (CLI) is discussed in Appendix A, “HMC command list” on page 473.

6.6.1 Web SM

Web SM is short for Web-based System Manager Remote Client. It allows remote access to most HMC functions from a network attached PC client. This function is very useful for remote technical support and planning.

Obtaining the client

In order to obtain the Web SM client, open a Web browser session to the following URL:

`http://hostname/remote_client.html`

The *hostname* needs to be replaced with the hostname or IP address of your HMC. Once connected, you will be prompted for HMC user ID and password. Figure 6-39 shows a sample connection to an HMC to download the Web SM client.

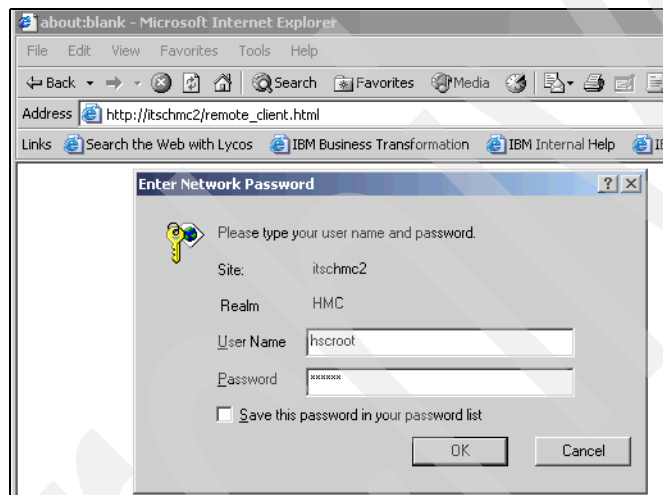


Figure 6-39 Connecting to an HMC to download the Web-based System Management remote client

For the remote client type, select **Static Remote Client**. Select the version of the client based on the operating system level of the client.

Installing the client

To install the client, run the setup executable that was downloaded from the HMC. A wizard is started to guide you through the install of the Web-based System Management remote client. In most cases, the default values are OK.

Starting the Web-based client

The Windows version of the Web-based System Manager Remote Client is started as follows from the desktop:

Start → Programs → Web-based System Manager → Web-based System Manager Remote Client

The first piece of information required in the signon process is the hostname or IP address of the HMC (see Figure 6-40). At the moment, the HMC user and password cannot be specified.

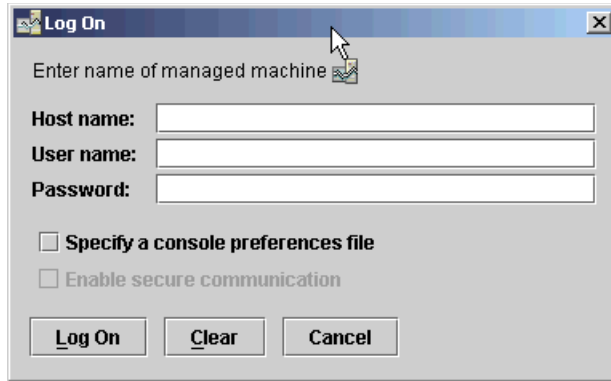


Figure 6-40 Web SM logon dialog

Next, the Web-based System Management remote client and the HMC sync up via handshaking as shown in Figure 6-41.

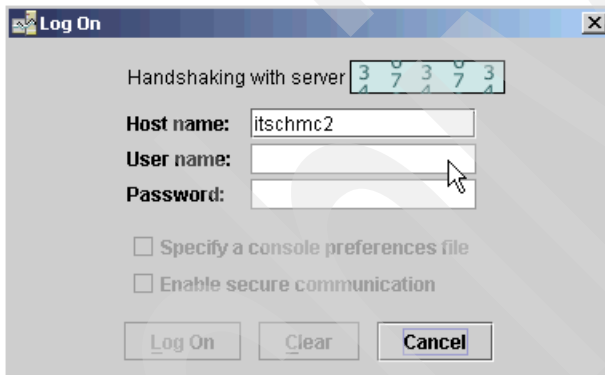


Figure 6-41 Handshaking between Web-based System Management remote client and the HMC

Once the handshaking completes, the HMC user name and password needs to be specified, as shown in Figure 6-42. Click **Log On** to continue the connection process.

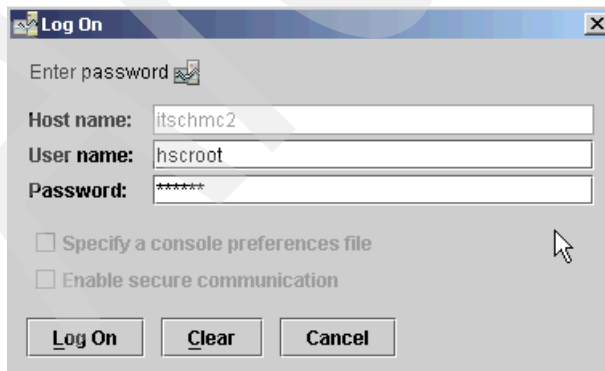


Figure 6-42 Providing the HMC user and password information

Once the signon process is complete, a Web-based System Management remote client display is shown that is similar to the one in Figure 6-43 below. Except for security functions, the display is practically identical to the local HMC display. A sample of the local HMC display is shown earlier in this chapter in Figure 6-1 on page 178.

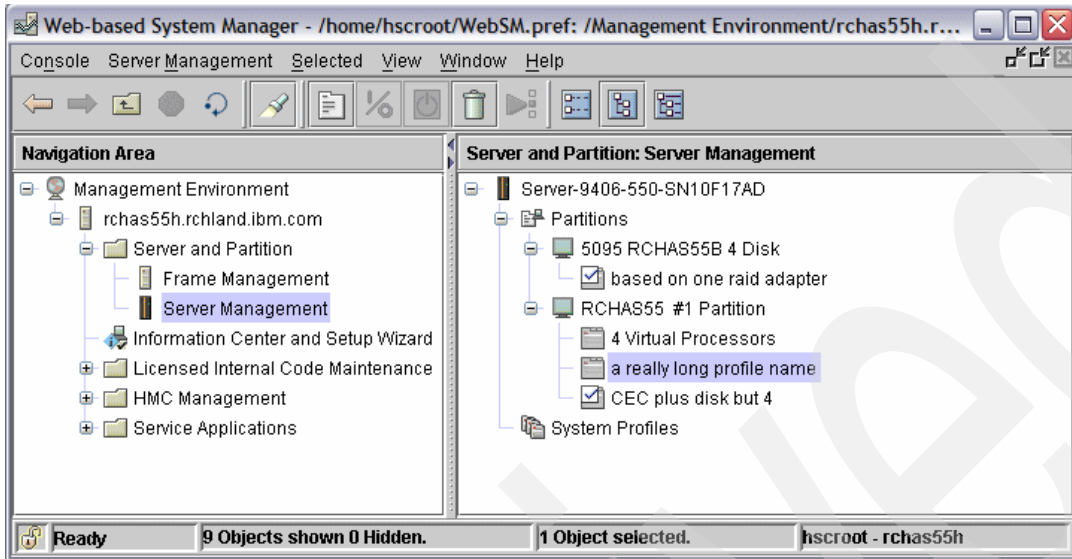


Figure 6-43 The main Web-based System Management remote client display

6.7 iSeries control panel functions versus HMC options

Table 6-2 provides a cross reference between the iSeries control panel functions and their counterparts on the HMC menu.

Table 6-2 iSeries control panel to HMC option cross reference

iSeries control panel function	Description	HMC option
White button power on (virtual panel option 1)	power on, start IPL	Activate partition profile 6.3.2, "Starting a powered off partition" on page 192
3	IPL Restart	Immediate reboot option 6.3.3, "Restarting a running partition" on page 193
double white button (virtual panel option 7)	Delayed power off	Delayed shutdown option 6.3.4, "Stopping (powering down) a running partition" on page 194
8	Immediate power off	Immediate shutdown option 6.3.4, "Stopping (powering down) a running partition" on page 194
11-19	System Reference Codes	System Reference Code tab on partition properties dialog 6.5.1, "Displaying reference code information" on page 199

iSeries control panel function	Description	HMC option
20	Machine Type/Model	Partition properties - General tab 6.2.1, "Partition properties" on page 179
21	Post DST to Console	Enable DST 6.5.2, "Posting DST to the console" on page 199
22	Force MSD	6.5.3, "Performing main storage dumps" on page 200
34	Retry MSD IPL	"Retry MSD IPL" on page 200
65	Deactivate remote service	Deactivate remote service
66	Activate remote service	Activate remote service
67	Disk unit IOP reset	IOP Reset "IOP reset (disk unit IOP reset/reload)" on page 203
68	Concurrent Maintenance power domain Power OFF	Power off I/O Domain
69	Concurrent Maintenance power domain Power ON	Power on I/O Domain
70	Service Processor / MFIOIP dump	IOP Control Storage Dump "IOP control storage dump" on page 203

6.8 Licensed Internal Code Updates

Figure 6-44 shows that this option is divided into three main areas of Licensed Internal Code Updates: Change Licensed Internal Code for the current release, Flash Side selection, and Upgrade Licensed Internal Code to a new release.

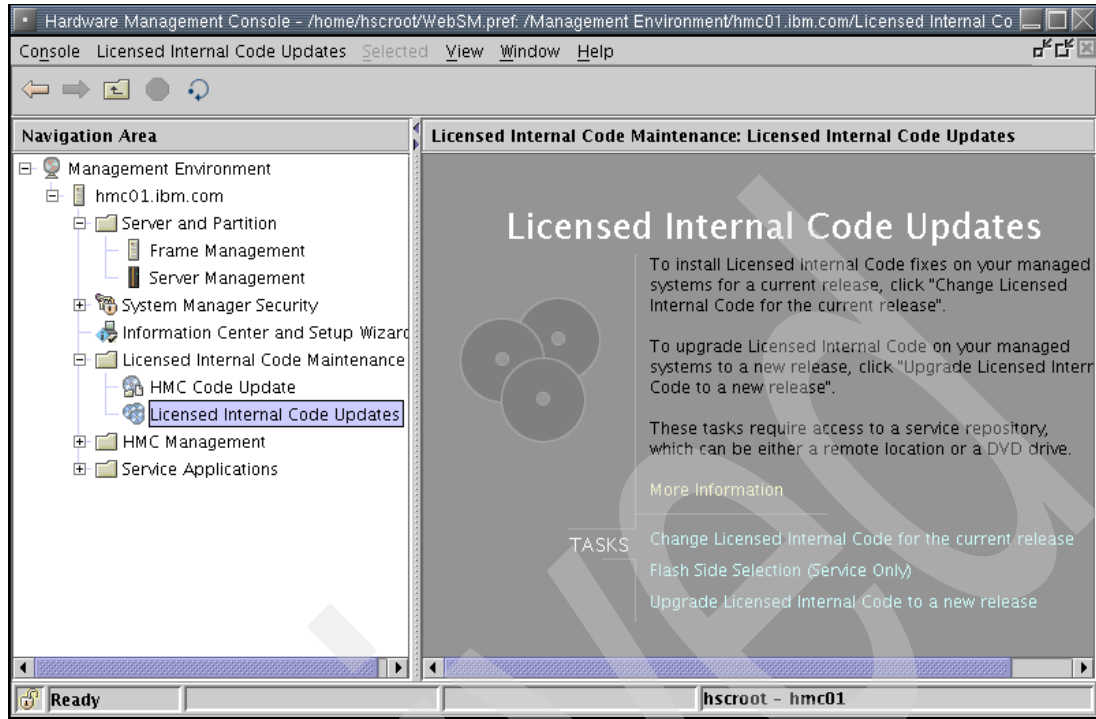


Figure 6-44 Licensed Internal Code Maintenance option

6.8.1 HMC Code Update

Clicking HMC Code Update presents four options:

- ▶ Backup Critical Console Data.
- ▶ Save Upgrade Data.
- ▶ Install Corrective Service.
- ▶ Format Removable Media.
- ▶ Remote Restore of Critical Console Data.

The option to backup critical console data is important, as the name implies and backs up data such as the following:

- ▶ User preference files
- ▶ User information
- ▶ HMC platform-configuration files
- ▶ HMC log files

The backup saves the HMC data stored on the HMC disk drive to the DVD-RAM and is critical to support HMC operations. You should back up the HMC data after you have made any changes to the HMC or to information associated with partitions.

The option to save upgrade data saves system information onto the HMC hard disk in preparation for an upgrade of the HMC software. It should only be performed immediately prior to upgrading the HMC software from one release to the next. Any configuration changes made after performing this task will not be migrated to the new HMC software release.

The save upgrade data option saves configuration data, including the following items:

- ▶ System preferences
- ▶ Profile information
- ▶ Service Agent files
- ▶ Inventory Scout Service Files

The option to install corrective service fixes on the HMC is similar to the same option for the Frame. It allows you to update the level of code on the HMC either from removable media or from a remote site.

The option to format removable media allows you to format the diskettes with the DOS file system or format the DVD-RAM with the UDF file system.

6.9 Troubleshooting

In this section we cover some possible problems with the Management Server.

6.9.1 Management Server State

If the **Management Server** hangs after power on, you should check **State** first.

Figure 6-45 shows an example to view **State** with HMC.

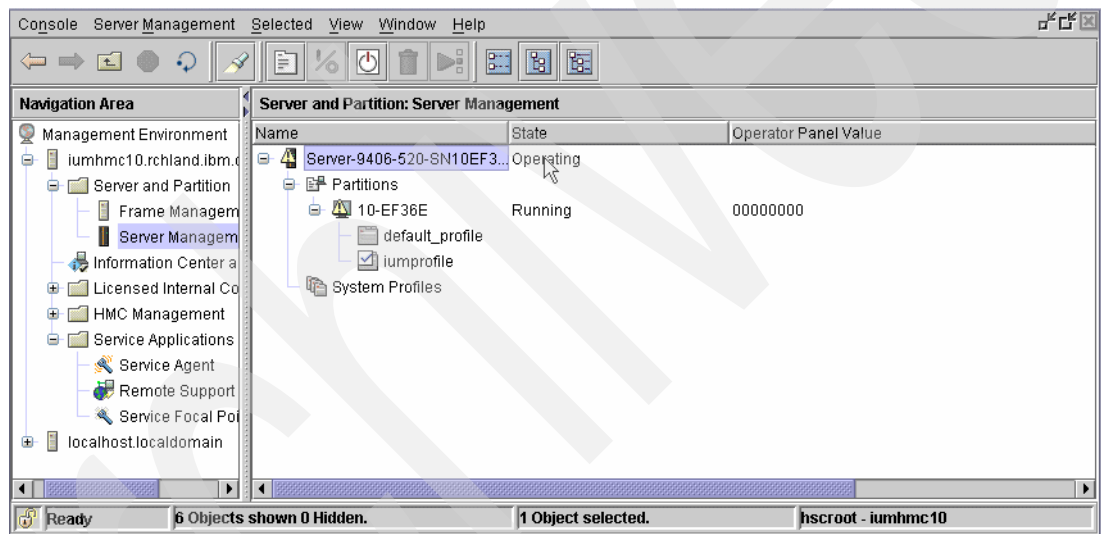


Figure 6-45 HMC-to view Management Server State

State is the value shown on the overview screen and indicates the current server status. Possible values for “State” are shown in Table 6-3.

Table 6-3 Possible values for State

Value	Description
Operating	The server is initialized and is working.
Power off	The server is powered off.
Initializing	The server is powered on and is initializing. The initialization time may vary depending on the hardware and the partition configuration of the server.
Pending authentication	The system is waiting for the password to be authenticated.
Failed authentication	The service processor password is not synchronized with the HMC password

Value	Description
Error	The operating system or the hardware of the managed system is experiencing errors.
Error - Terminated	Power On, Dump in progress.
CoD Click to Accept	Power On, operational and waiting for CUoD. Click to Accept.
Powering Off	Power Off in progress.
Standby	The managed system is powered on using the Power on Standby option. It will stay in this state until a partition or system profile is activated. You can create and modify profiles while the managed system is in this state.
Incomplete	The HMC failed to get all the information from the managed system.
No connection	The HMC is not connected or the handshake with the managed system failed.
Recovery	The save area in the Service Processor is not synchronized with the HMC.

6.9.2 Correcting the managed system operating state

If you find the following status, use these procedures to put the managed system back to a correct state:

- ▶ No Connection state
- ▶ Incomplete state
- ▶ Recovery state
- ▶ Error state

No Connection state

In the *No Connection* state, the HMC is not connected, or the handshake with the managed system failed. Perform the following steps:

1. Ensure that the network setting is correctly set up on the HMC. You can verify this by pinging the HMC.
2. Ensure that the network setting is correct on the managed system. Verify this by opening an ASMI session through the web browser and pointing at the managed system.
3. Perform the reset or remove connection task on the HMC:
 - a. In the navigation area, select **Server and Partition** → **Server Management**.
 - b. In the Content area, right-click the managed system and select the **Reset or Remove Connection**.
 - c. Select **Reset Connection**, then click **OK**.
 - d. Wait for 5 minutes for the HMC to re-establish contacts.
4. Check the control panel on the managed system to verify that the power is on.
 - a. If the power light indicates that there is no power on the managed system control panel, then power on the managed system manually or with AMSI.
 - b. After the power is restored, wait 5 minutes for the service processor to re-IPL and the HMC to re-establish contact. If the HMC can access partitions using a virtual terminal (VTERM) but the state is still *No Connection*, contact your next level of support or your hardware service provider.

5. If the power indicator is on, wait 5 minutes for the HMC to attempt to reestablish contact. The service processor in the managed system may be in the process of turning power on. If partitions no longer respond, the system power is off.
6. From a telnet session from another system, attempt to ping or contact active partitions on this managed system.

If the partitions are active, perform the following steps:

- a. Verify that the HMC and the flexible service processor are correctly connected to your Ethernet network, and that your Ethernet network is functioning correctly.
 - b. Restart the HMC.
 - c. Reset the service processor. See “Service processor reset” on page 225.
7. If the managed system is running, and restarting the HMC did not resolve the problem, contact your next level of support or your hardware service provider.

Incomplete state

In the *Incomplete state*, the HMC failed to get all of the necessary information from the managed system. Perform the following steps:

1. In the contents area, select the managed system.
2. From the menu, click **Selected** → **Rebuild Managed System**. If the state goes to *Recovery*, see “Recovery state” on page 215. If the state does not go to *Recovery*, continue with the next step.
3. Reset the connection from HMC. If it still fails, continue with the next step.

Notes: Step 4 may take several minutes to complete.

4. Restart the HMC. Choose from the following options:
 - If the state goes to *Recovery*, see *Recovery state*.
 - If the state remains *Incomplete*, perform the following steps:
 - i. Verify that there is a redundant HMC.
 - ii. Verify that no one is entering commands from the alternate HMC.
 - iii. Repeat step 1. If it still fails, continue with the next step.
5. Verify that the HMC and the flexible service processor are correctly connected to your Ethernet network, and that your Ethernet network is functioning correctly.
6. Reset the service processor. See “Service processor reset” on page 225.
7. If the problem persists, contact your next level of support or your hardware service provider.

Recovery state

In the *Recovery state*, the save area in service processor is not synchronized with the HMC Database. Perform the following steps:

1. Recover partition data. See Recovering partition data on a managed system for details. Then return here and continue with the next step.
2. After restoring profile data, choose the option that describes what happened:
 - If the state changed to *Incomplete*, refer to the *Incomplete state* procedure.
 - If the state changed to *No Connection*, refer to the *No Connection state* procedure.

- If the restore failed, reset the service processor. See “Service processor reset” on page 225. Then continue with the next step.
3. If the problem persists, contact your next level of support or your hardware service provider.

Error state

The Error state automatically generates a call to the service support center if the function is enabled. If the function is not enabled, contact your next level of support or your hardware service provider. You also can follow the next section to correct it.

6.9.3 List of progress codes

If the **State of Management Server** displays an *Error* status, then we should check the **Operator Panel Value**.

Figure 6-46 shows an example to view **Operator Panel Value** with HMC.

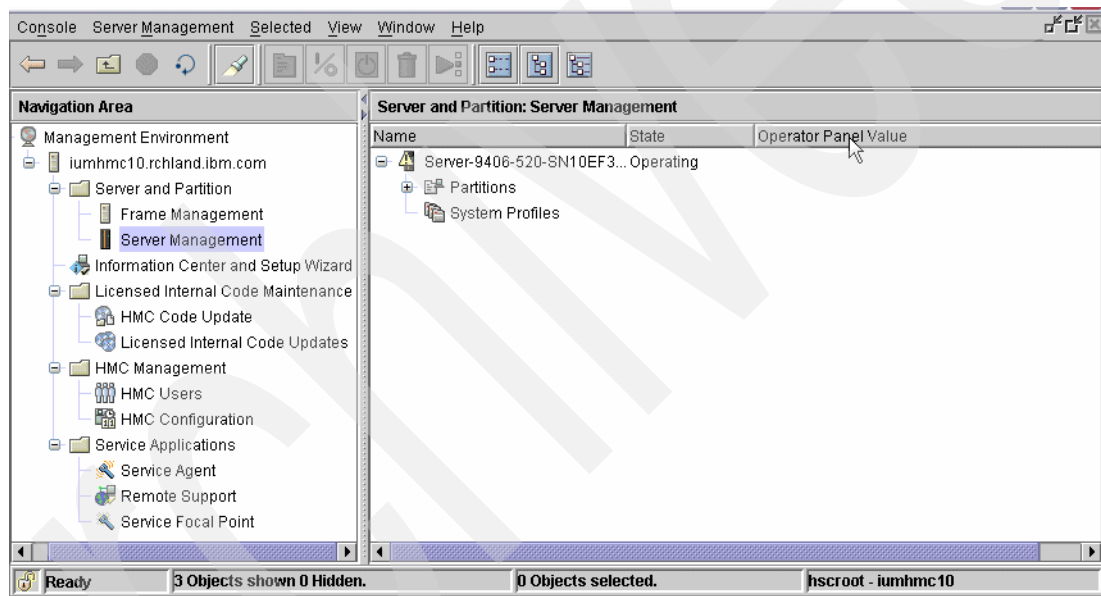


Figure 6-46 HMC-to view Management Server Operator Panel Value

The following progress codes (or checkpoints) displayed in the **Operator Panel Value** field represent different processes that occur during initial program load (IPL). The list of progress codes is in numeric order. A general description of the processes represented by the progress code appears beside each number (or range of numbers). Progress codes are 8-characters long, unless otherwise specified.

Use the following list (Table 6-4) to find information about specific progress codes. The first column represents the first 4 characters of the 8 character progress code, unless otherwise specified. The second column describes the progress code and links to information that includes one or more tables. If you do not find the progress code that you are looking for in this list, ask your next level of support for assistance.

Table 6-4 Progress code list and their meaning

Progress codes:	Refer to these topics:
4-character codes (including those that begin with a space character or a zero)	AIX IPL progress codes
C1xx	Service processor progress codes
C2xx	(C2xx) Virtual service processor progress codes
C3yx, C500, C5yx, C600, C6xx	IPL status progress codes
C700	(C700) Server firmware IPL status progress codes
C900	(C900) IPL status progress codes
CAxx	(CAxx) Partition firmware progress codes
D1xx	Service processor progress codes
D2xx	(D200) Partition status progress codes
D6xx	(D6xx) General status progress codes
D9xx	(D9xx) General status progress codes

6.9.4 Using progress codes

Progress codes (or checkpoints) offer information about the stages involved in powering on and performing initial program load (IPL). Although progress codes do not always indicate an error, in some cases your server can pause indefinitely (hang) without displaying an 8-character reference code. Only when you experience such a hang condition should you take any action with respect to a progress code. When you look up a progress code, the information indicates the most appropriate action for that progress code.

Progress code formats

Progress codes appear in several formats:

- ▶ An 8-digit hexadecimal number that begins with C or D
- ▶ A 3-digit hexadecimal number
- ▶ A 4-digit hexadecimal number that has a value between 2000 and 27FF
- ▶ A 4-digit hexadecimal number that begins with a 0 or 2

When the control panel or **Operator Panel Value** displays more than eight characters, use only the first eight characters to find the progress code. The characters that display beyond the first eight represent a location code that can assist you in diagnosing the problem.

Notes: Some operating systems, including Linux, do not produce progress codes when performing an IPL. Your server might display progress codes in the system control panel (for systems with a physical control panel) or in the virtual console on the Hardware Management Console.

Using the ASMI to view progress code logs

The service processor keeps two logs of progress codes. Use the Advanced System Management Interface (ASMI) to view the progress code logs. To view the progress logs, complete the following steps:

1. Access the ASMI and log in.

2. Select **System Information**, then select either **Previous Boot Progress Indicators** or **Progress Indicator History**

Selecting **Previous Boot Progress Indicators** shows the progress codes that the server displayed in the control panel during the last system boot. **Progress Indicator History** shows the latest progress codes that the server is currently displaying in the control panel.

Using the list of progress codes

The list of progress codes is in numeric order. The first column represents either an entire 4-character progress code or the first 4 characters of an 8-character progress code. The second column describes the progress code and links to information that includes one or more tables and, when necessary, a description of the table columns.

To use the list of progress codes, locate the code that you want to find in the first column and click the link in the second column.

Figure 6-47 shows an example to view progress code logs with ASMI.

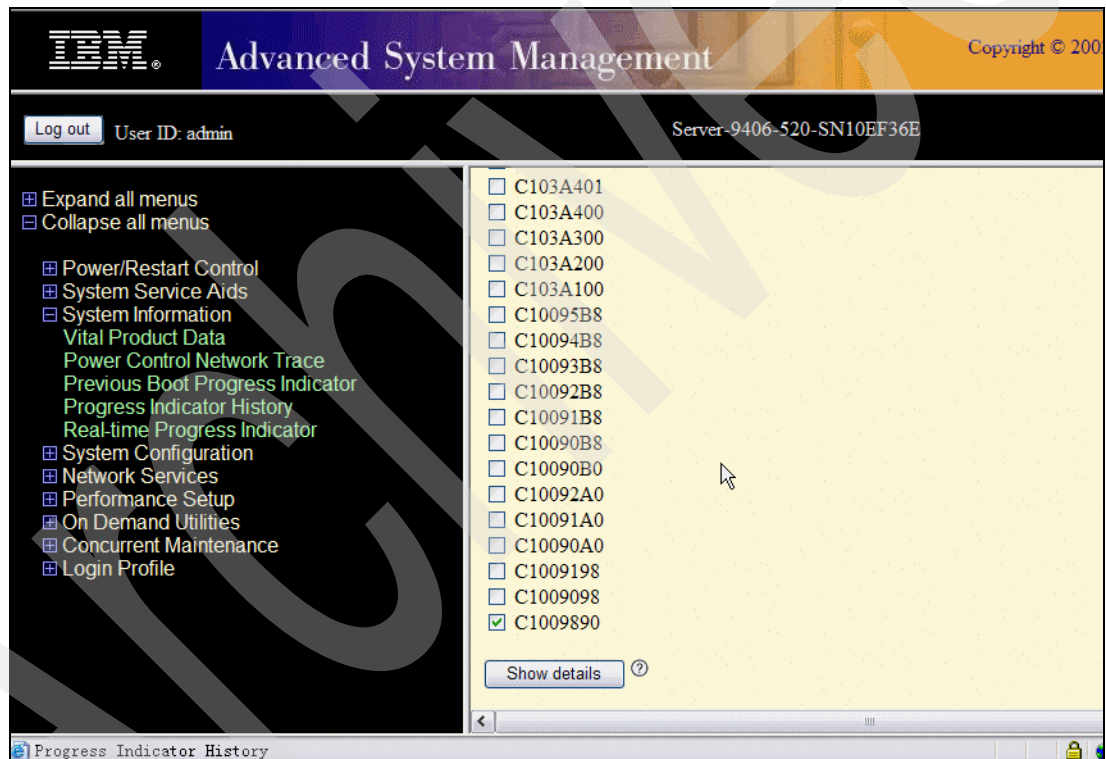


Figure 6-47 ASMI-to view the process code logs

6.9.5 Service processor progress codes

Service processor progress codes are typical reference codes that occur during the initial program load (IPL) of the server.

The tables for the service processor progress codes use the following format:

- ▶ The first column, Progress Code, contains characters that represent a particular action that the service processor performs to initialize your server.
- ▶ The second column, Description/Action, describes the type of initialization being performed.

- ▶ The third column, Failing Item, offers instructions for recovering from a hang on a progress code. Click the link in this column to view the service actions only when you experience a hang condition on a progress code. A hang condition exists when the code in the control panel display does not change for several minutes and the service processor appears to be stalled (hung). In some cases, you might notice that the server does not power down normally.

Use the following list to display the different types of service processor progress codes.

- ▶ (C1xx) Service processor progress codes (checkpoints)
- ▶ (D1xx) Service processor progress codes (service processor dump and platform dump)
- ▶ (D1xx) Service processor status progress codes

(C1xx) Service processor progress codes (checkpoints)

C1xx progress codes, or checkpoints, offer information about the initialization of both the service processor and the server. Table 6-5 gives a list of C1xx progress codes.

Table 6-5 (C1xx)Service processor progress codes(checkpoints)

Progress Code	Description/Action Perform all actions before exchanging Failing Items
C10010xx	Pre-standby
C1001F00	Pre-standby: starting initial transition file
C1001FFF	Pre-standby: completed initial transition file
C1009x02	Hardware object manager: (HOM): build cards IPL step in progress
C1009x04	Hardware object manager: (HOM): build processors IPL step in progress
C1009x06	Hardware object manager: (HOM): build chips IPL step in progress
C1009x08	Hardware object manager: (HOM): build HOM IPL step in progress
C1009x10	Hardware object manager: (HOM): configure HOM IPL step in progress
C1009x12	Hardware object manager: (HOM): initialize HOM IPL step in progress
C1009x18	Hardware object manager: (HOM): validate HOM IPL step in progress
C1009x1B	GARD IPL step in progress
C1009x20	Get LBIST values IPL step in progress
C1009x23	Clock test IPL step in progress
C1009x28	Frequency control IPL step in progress
C1009x30	Initialize JTAG IPL step in progress
C1009x38	Mask attentions IPL step in progress
C1009x40	CEC flush IPL step in progress
C1009x48	Scan interface basic assurance tests (SIBATs) IPL step in progress
C1009x50	LBIST IPL step in progress
C1009x58	ABIST IPL step in progress
C1009x60	Asset protection IPL step in progress

Progress Code	Description/Action Perform all actions before exchanging Failing Items
C1009x68	Wire test IPL step in progress
C1009x70	Memory size IPL step in progress
C1009x78	Long scan initialization IPL step in progress
C1009x80	Start clocks IPL step in progress
C1009x88	SCOM initialization IPL step in progress
C1009x90	Run interface alignment procedure IPL step in progress
C1009x98	DRAM initialization IPL step in progress
C1009x9B	Random data test IPL step in progress
C1009xA0	Memory diagnostic IPL step in progress
C1009xA8	Dump initialization IPL step in progress
C1009xB0	Processor runtime diagnostics (PRD) IPL step in progress
C1009xB8	Enable machine check IPL step in progress
C1009xC0	Message passing wait in progress
C1009xC8	Message passing start in progress
C100C1xx	Power control
C100D009	Licensed Internal Code (system) running initialization
C1011F00	Pre-standby: starting independent initial transition file (primary/secondary)
C1011FFF	Pre-standby: completed independent initial transition file (primary/secondary)
C1021F00	Pre-standby: starting primary Initial transition file (primary)
C1021FFF	Pre-standby: completed primary Initial transition file (primary)
C1031F00	Pre-standby: starting secondary Initial transition file (secondary)
C1031FFF	Pre-standby: completed secondary Initial transition file (secondary)
C103A100	Hypervisor code modules are being transferred to system storage
C103A200	Hypervisor data areas are being built in system storage
C103A300	Hypervisor data structures are being transferred to system storage
C103A400	Special purpose registers are loaded and instructions are started on the system processors
C103A401	Instructions have been started on the system processors
C103C2xx	The service processor is waiting for the batteries in the uninterruptable power supply (UPS) to charge prior to automatic power on-IPL. The last byte (xx) will increment while waiting on the UPS batteries.
C1041F00	Pre-standby: starting GardedInitial transition file (secondary)
C1041FFF	Pre-standby: completed GardedInitial transition file (secondary)
C1112000	Power on: starting Standby-PowerOnTransition transition file (primary)

Progress Code	Description/Action Perform all actions before exchanging Failing Items
C11120FF	Power on: completed Standby-PowerOnTransition transition file (primary)
C1122000	Power on: starting PowerOnTransition-PoweredOn transition file (primary)
C11220FF	Power on: completed PowerOnTransition-PoweredOn transition file (primary)
C1132000	Power on: starting PoweredOn-IplTransition transition file (primary)
C11320FF	Power on: completed PoweredOn-IplTransition transition file (primary)
C116C2xx	System power interface is listening for power fault events from SPCN. The last byte (xx) will increment up from 00 to 1F every second while it waits.
C1202000	IPL transition: starting PowerOn/IplTransition-Ipl transition file (primary)
C12020FF	IPL transition: completed PowerOn/IplTransition-Ipl transition file (primary)
C12040xx	IPL lock time left until expiration. The last byte (xx) will count down as the IPL lock time runs out (FF-00).
C1212000	IPL transition: starting Standard/IplTransition-Ipl transition file (primary)
C12120FF	IPL transition: completed Standard/IplTransition-Ipl transition file (primary)
C1222000	IPL transition: starting Flash/IplTransition-Ipl transition file (primary)
C12220FF	IPL transition: completed Flash/IplTransition-Ipl transition file (primary)
C1232000	IPL transition: starting PostDump/IplTransition-Ipl transition file (primary)
C12320FF	IPL transition: completed PostDump/IplTransition-Ipl transition file (primary)
C1242000	IPL transition: starting Idle/IplTransition-Ipl transition file (secondary)
C12420FF	IPL transition: completed Idle/IplTransition-Ipl transition file (secondary)
C1252000	IPL transition: starting Standby/IplTransition-Ipl transition file (secondary)
C12520FF	IPL transition: completed Standby/IplTransition-Ipl transition file (secondary)
C1382000	IPL: starting HostStarted-BcuSwitched transition file (primary)
C13820FF	IPL: completed HostStarted-BcuSwitched transition file (primary)
C1392000	IPL: starting BcuSwitched-Runtime transition file (primary)
C13920FF	IPL: completed BcuSwitched-Runtime transition file (primary)
C1402000	IPL: starting Normal/fast/Ipl-HostStarted transition file (primary)
C14020FF	IPL: completed Normal/fast/Ipl-HostStarted transition file (primary)
C1412000	IPL: starting Normal/slow/Ipl-HostStarted transition file (primary)
C14120FF	IPL: completed Normal/slow/Ipl-HostStarted transition file (primary)
C1422000	IPL: starting PostDump/Ipl-HostStarted transition file (primary)
C14220FF	IPL: completed PostDump/Ipl-HostStarted transition file (primary)
C1432000	IPL: starting Ipl-IdleTransition transition file (secondary)
C14320FF	IPL: completed Ipl-IdleTransition transition file (secondary)
C1442000	IPL: starting IdleTransition-Idle transition file (secondary)

Progress Code	Description/Action Perform all actions before exchanging Failing Items
C14420FF	IPL: completed IdleTransition-Idle transition file (secondary)
C1452000	IPL: starting Ipl-StandbyVerificationTransition transition file (secondary)
C14520FF	IPL: completed Ipl-StandbyVerificationTransition transition file (secondary)
C1462000	IPL: starting StandbyVerificationTransition-Standby transition file (secondary)
C14620FF	IPL: completed StandbyVerificationTransition-Standby transition file (secondary)
C162Exyy	VPD is being collected, where x indicates the phase of VPD collection and yy indicates the type of device from which VPD is being collected
C1802000	Termination: starting TerminationTransition-Termination transition file (primary)
C18020FF	Termination: completed TerminationTransition-Termination transition file (primary)
C1902000	Power off: starting Any-Dpo transition file (primary)
C19020FF	Power off: completed Any-Dpo transition file (primary)
C1912000	Power off: starting Any-PowerOffTransition transition file (primary)
C19120FF	Power off: completed Any-PowerOffTransition transition file (primary)
C1922000	Power off: starting PowerOffTransition-PoweredOff transition file (primary)
C19220FF	Power off: completed PowerOffTransition-PoweredOff transition file (primary)
C1C02000	Secondary VERIFICATION: starting Standby-StandbyVerification transition file (primary)
C1C020FF	Secondary verification: completed Standby-StandbyVerification transition file (primary)
C1C12000	Secondary verification: starting StandbyVerification-Standby transition file (primary)
C1C120FF	Secondary verification: completed StandbyVerification-Standby transition file (primary)
C1C22000	Secondary verification: starting Runtime-secondaryVerification transition file (primary)
C1C220FF	Secondary verification: completed Runtime-secondaryVerification transition file (primary)
C1C32000	Secondary verification: starting secondaryVerification-Runtime transition file (primary)
C1C320FF	Secondary verification: completed secondaryVerification-Runtime transition file (primary)
C1D22000	Dump: starting DumpTransition-Dump transition file (primary)
C1D220FF	Dump: completed DumpTransition-Dump transition file (primary)
C1E82000	Exit error: starting ExitError/Ipl transition file (primary)
C1E820FF	Exit error: completed ExitError/Ipl transition file (primary)
C1F22000	Reset/reload: starting Reset/Ipl-LimitedRuntime transition file (primary)
C1F220FF	Reset/reload: completed Reset/Ipl-LimitedRuntime transition file (primary)
C1F32000	Reset/reload: starting Reset/Ipl-Runtime transition file (primary)
C1F320FF	Reset/reload: completed Reset/Ipl-Runtime transition file (primary)

Progress Code	Description/Action Perform all actions before exchanging Failing Items
C1F42000	Reset/reload: starting Reset/Ipl-TermTransition transition file (primary)
C1F420FF	Reset/reload: completed Reset/Ipl-TermTransition transition file (primary)

(D1xx) Service processor progress codes (SP dump & Platform dump)

Service processor dump status codes use the format of D1yy1xxx, where yy indicates the type of data that is being dumped, xxx is a counter that increments each time the server stores 4K of data.

When these codes occur during a service processor dump or a platform dump, they appear in the control panel display. Table 6-6 gives a list of D1xx progress codes (SP dump and Platform dump).

Table 6-6 (D1xx) Service processor progress codes(service processor dump and platform dump)

Progress Code	Description/Action Perform all actions before exchanging Failing Items
D1001xxx	Dump error data
D1011xxx	Dump sai_header Hardware Management Console (HMC) file
D101C00F	No power off to allow debugging for CPU controls
D1021xxx	Dump sai_header directory
D1031xxx	Dump sai_header fips header
D1041xxx	Dump sai_header entry header
D1051xxx	Dump core file for failing component
D1061xxx	Dump all NVRAM
D1071xxx	Dump component trace for failing component
D1081xxx	Dump component data from /opt/p0
D1091xxx	Dump /opt/p1/*
D10A1xxx	Dump safety net trace: SND_LTT_TRACE
D10B1xxx	Dump safety net trace: SND_LTT_PROC
D10C1xxx	Dump safety net trace: SND_PRINTK_BUF
D10D1xxx	Dump safety net trace: SND_PRINTK_IDX
D10E1xxx	Dump safety net trace: SND_INSMOD_BUF
D10F1xxx	Dump safety net trace: SND_LTT_INTR
D1101xxx	Dump safety net trace: SND_SP_TRACE
D1111xxx	Dump /opt/p0/*
D1121xxx	Dump /opt/p1/*
D1131xxx	Dump all traces
D1141xxx	Dump code version

Progress Code	Description/Action Perform all actions before exchanging Failing Items
D1151xxx	Dump all /opt/p3 except rtbl
D1161xxx	Dump pddcustomize -r command
D1171xxx	Dump registry -l command
D1181xxx	Dump all /core/core.* files
D1191xxx	Dump BDMP component trace (after dump if enough space)
D11A1xxx	Dump any state information before dumping starts
D1231xxx	Dump update dump headers
D1241xxx	Dump CRC1 calculation off
D1251xxx	Dump CRC1 calculation on
D1261xxx	Dump CRC2 calculation off
D1271xxx	Dump CRC2 calculation on
D1281xxx	Dump output the calculated CRC1 (sai_headers)
D1291xxx	Dump output the calculated CRC2 (data and data headers)
D12A1xxx	Jump to the position in dump directly after CRC1
D12B1xxx	Initialize the headers dump time and serial numbers
D12C1xxx	Display final SRC to panel
D12D1xxx	Remove /core/core.app.time.pid
D12E1xxx	Remove /core/core.*
D12F1xxx	Display beginning SRC to panel
D1301xxx	Turn off error log capture into dump
D1311xxx	Turn on error log capture into dump
D1321xxx	Store information about existing core files
D1381xxx	Invalidate the dump
D1391xxx	Check for valid dump sequence
D13A1xxx	Get dump identity sequence
D13B1xxx	Get dump length sequence
D1FF1xxx	Dump complete
D1xx3000 - D1xx3FFF	Platform dump status codes

(D1xx) Service processor status progress codes(Platform power off)

D1xx reference codes, posted by the service processor, offer information about the state of the service processor during a power-off operation. Table 6-7 gives a list of D1xx progress codes(Platform power off).

Table 6-7 (D1xx) Service processor status progress codes(Platform power off)

Progress Code	Description/Action Perform all actions before exchanging Failing Items
D1xx900C	Breakpoint set in CPU controls has been hit
D1xxB0FF	Request to initiate power-off program has been sent
D1xxC000	Indicates a message is ready to send to the server firmware to power off
D1xxC001	Waiting for the server firmware to acknowledge the delayed power off notification
D1xxC002	Waiting for the server firmware to send the power off message
D1xxC003	Server firmware handshaking is complete

6.9.6 Service processor reset

If the system hangs after the code that sent you to this procedure appears in the control panel or **Operator Panel Value** in HMC, perform these steps to reset the service processor. If the platform is already powered off, you can reset SP with ASMI.

Attention: Resetting SP will lost customized configuration on SP! You need to synchronize from the HMC after resetting. Before resetting SP, please make sure that you keep a backup of the configuration data.

Reset SP with ASMI

Follow this procedure:

1. Access the ASMI and log in
2. Select **System Server Aids**, then select either **Reset Service Processor**

Figure 6-48 shows an example to reset SP with ASMI.

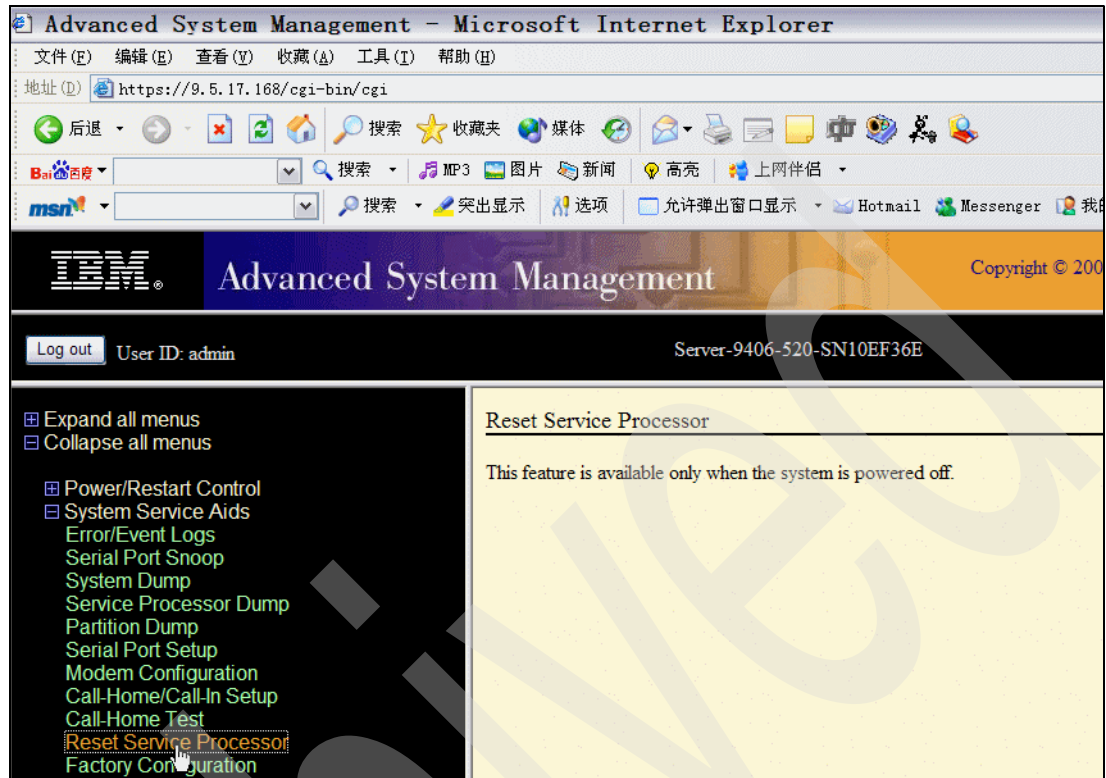


Figure 6-48 ASMI-to reset SP

Notes: This feature is available only when the system is powered off.

Reset SP with reset button

Follow this procedure:

1. Activate the service processor pinhole reset switch on the system's operator panel by carefully performing these steps:
 - a. Using an insulated paper clip, unbend the clip so that it has a straight section about two inches long.
 - b. Insert the clip straight into the hole, keeping the clip perpendicular to the plastic bezel.
 - c. When you engage the reset switch, you should feel the detent of the switch.
 - d. After you press the switch, the service processor is reset, then the system shuts down.
2. Reboot the system in slow mode from the permanent side using control panel function 02.
3. If the hang repeats, check with service support to see if there is a firmware update that fixes the problem. Skip to **Getting fixes** in this section for details.

4. Choose from the following options:
 - If there is no firmware update available, continue with the next step.
 - If a firmware update is available, apply it using the **Service Focal Point** in the HMC.
 - Did the update resolve the problem and the system now boots?

Yes: This ends the procedure.

No: You are here because there is no HMC attached to the system, the flash update failed, or the updated firmware did not fix the hang. Continue with the next step.

5. Choose from the following options:
 - If you are a customer, contact your hardware service provider. **This ends the procedure.**
 - If you are an hardware service provider, continue with the next step.

Notes: You can try to get help from IBM service engineer to reset SP with toggle jumpers before replacing SP.

6. Replace the service processor (skip to **FRU SVCPROC** in this section for details.).
7. If replacing the service processor does not fix the problem, contact your next level of support. **This ends the procedure.**

Getting fixes

Fixes provide changes to your software, Licensed Internal Code, or machine code that fix known problems, add new function, and keep your server or Hardware Management Console operating efficiently. For example, you might install fixes for your operating system in the form of a PTF (program temporary fix). Or, you might install a server firmware (Licensed Internal Code) fix with code changes that are needed to support new hardware or new functions of the existing hardware.

A good fix strategy is an important part of maintaining and managing your server. If you have a dynamic environment that changes frequently, you should install fixes on a regular basis. If you have a stable environment, you do not have to install fixes as frequently. However, you should consider installing fixes whenever you make any major software or hardware changes in your environment.

You can get fixes using a variety of methods, depending on your service environment. For example, if you use an HMC to manage your server, you can use the HMC interface to download, install, and manage your HMC and firmware (Licensed Internal Code) fixes. If you do not use an HMC to manage your server, you can use the functions specific to your operating system to get your fixes.

In addition, you can download or order many fixes through Internet Web sites. You can search by server and product to find the latest fixes for your system's software, hardware, and operating system. You must manage several types of fixes to maintain your hardware.

Figure 6-49 shows the hardware and software that might require fixes, including the HMC, I/O adapters and devices firmware, server firmware, power subsystem firmware, and operating systems.

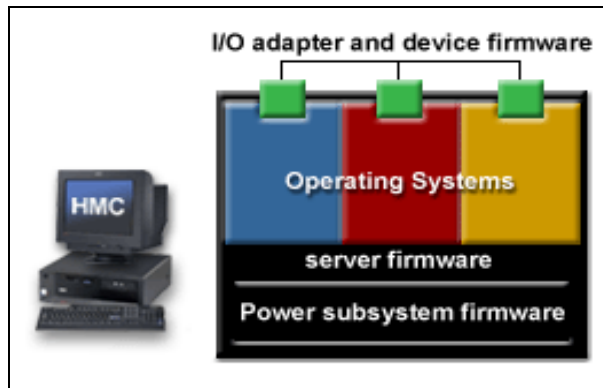


Figure 6-49 hardware and software that might require fixes

Read about each type of fix to learn more about them and to determine the best method to get fixes in your environment.

Notes: For all models except 57x and 59x model servers, if you use an HMC to manage your system and you are setting up the server for the first time or upgrading to a new server firmware release, we recommend that you install the HMC fixes before you install server firmware fixes so the HMC can handle any fixes or new function updates that you apply to the server.

For 57x and 59x model servers, you must install **HMC fixes** before you install server or power subsystem firmware fixes so the HMC can handle any fixes or new function updates that you apply to the server.

HMC fixes

Fixes are periodically released for the HMC. Configuration steps are as follows:

1. (With Internet) Ensure that you have a connection to service and support.

If you want to download HMC fixes from the service and support system or Web site to your HMC or server, you must set up a connection to service and support either through a local or remote modem or through a VPN connection. You typically set up the service connection when you first set up your server. However, the service connection is not required for initial server setup. Therefore, you must verify that the service connection exists.

To verify the service connection, follow these steps:

- a. In the navigation area, open **Service Applications**.
- b. Select **Remote Support**.
- c. Select **Customize Outbound Connectivity**.
- d. Select the tab for the type of outbound connectivity you chose for your HMC (Local Modem, Internet VPN, or Pass-Through Systems).

Notes: If a connection to service and support does not exist, set up the service connection before proceeding with this procedure.

- e. Click **Test**.
 - f. Verify that the test completes successfully. If the test is not successful, troubleshoot your connectivity and correct the problem before proceeding with this procedure. If you prefer, you can follow the “**Without Internet**” path in this procedure. You will need to obtain the fix on CD-ROM.
2. Determine existing and available HMC levels.

To determine the existing level of HMC machine code:

- a. In the Navigation Area, open the **Licensed Internal Code Maintenance** folder.
- b. Select **HMC Code Update**.
- c. In the right pane, look for the version and release of your HMC machine code in the Status area.

To find out if there are new HMC machine code fixes available, choose from the following options:

- Contact service and support.
- Use the HMC interface.
- Go to the Fix Central Web site: <http://www.ibm.com/eserver/support/fixes/>
 - i. From the Web site, click the appropriate family in the **Server** list.
 - ii. Click **Hardware Management Console** in the **Product** or fix type list.
 - iii. Click **Continue**. Information on HMC fixes appears.

3. (Without Internet) Obtain fixes.

If you do not have an Internet connection from the HMC, you must obtain the HMC fixes either on CD-ROM or on an FTP server.

You can get HMC fixes in several ways, depending on how your HMC is set up. Read about each of the following methods to get the HMC fixes, and choose one of the methods:

- Order optical media (CD-ROM).

There are two options for ordering the fix on CD.

- Contact service and support to order optical media (CD-ROM) with the fixes you need. After you order and receive the optical media, go to Step 4. Back up the HMC.
- Go to the Fix Central Web site: <http://www.ibm.com/eserver/support/fixes/>
 - i. From the Web site, click the appropriate family in the **Server** list.
 - ii. Click **Hardware Management Console** in the **Product** or fix type list.
 - iii. Click **Continue**. Information on HMC fixes appears.
 - iv. Select the version under fixes and updates.
 - v. In the **Order CD** column, select **Go**.
 - vi. Click **Continue** to sign in with your IBM ID.
 - vii. Follow the on-screen prompts to submit your order.

- Download fixes from a Web site to an FTP server that can accept an FTP request from your HMC.

To use this method, your HMC must be connected to an open network. This method requires two steps. First, you go to a Web site from which you download the fixes to the FTP server. Second, you use the HMC interface to install the fixes from the FTP server to the HMC. Follow these steps to download the HMC machine code fixes to an FTP server:

- Go to the Fix Central Web site: <http://www.ibm.com/eserver/support/fixes/>
- Click the appropriate family in the **Server** list.
- Click **Hardware Management Console** in the **Product** or fix type list.
- Click **Continue**.
- After you download the fixes from the Web site to your FTP server, go to Step 4. Back up the HMC.

4. Back up the HMC.

Before you install the fixes, you should back up critical console information on your HMC.

5. Use the Install Corrective Service wizard to update the HMC machine code.

To install HMC machine code fixes, follow these steps:

- If you created a CD-ROM for this procedure, insert it into the optical media drive on the HMC.
 - In the navigation area, expand the **Licensed Internal Code Maintenance** folder.
 - Select **HMC Code Update**.
 - Select **Install Corrective Service**. Corrective service refers to the HMC machine code fix.
 - On the Install Corrective Service window, choose from the following options:
 - If you created a CD-ROM for this procedure, select **Apply corrective service from removable media** and continue with step 6.
 - If you want to download the fix from service and support's Web site directly to your HMC, or from an FTP server to your HMC, then select **Download the corrective service file from a remote system**, and then apply downloaded service file and type related information about the Web site or FTP server.
 - Click **OK**.
 - Follow the instructions to install the fixes.
 - Reboot the HMC for the fixes to take effect.
6. Verify that the fix installed successfully.

To verify that the HMC fix installed successfully, follow these steps:

- In the Navigation Area, open the **Licensed Internal Code Maintenance** folder.
- Select **HMC Code Update**.
- In the right pane, look for the version and release of your HMC machine code in the Status area. Verify that the version and release match the fix that you installed.

Firmware (Licensed Internal Code) fixes

This topic describes the following types of firmware (Licensed Internal Code) fixes:

► **Server firmware:**

Server firmware is the part of the Licensed Internal Code that enables hardware, such as the service processor. Check for available server firmware fixes regularly, and download and install the fixes if necessary. Depending on your service environment, you can download, install, and manage your server firmware fixes using different interfaces and methods, including the HMC or by using functions specific to your operating system. However, if you have a 57x or 59x model server, that is managed by an HMC, you must use the HMC to install server firmware fixes.

► **Power subsystem firmware:**

Power subsystem firmware is the part of the Licensed Internal Code that enables the power subsystem hardware in the model 57x or 59x servers. You must use an HMC to update or upgrade power subsystem firmware fixes.

► **I/O adapter and device firmware fixes:**

I/O adapter and device firmware is the part of the Licensed Internal Code that enables hardware, such as Ethernet PCI adapters or disk drives.

i5/OS I/O adapter and device firmware PTFs for i5/OS partitions are ordered, packaged, delivered, and installed as part of the Licensed Internal Code using the same processes that apply to i5/OS PTFs. Regardless of whether you use an HMC to manage your server, you use the usual i5/OS PTF installation functions on each logical partition to download and install the I/O adapter and device firmware fixes.

Operating system fixes

To keep your partitions running smoothly, it is important that you install fixes for your operating system code when fixes are available.

You can access fixes for your operating systems using the Fix Central Web site:

<http://www.ibm.com/eserver/support/fixes/>

From the Web site, follow these steps:

1. In the **Server** list, click the appropriate family. For example, iSeries family or pSeries family.
2. In the **Product** or fix type list, click the operating system for which you want to get a fix. Depending on your selections for **Server** and **Product** or fix type, you might see additional lists from which you can select specific options.
3. Click **Continue**.

For more information about fixes for the operating systems, see the following Web sites:

- **AIX:** Go to the Support for eServer p5 and pSeries products Web site:

<http://www-1.ibm.com/servers/eserver/support/pseries/>

- **i5/OS:** Go to *Maintain and manage i5/OS and related software*.

- **Linux:** Go to the Support for eServer p5 and pSeries products Web site:

<http://www-1.ibm.com/servers/eserver/support/pseries/>

FRU SVCPROC-replace SP

The service processor is failing. After you have replaced the part, set the configuration ID for SPCN before powering up, otherwise the machine will not IPL.

You can change the processing unit identifier, also referred to as the processing unit SPCN (system power control network) ID. The processing unit SPCN ID is used by the SPCN firmware to identify the system type. It is also used to identify the primary service processor if there are two service processors in the system.

To perform this operation, your authority level must be one of the following:

- ▶ Administrator
- ▶ Authorized service provider

Notes: This feature is available only when the system is powered off. This operation resets the service processor

To change the processing unit identifier, do the following operations:

- ▶ On the ASMI Welcome pane, specify your user ID and password, and click **Log In**.
- ▶ In the navigation area, expand **System Configuration**.
- ▶ Select **Processing Unit Identifier**.
- ▶ Enter the desired information into the 2-character text area. Then click **OK**.

6.10 HMC boot up fails with “fsck”

There are certain circumstances where the HMC can fail and leave the file subsystem corrupt or damaged. Power failure is the most common form of incident, caused by inadvertently hitting the white power button, pulling the power cable, or power supply failure.

The HMC has the ability to recover itself in most incidents. Consider a failure when the HMC was performing a normal shutdown. During boot up, the HMC could stop with the following message displayed in the HMC windows:

```
'fsck failed please repair manually. login with root
```

In this case, *root* is the operating system Root password. If you have not completed the Guided Setup, the default password will still be active, the default password is '**passw0rd**' and is case sensitive. If you have completed the Guided Setup, you will have changed the default root password. You should use this new password.

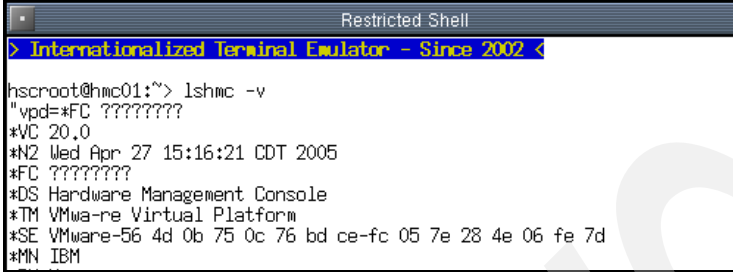
Once you have logged in with root password, you will be presented with a command prompt. At the prompt type in '**fsck**' and press Enter. You will see the **fsck** process running in line mode. You will be asked if the OS should repair certain components. You should say **yes** to all prompts unless you are instructed otherwise. Once **fsck** completes, you may automatically return to the GUI signon screen. If you are not returned to the GUI, you should type '**reboot**' at the prompt and press Enter. The HMC GUI signon screen will appear.

When you sign onto the HMC and the management pane appears, you may have to reconnect or find the managed server if it does not automatically appear.

6.11 Determining the HMC serial number

For some HMC or Service Processor troubleshooting situations, a Product Engineer (PE) will have to sign on to the HMC. The PE password changes daily and is not available for normal customer use. If the PE determines that a local service engineer can sign on to the HMC, the PE may request the HMC serial number.

To find the HMC serial number, open a restricted shell window and run the following command, `lshmc -v`. Figure 6-50 is an example of the information displayed.



```
Restricted Shell
> Internationalized Terminal Emulator - Since 2002 <
hscroot@hmc01:~> lshmc -v
"vpc=#FC ????????"
*VC 20.0
*N2 Wed Apr 27 15:16:21 CDT 2005
*FC ????????"
*DS Hardware Management Console
*TM VMwa-re Virtual Platform
*SE VMware-56 4d 0b 75 0c 76 bd ce-fc 05 7e 28 4e 06 fe 7d
*MN IBM
```

Figure 6-50 `lshmc -v` command detail

The `*SE` field is the HMC serial number.

Archived



HMC security and user management

In this chapter we discuss security implementation within the HMC environment.

The following topics are described:

- ▶ Certificate authority
- ▶ Server security
- ▶ Object manager security
- ▶ HMC User management

7.1 System Manager Security

System Manager Security ensures 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. Each HMC System Manager server has its own private key and a certificate of its public key signed by a Certificate Authority (CA) that is trusted by the System Manager clients. The private key and the server certificate are stored in the server's private key ring file. Each client must have a public key ring file that contains the certificate of the trusted CA.

You define one HMC as a Certificate Authority. You use this HMC to generate keys and certificates for your HMC servers and client systems. The servers are the HMCs you want to manage remotely. A unique key must be generated and installed on each server. You can generate the keys for all your servers in one action on the CA and then copy them to diskette, install them at the servers, and configure the servers for secure operation.

The client systems are the systems from which you want to do remote management. Client systems can be HMCs, AIX, or PC clients. Each client system must have a copy of the CA's public key ring file in its System Manager codebase directory. You can copy the CA public key ring file to the diskette on the CA and copy it from the diskette to each client.

Note: To configure an AIX client correctly, you must install a security fileset. For more information, see your AIX documentation.

To use the System Manager Security application, you must be a member of the System Administrator role. To ensure security during configuration, users of this application must be logged in to the HMC locally.

Figure 7-1 shows the functions available from the System Manager Security.

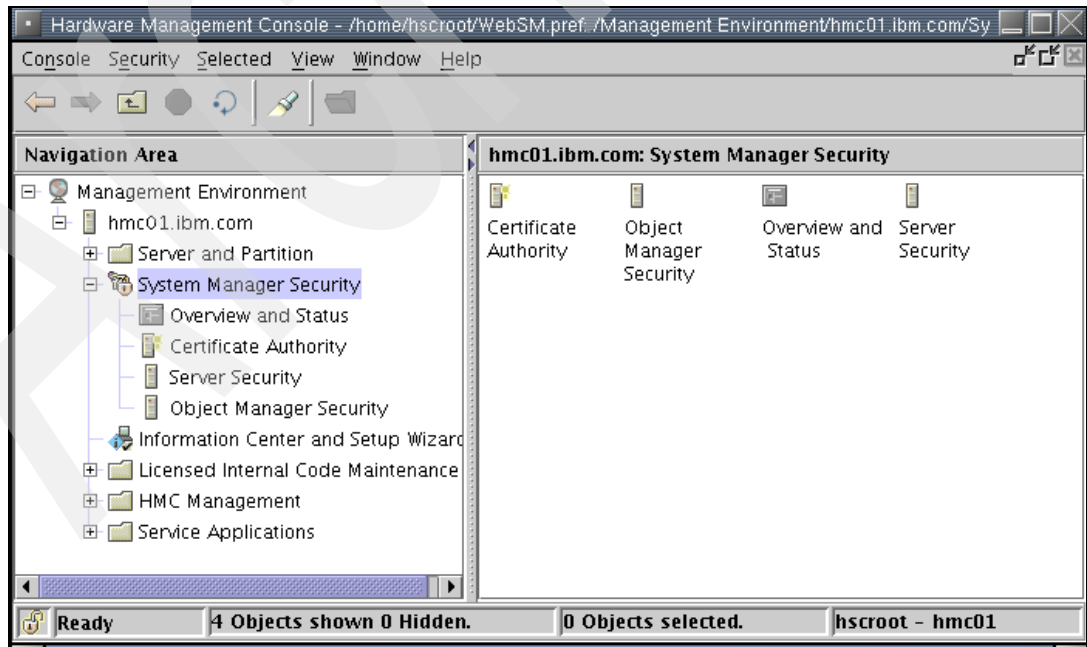


Figure 7-1 System Manager Security available functions

The following sections describe the System Manager Security functions that are available.

7.2 Overview and status

The overview and status window displays the following information about the secure system manager server:

- ▶ Whether the secure system manager server is configured
- ▶ Whether the private key for this system manager server is installed
- ▶ Whether this system is configured as a Certificate Authority

7.2.1 Certificate Authority (CA)

Note: You cannot perform the following function using a remote client.

Define one HMC as a Certificate Authority (CA) to generate keys and certificates for your HMC servers and clients.

A Certificate Authority verifies the identities of the HMC servers to ensure secure communications between clients and servers. To define a system as a Certificate Authority, you must be logged in as the **hscroot** user at the machine being defined as the internal Certificate Authority. This procedure defines a system as an internal Certificate Authority for HMC security and creates a public key ring file for the Certificate Authority that you can distribute to all of the clients that access the HMC servers.

A wizard guides you through configuring the Certificate Authority.

After you define the internal Certificate Authority, you can use the CA to create the private key ring files for the HMCs that you want to manage remotely. Each HMC server must have its private key and a certificate of its public key signed by a Certificate Authority that is trusted by the HMC clients. The private key and the server certificate are stored in the server's private key ring file. There is an option to copy the private key ring files to a diskette so you can install them on your servers.

7.2.2 Server Security

This option allows you to install the private key ring file that you have copied to diskette from the HMC server that is acting as the Certificate Authority. Once you have copied the private key file, there is another option to configure the HMC as a secure server so that secure, remote clients can be used to remotely manage the HMC.

There is a remote client available for download from the HMC itself. It is called the Web-based System Management remote client, and there is a Windows based version and a Linux based version. To run in secure mode, a second file needs to be downloaded to the client. This is also available for download from the HMC.

To download the Web-based System Management remote client to your Windows based or Linux based PC, type in the following address from your Web Browser:

hostname/remote_client.html

Here, *hostname* is the name of the HMC you are downloading the Web-based System Management remote client from. You choose whether you want the Windows based version or the Linux based version.

To download the security package so that the client/server connection, that is, the PC to HMC, is secure, type in the following address in your Web Browser:

`hostname/remote_client_security.html`

Once again you choose whether you want the Windows based version or the Linux based version.

7.2.3 Object Manager Security

The HMC Object Manager Security mode can be configured as either Plain Socket or Secure Sockets Layer (SSL). By default, the Plain Sockets mode is used. For SSL mode, the Object Manager reuses the HMC System manager server's private key ring. The server private ring and the Certificate Authority's public key ring must be installed when establishing the SSL connection.

7.3 HMC User Management

The HMC management option allows you to create and manage HMC user profiles and to configure the HMC. Some of this configuration is done when the setup wizard is initially run to set up the HMC. The options under HMC Management allow you to go in and modify the configuration that was initially set up.

7.3.1 HMC Users

To use this function, select **HMC Management** → **HMC Users**.

This option allows you to perform the following functions for users:

- ▶ Creating a user
- ▶ Editing user information
- ▶ Viewing user information
- ▶ Deleting a user
- ▶ Changing passwords

You must be the system administrator to perform the functions listed above. Each HMC user can be a member of one to six different roles. Each of these roles allows the user to access different parts of the HMC. The user roles as specified by the HMC are as follows:

- ▶ Super Administrator
- ▶ Operator
- ▶ Service Representative
- ▶ Product Engineer
- ▶ Viewer

Each role is described in the following sections.

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.

Operator

An operator is responsible for the daily system operation.

Service Representative

A service representative is the person who installs or repairs the system.

Product Engineer

The product engineer will assist in support situations, but the profile cannot be used to access user management functions in HMC.

Viewer

A viewer can view HMC information, but cannot change any configuration information.

7.3.2 Add, modify, copy, or remove user profile

This section shows you how to add, modify, copy or remove a user profile with various user roles as described in 7.3.1, “HMC Users” on page 238.

To access the User Profiles configuration for each user, select **HMC Management** → **HMC Users**. The HMC will display the User Profiles window as shown in Figure 7-2.

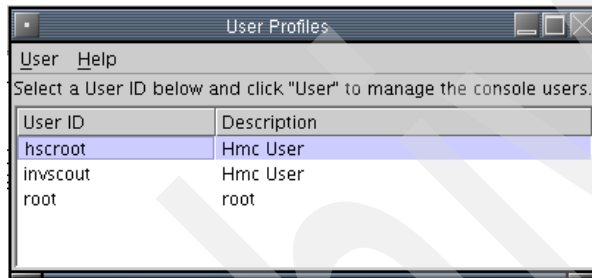


Figure 7-2 User Profiles window

Add a new user with Super Administrator role

To add a new user with the Super Administrator role, perform the following steps:

1. Select the **User** menu, and click **Add**. to add a new user. The HMC will display the Add User window as shown in Figure 7-3.

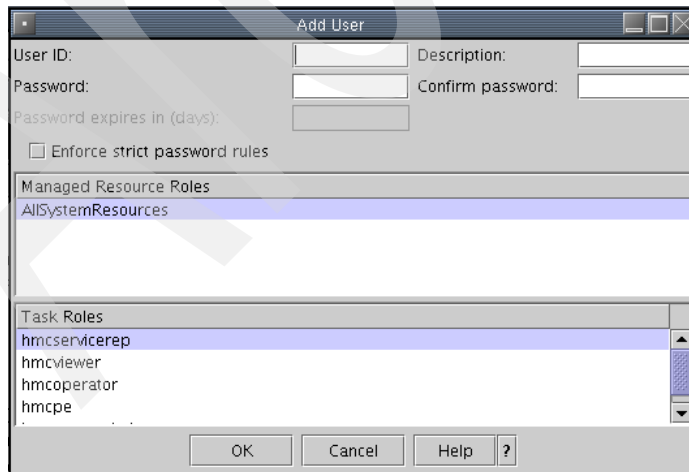


Figure 7-3 Add User window

2. Insert the new user ID, a description of the user ID, the password for the new user ID, and re-type the new password (Figure 7-4).
3. Select **hmcsuperadmin** from Task Roles in order to create a new user with the System Administrator role. You may select **Enforce strict password rules** to give the password expiration and type the number of the expiration day as shown in Figure 7-5. **Enforce strict password rules** will set the password to expire after the number of the day specified.




Figure 7-4 Add a new user with Super Administrator role

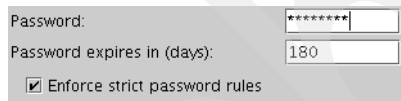


Figure 7-5 Enforce strict password rules

4. Click **OK** to create a new user. The new user ID will be added in User profiles window as shown in Figure 7-6.




Figure 7-6 User Profiles user ID list updated

Add a new user with Viewer role

If you want to grant a person permission to sign in to the HMC and view the HMC and system configuration and status, but not to make any changes, you can create a User ID with the Viewer role, which has the capability to view only the HMC and system configuration and status.

To add a new user with the Service Representative role, perform the following steps:

1. Select **User** menu, click **Add**. to add a new user. The HMC will display the Add User window as shown in Figure 7-3.
2. Insert the new user ID, a description of the user ID, the password for the new user ID, and re-type the new password.
3. Select **hmcviewer** from Task Roles in order to create a new user with the Viewer role as shown in Figure 7-7. You may select **Enforce strict password rules** to give the password expiration and type the number of the expiration day as shown in Figure 7-5 on page 240. **Enforce strict password rules** will set the password to expire after the number of the day specified.

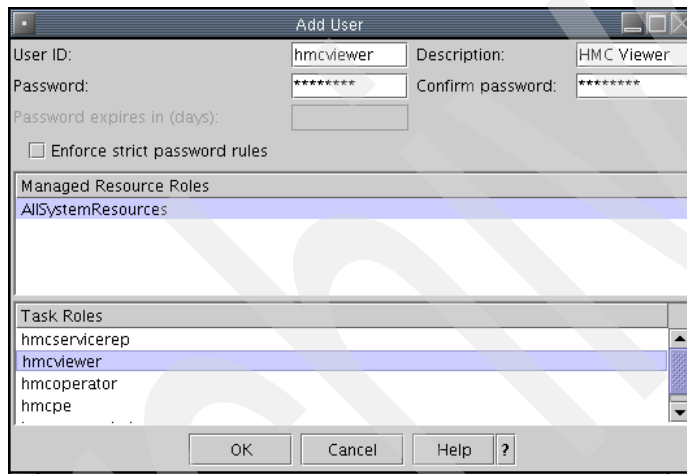


Figure 7-7 Add a new user with Viewer role

4. Click **OK** to create a new user. The new user ID will be added in the User profiles window.

Once created and used to sign in to the HMC, the HMC Viewer user can only view the configurations of managed system and logical partition configuration status. It has no authority to run any administrative tasks such as HMC Configuration or Licensed Internal Code updates. The HMC will notify the HMC Viewer user if it has no sufficient authority to access administrative interfaces, as shown in Figure 7-8.




Figure 7-8 HMC notification of lack of sufficient authority

The HMC Viewer is only given very limited access to functions in the HMC. Figure 7-9 shows the limited menu for the HMC Viewer.

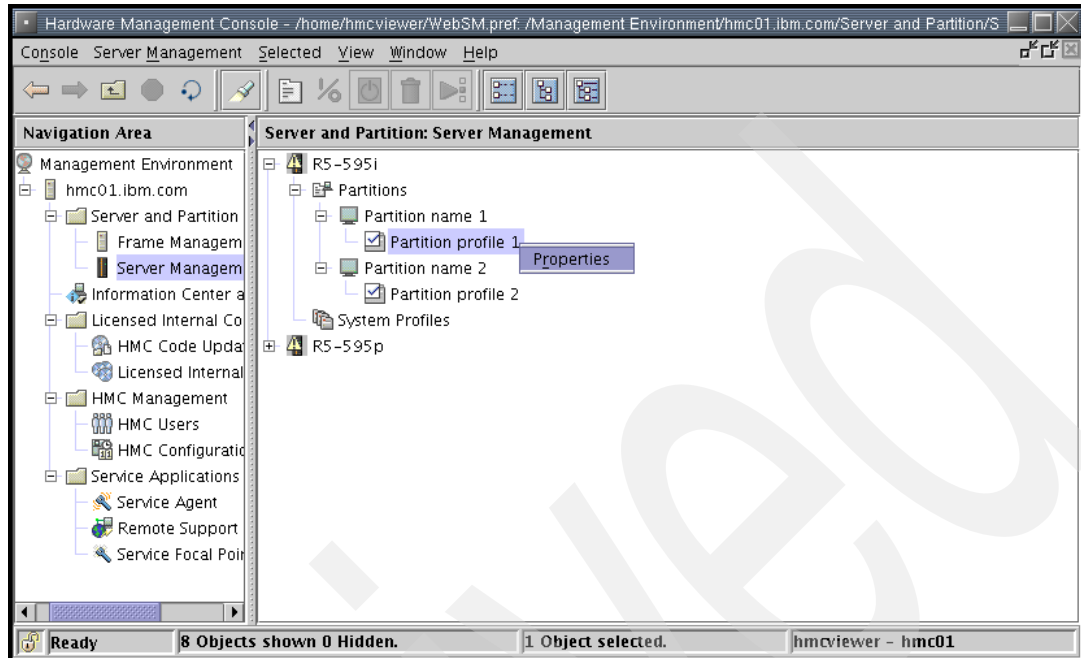


Figure 7-9 Very limited menu available for HMC Viewer user

7.3.3 Customizing user task roles and managed resource roles

You can customize HMC Task Roles and Managed Resource Roles via the HMC console. You can add new Task Roles and Managed Resource Roles based on existing roles in the HMC. System defined roles cannot be modified, but you can create a new role based on system defined role or existing role. To manage access task and resource roles, select **Manage Access Task Roles and Managed Resource Roles** on the HMC Users pane. The Customize User Controls window will be displayed as shown in Figure 7-10.

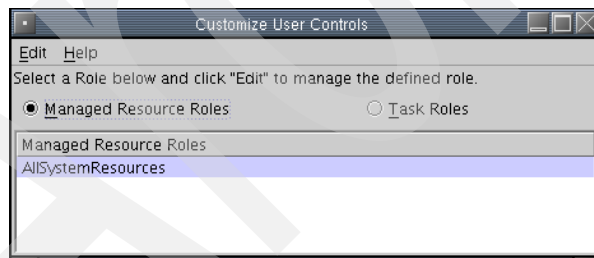


Figure 7-10 Customize User Controls window

Create, copy, modify, or delete managed resource roles

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

You can create a new managed resource role, copy 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 desired operation from the **Edit** menu. By default, there is only one managed resource role: it is *AllSystemResources*.

To create a new managed resource role, click **Edit** → **Add**, and the Add Role window will be displayed. Fill in the name for the new managed resource role, and choose from which resource role the new managed resource role objects will be based on. Select which object will be 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 new managed resource role. An example of creating a new managed resource role is shown in Figure 7-11.

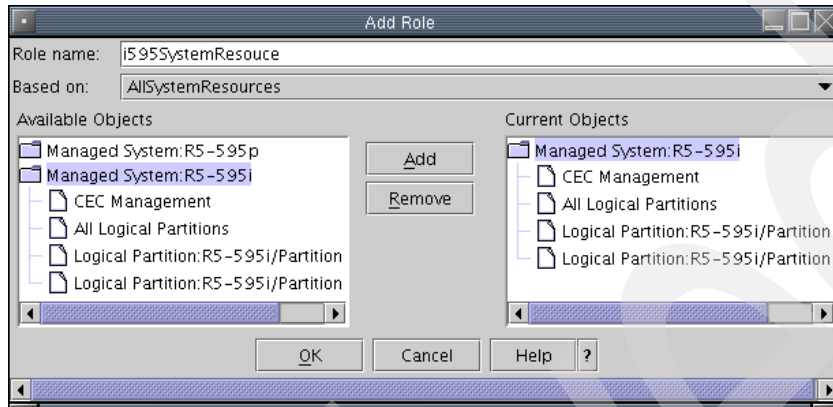


Figure 7-11 Add a new managed resource role

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

To delete a managed resource role, select desired managed resource role and select **Edit** → **Remove**. A verification window will be shown as in Figure 7-12.




Figure 7-12 Delete managed resource role verification

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

Create, copy, modify, or delete task roles

A task role defines the access level for a user to perform tasks on the managed object or group of objects, such as a managed system or logical partition. There are five system defined task roles: *hmcservicerep*, *hmcviewer*, *hmcoperator*, *hmcpe*, and *hmcsuperadmin*.

You can create a new 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 desired operation from the **Edit** menu.

To create a new user task role, click **Edit** → **Add**, and the Add Role window will be displayed. Fill in the name for the new managed resource role, and choose from which task role the new task role objects will be based on. Select which object will be available for the new task role, then click **Add**. to add them to new task role current objects. Click **OK** to create a new task role. An example of creating a new task role is shown in Figure 7-13.

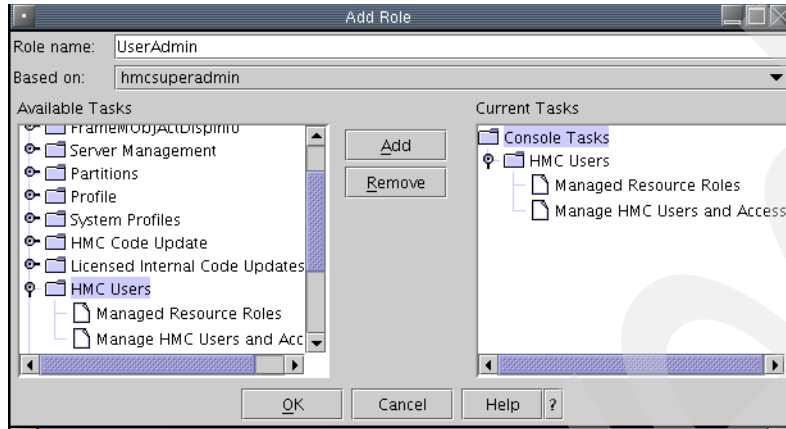


Figure 7-13 Create new task role window

To copy a task role, select the desired 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 verification window will be shown as in Figure 7-12 on page 243. System defined task roles cannot be removed.

To modify existing task roles, select a task role you wish to change, and select **Edit** → **Modify**. You may change the objects' configuration, then click **OK** to save the changes. Only user defined task roles that are created by HMC users can be modified.

HMC duplication and redundancy

In this chapter we introduce HMC cloning — attaching redundant HMC devices to a single i5 managed system, multiple managed systems, or a single HMC to multiple managed systems. Although the System i5 systems will function properly with the HMC disconnected, it is the only interface for managing a partitioned system and thus is a key component.

This introduction provides a closer look at:

- ▶ Redundant HMC configurations
- ▶ Multiple managed system configurations
- ▶ Cloning HMC configurations
- ▶ Redundant HMC configuration considerations

8.1 Redundant HMC configurations

From late August 2004, there have been improvements in HMC options and support. One HMC can manage 48 servers, and each HMC can manage 254 i5/OS partitions.

For redundancy of the system management control point, you can configure a redundant HMC configuration. As shown in Figure 8-1, there are two HMC devices directly attached to the ethernet ports of the HMC. This configuration provides a very secure connection for both devices.

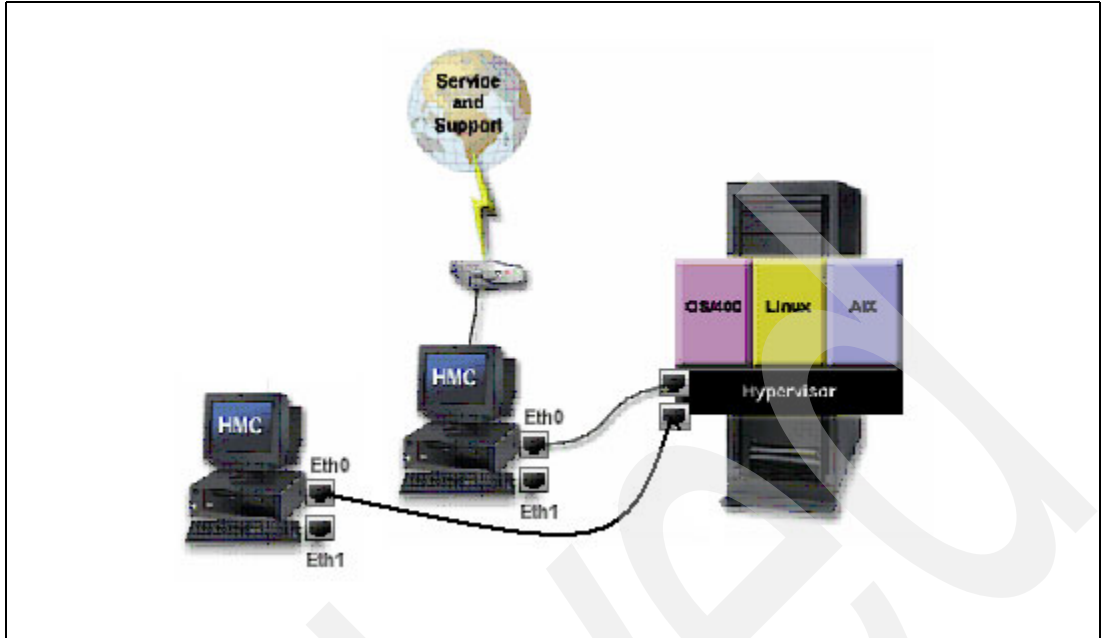


Figure 8-1 Two HMCs directly attached

Figure 8-2 shows the same redundant HMCs connected via a hub. This configuration allows attachment of other devices. We recommend using a private network.

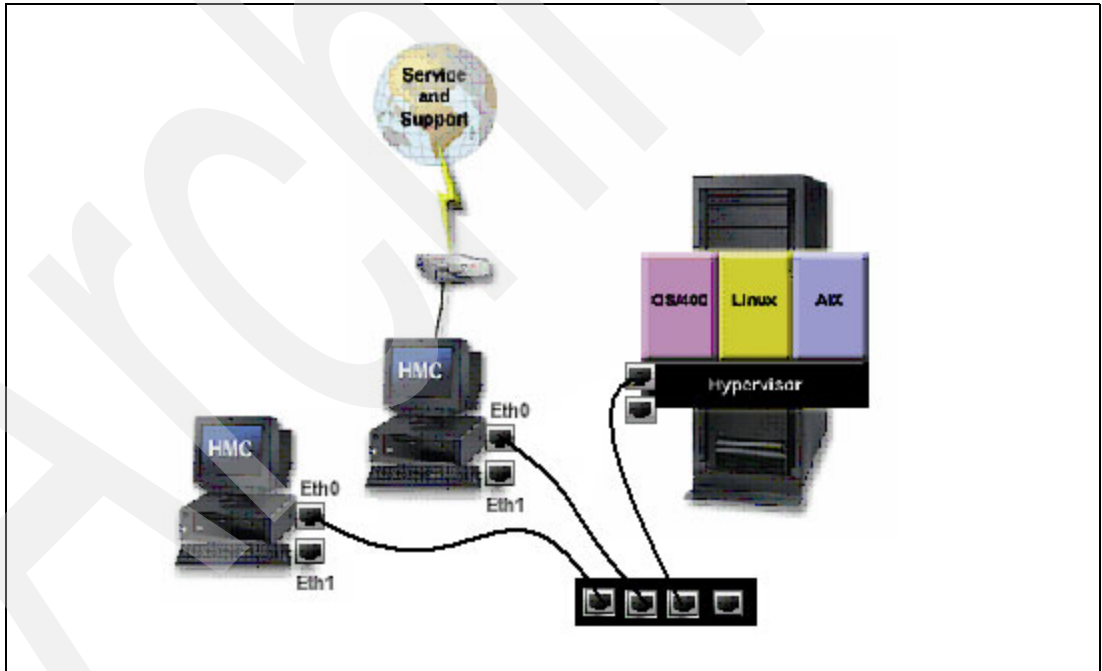


Figure 8-2 Two hub attached HMCs

8.2 Multiple managed system configuration

To save space and to centralize multiple system management control points, you can configure up to 48 managed systems on a single HMC. In Figure 8-3 and Figure 8-4, we show the current HMC management scenario of one HMC managing two servers and then one HMC managing three servers.

As you increase the number of managed servers, you will definitely need to introduce hubs/switches into your private network. You should plan for redundancy and maintenance in these hubs and switches, as the network is crucial to system management.

Figure 8-3 shows one HMC managing two iSeries servers.




Figure 8-3 HMC with two managed systems

Figure 8-4 shows one HMC managing three iSeries servers.




Figure 8-4 HMC with three managed systems

8.3 Cloning HMC configurations

System Profiles and their associated Partition Profiles are stored in NVRAM of the Service Processor (SP). When a redundant HMC is connected to an System i5 system, with valid partitions, the System Profile and Partition Profile information is automatically downloaded to the redundant HMC when the HMC is powered on. The System i5 system SP detects the additional HMC and moves the data from NVRAM. User preferences and user profiles are not stored in NVRAM. This information is stored on the hard drive of the HMC. If you want a mirror image of this data, you must restore it from your Critical Console Backup taken from the HMC you want to clone.

In a multiple managed system environment, each system will download its current System Profile and Partition Profile to the HMC. This new information does not overlay what is already on the HMC. A new system profile is added to the existing profiles listed on the HMC Server Management window, as seen in Figure 8-5.

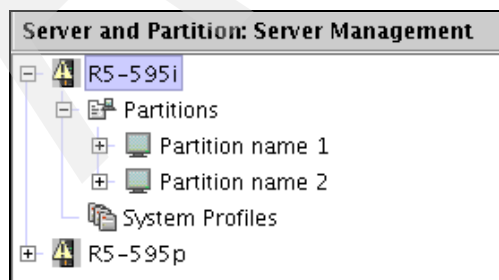


Figure 8-5 Server and Partition: Server Management window

8.4 Redundant HMC configuration considerations

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

Both HMCs can be used concurrently. You have to 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 must 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 allows one HMC to run at the new fix level, while the other HMC can continue to run at the previous fix level. However, the best 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 issued from two HMCs in the redundant HMC configuration. A locking mechanism provided by the service processor allows inter-operation in a parallel environment. This 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 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 issued by one HMC are visible in the other. For example, if you choose to activate a partition from one HMC, you will 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 will power off. Effectively, any sequence of commands that you can do from a single HMC is also permitted 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 may be necessary to manually release the lock.

When running two HMCs to the same server, you should also be careful with long running functions, as they may be impacted if they have not completed before an additional function is run on the second HMC.

Archived

Migration of existing LPAR profiles to HMC

Users of logical partitioning have a new interface for managing LPAR on System i5 systems. The *Work with System Partitions* option available in V4R4/V4R5/V5R1/V5R2 is no longer available in V5R3 running on the new System i5 hardware. This option has been moved from DST/SST to the Hardware Management Console (HMC). The HMC is required for any System i5 that is running logical partitioning. All LPAR configuration profiles must be created and managed from the HMC.

In this chapter the following topics are discussed:

- ▶ Migration planning
- ▶ Preparing your system for migration to HMC
- ▶ Migrating Linux partitions
- ▶ Migration example scenarios
- ▶ Backing up HMC profiles

9.1 Migration planning

This section focuses on the planning activities required to move POWER4™ LPAR configurations to System i5 HMC LPAR profiles. You must understand your current LPAR environment and have accurate documentation of how your system is configured and an inventory of the resources that are presently in use. There are several tools and tasks that can be used to document the current system environment and plan for a successful migration. The following ones are recommended:

- ▶ Use DST or SST to print a copy of your current system configuration.
- ▶ Get a printout of your current system resource allocations.
- ▶ Review the resources in the Primary partition.
- ▶ Review bus numbering.
- ▶ Use the LPAR Validation Tool (LVT) to create a graphical worksheet of your partitions.
- ▶ Review your current system configuration to identify unsupported hardware at V5R3 and determine a migration plan.
- ▶ Review this redbook to understand V5R3 changes and System i5 enhancements.
- ▶ Create a customized migration checklist.
- ▶ Schedule a pre sale Solution Assurance Review (SAR).
- ▶ Order the System i5 system.

9.1.1 Use DST or SST to print a copy of your current system configuration

Before upgrading an existing LPAR system, it is necessary to get the rack configuration of the system with the details of each logical partition. This information is available through the *Display Partition Information* screens which are available in *Dedicated Services Tools* (DST) or *System Service Tools* (SST).

The following steps document the process to print a rack configuration for a partitioned system running V5R1/V5R2. It is a single pass that generates a multi-part report. The first section looks very much like a normal rack configuration and can be used to verify the system inventory. The remaining sections display the hardware assigned to each partition.

Note: The following procedure includes certain steps to clear out the non-reporting hardware prior to printing the rack configuration from DST. This is necessary because there is no way to determine from the printout if the hardware listed is actually there or not. If you feel comfortable that there is no non-reporting hardware, you can perform these tasks from SST. Skip steps 1, 2, and 3 and simply start with step 4.

Perform the following steps:

1. Access the DST console environment. Place the Control Panel in **B Manual** mode and then use the up/down arrow keys to select function **21**. Press Enter. The DST menu displays on the system console. Log on the system with a valid DST user ID and password.
2. From the Work with Partitions screen, select Option **4, Recover Configuration Data**.
3. Clear out old hardware from the partition table by selecting Option **4, Clear Non-reporting Logical Partitioning Resources**.

Note: All partitions must be active to select this option.

4. On the Work with Partitions screen, select Option 1, **Display Partition Information**.
5. Select Option 5, **Display System I/O Resources**.
6. Use the F6 Print Key to print a copy of the rack configuration. This printout contains all of the information usually collected from the *Hardware Service Manager* plus all of the partition information.

9.1.2 Get a printout of your resource allocations

To document your resource allocations, perform the following steps:

1. Start SST by entering **STRSST**.
2. Log on using a valid Service Tools user ID and password.
3. From the Service Tools menu, select Option 5, **Work with system partitions**.
4. From the Work with System Partitions menu, select Option 1, **Display partition information**.
5. Press F6 from the Display Partition Information menu to get a printout of your resource allocation information. F6 prints information from each of the listed options on the Display Partition Information menu.
6. Use the F12 key to logout of SST.

9.1.3 Review the resources in the Primary partition

Review the resources in the Primary partition. A Primary/P0 partition no longer exists in a System i5 system at V5R3. If the Primary/P0 partition is migrated, it will be assigned the next available partition number.

For example, a 4 LPAR system (P0, P1, P2, P3), P0 will be migrated as P4. Or in a 2 processor system (maximum is 10 LPAR per processor) with 20 LPAR (P0 through P19), P0 will be migrated as P20. Care must be taken to assure that the new system is ordered with sufficient resources to account for any loss of hardware resources as a result of the migration (replace unsupported hardware like migrations towers, etc.). Planning is critical for customers who use the Primary partition as a production partition. The best migration plan would relocate the Primary partition resources to the new System i5 system intact so that a scratch install is avoided.

9.1.4 Bus numbering

It is important to understand that the bus numbering of the system unit will change after the upgrade. This is critical to preserve or to update logical partition configurations after the migration. For example, let us consider an 810 to System i5 520 upgrade. The 810 system unit has bus numbers 1 and 2. After the upgrade to the System i5 520, the new system unit will contain buses 2 and 3. At installation time, the IBM Customer Service representative (CSR) will take steps to prevent the re-numbering of the other existing buses.

9.1.5 Use the LVT tool to create a graphical worksheet

This task is optional but recommended. If you would like to create a graphical layout of your partitioned system, use the *Logical Partitioning Validation Tool* (LVT). The LVT is available to assist the user in the design of LPAR systems and to provide an LPAR validation report that reflects the user's system requirements while not exceeding LPAR recommend. This is also a good method of inventorying your system and it provides a picture of slot layout and usage.

To download the tool, visit the *IBM eServer iSeries Support Web* site and look for *iSeries Tools*. Use the following URL:

<http://www-912.ibm.com>

Note: The LVT is *not* a marketing configurator. It does not automatically add hardware features (except base and primary CD or DVD). It will not prevent inefficient system design so long as the design itself meets manufacturing card placement rule and minimum LPAR recommendations.

9.1.6 Review the list of supported and unsupported hardware

With every new release of the OS400 Operating System and new hardware announcement there are changes that inevitably cause some products and features to become obsolete or unsupported. An example of this is the removal of support of migration towers. The System i5 platform does not support the attachment of migration towers. You must plan for this change in case replacing the migration tower requires additional expansion towers and features to be purchased. For more hardware and software planning information, visit the *IBM eServer iSeries Support Web* site under *Planning*:

<http://www-912.ibm.com>

9.1.7 Create a customized planning checklist

Create a customized planning checklist to use prior to and during the migration. Go to the *V5R3 iSeries Information Center*. From there you can run an interview wizard through EZ-Setup for iSeries. You can create a customized planning checklist by entering your proposed system data and requirements. For more information, visit the *IBM eServer iSeries Support Web* site under *Technical Databases*.

<http://www-912.ibm.com>

9.1.8 Schedule a Solutions Assurance Review

Note: This is an IBM Business Partner or IBM Sales Representative task.

Partition planning is critical to ensure that each partition is configured with sufficient CPU, memory and I/O resources to support the proposed solution. Applications may have special licensing requirements to run in an LPAR environment. Partitioned systems have additional installation and systems management requirements that need to be carefully reviewed.

If logical partitioning is part of your solution design, completion of an expert level Pre-Sale *Solution Assurance Review (SAR)* is recommended before you order your new System i5 hardware. This formal review meeting includes a technical inspection by Subject Matter Experts who were not involved in designing the solution. For more information, visit one of the following Web sites:

IBM Business Partners:

<http://www-100.ibm.com/partnerworld/software/Sa1/assur30i.nsf/Web/Learn>

IBM:

<http://w3-1.ibm.com/support/assure/assur30i.nsf/Web/SA>

9.1.9 Order the System i5

Order the necessary hardware and software based on the output from the LVT tool or validated work sheets/configurator output.

9.2 Preparing your system for migration to System i5 and HMC

Before you start the migration, complete the following required tasks:

- ▶ Review your LPAR configuration. Clean up non-reporting resources and unsupported hardware, including attached migration towers.
- ▶ Load V5R3 on all OS400 partitions.
- ▶ Update the Linux kernel to 2.6, if you are running a Linux Partition.
- ▶ iSeries Navigator should be V5R2 or higher.
- ▶ Verify that you have System Administrator or Advanced Operator authority on the Hardware Management Console. This is done during the setup of the HMC.

9.2.1 Load V5R3

V5R3 is required in all System i5 partitions. Prior to moving to System i5, you must schedule some time to load V5R3 on each OS400 partition on your current system. Review Chapter 1.2, "Software requirements" on page 5. This section discusses licensing rule changes for V5R3.

9.2.2 Update the Linux kernel

System i5 requires Version 2.6 of the Linux kernel. A new redbook is in progress that will cover the migration of existing Linux partitions to V5R3 and System i5 systems.

9.2.3 iSeries Navigator

iSeries navigator is the recommended tool to use to move/export your current LPAR configuration to the System i5 platform. The export function of iSeries Navigator converts the LPAR configuration to a binary file which can be saved to 3 1/2 diskette or DVD/RAM if one is available on your PC. Use this saved file to restore the LPAR configurations to the HMC.

9.3 Migrating Linux partitions

This topic describes the options and requirements for migrating a Linux installation from an iSeries server to an @server System i5 system.

The first step in migrating a Linux installation from an iSeries server to an @server System i5 is to upgrade to a Linux version that supports the @server System i5 system.

Follow these steps to complete the Linux upgrade:

1. On your existing iSeries server, upgrade to a new Linux version that supports the @server System i5 systems. Refer to your Linux distributor documentation for detailed instructions.
2. Replace your existing I/O device drivers with System i5 I/O device drivers. If virtual I/O is being used, it must be configured in the OS/400 partition that provides the virtual I/O device.
3. From the new Linux distribution, retrieve an @server System i5 kernel and store it in the OS/400 file system.
4. Configure the @server System i5 partitioning. See Chapter 5, "Partition creation using the HMC" on page 139 for instructions.
5. Migrate the OS/400 hosting partition to the new system.
6. Boot the Linux partition from the new Linux kernel.

For a full discussion of Linux on @server System i5 systems, review the IBM Redpaper, *Linux on i5: Implementation Guide*, SG24-6388.

<http://www.redbooks.ibm.com/redpieces/pdfs/sg246388.pdf>

9.4 Migration scenarios

This topic presents the process and tasks involved in migrating to System i5 for two scenarios. The first is an upgrade from an 810 with two OS400 logical partitions to a 520 System i5 system. The second scenario is an upgrade from an 825 with three OS400 logical partitions to a 570 System i5 system.

Attention: Some IBM services are billable to the customer. Prior to starting the upgrade, all details of the upgrade should be discussed to determine the scope of the services performed by IBM. This can be done during the pre-sale Solution Assurance Review.

9.4.1 Information gathering and pre-sales planning

In the first scenario, the current system has two OS400 LPARs (P0 and P1) and consists of the 810 system unit and one 5094 tower. All of the P0 Primary partition resources are located in the 810 system unit. All of the resources for the P1 Secondary Partition are in the 5094 tower. A Primary partition no longer exist in the System i5 systems. However, in this case there is production work done in the Primary partition of the 810 which requires this partition to be migrated to the 520.

As the information gathered is reviewed, you should consider ways to migrate to the new System i5 system with the least amount of disruption or changes. Particularly consider ways to move over intact the essential resources of the Primary partition at migration time (Console and Disk IOA and all Disks including the load source). This will greatly simplify the migration process by avoiding the reload of the partitions. Figure 9-1 is a general illustration of the Current and Proposed Systems for both scenarios.

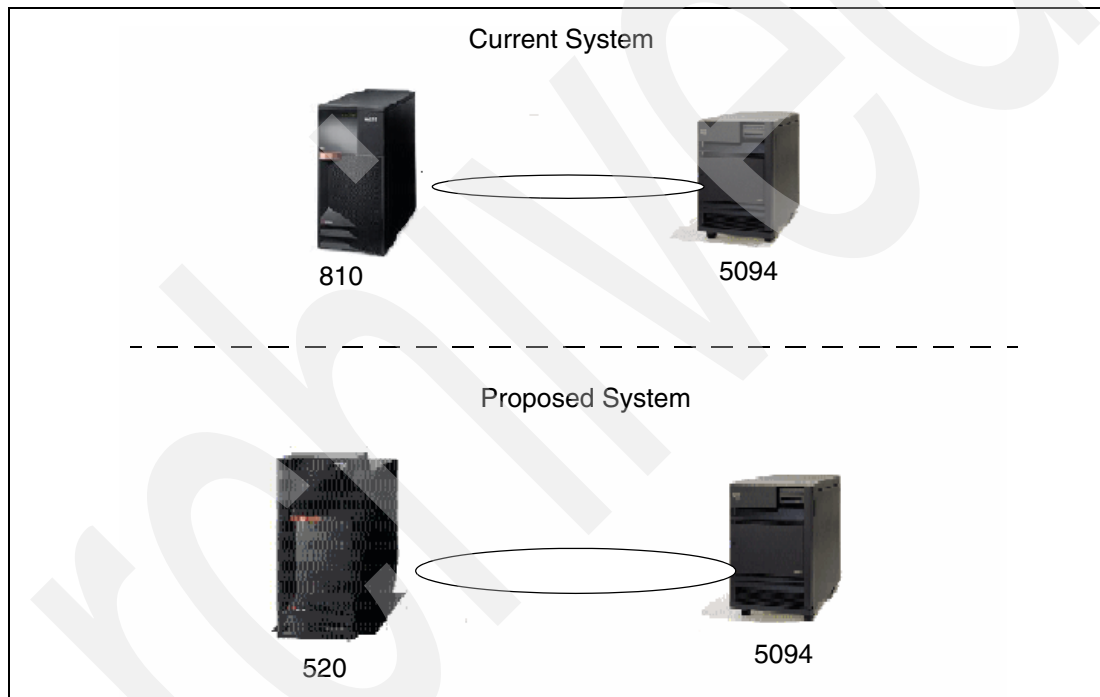


Figure 9-1 First Scenario 810 to 520 Upgrade

In the second scenario, the current system has three OS400 LPARs (P0, P1 and P2) and consists of the 825 system unit, a 5094 tower with a 5088, a 5074. The essential resources for the Primary partition are located in the 825 system unit (includes load source, disk and console IOAs). All of the resources for Secondary Partitions P1 are in the 5074 tower and all of the resources for P2 are in the 5094. The 5088 contains some switchable resources and other Primary partition resources. There is production work done in the Primary partition of the 825 which requires this partition to be migrated to the 570. Figure 9-2 is a general illustration of the current and proposed systems.




Figure 9-2 Second scenario 825 to 570 upgrade

Tips: There are new DVD features for 520 and 570. Previous DVD features do not physically fit in the system unit slots. The 520 has two DVD slots, but only one can be used for OS400. There is no longer a Console IOA placement requirement in the System i5 system unit with the exception of Ops and LAN Console. The system unit can only hold one logical partition.

The following tasks provide the necessary information to plan the migration. After reviewing this information, you will be able to plan the migration and identify any changes required to the current system configuration as well as the configuration of the proposed system.

1. Print a copy of the current LPAR configuration.
2. Print a copy of the current resource allocation.
3. Review the current LPAR configuration and resource allocation printouts to create a graphical worksheet for the current system configuration using the LVT Tool as follows:
 - a. Select **File, New** then the **System Type** from the pull down list (Figure 9-3).

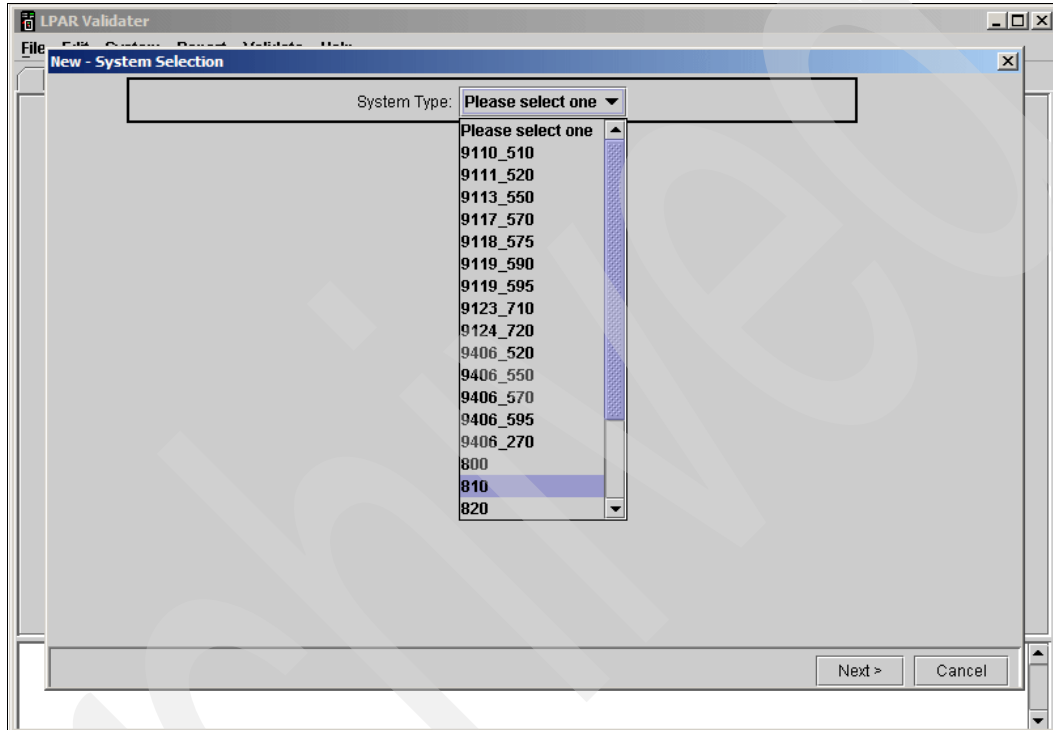


Figure 9-3 Example of LVT initial window

- b. Complete the **System Selection** window (each option is described below) and click **Next** (Figure 9-4):
 - i. **System Type:** Select the iSeries model (810 in this scenario).
 - ii. **Primary Partition OS Level:** Select Operating System Version and Release that will be used by the Primary partition (this option will not be available in System i5 models as a Primary partition no longer exists).
 - iii. **Interactive Feature:** Select Interactive Feature of your iSeries model (automatically selected for i800, 810, 825, 870, and some 890 models depending on the Server Edition).
 - iv. **System Memory:** Enter Total system memory of your iSeries.
 - v. **Number of Partitions:** Enter total number of Logical Partitions of your iSeries.
 - vi. **Secondary Partition Default Level:** Select Operating System Version and Release that will be used in the Secondary Partitions (select **None** if a different Release or Operating System will be used).
 - vii. **Total Number of Processors:** Select total number of processors available in the iSeries (this will vary on Capacity of Demand iSeries models).

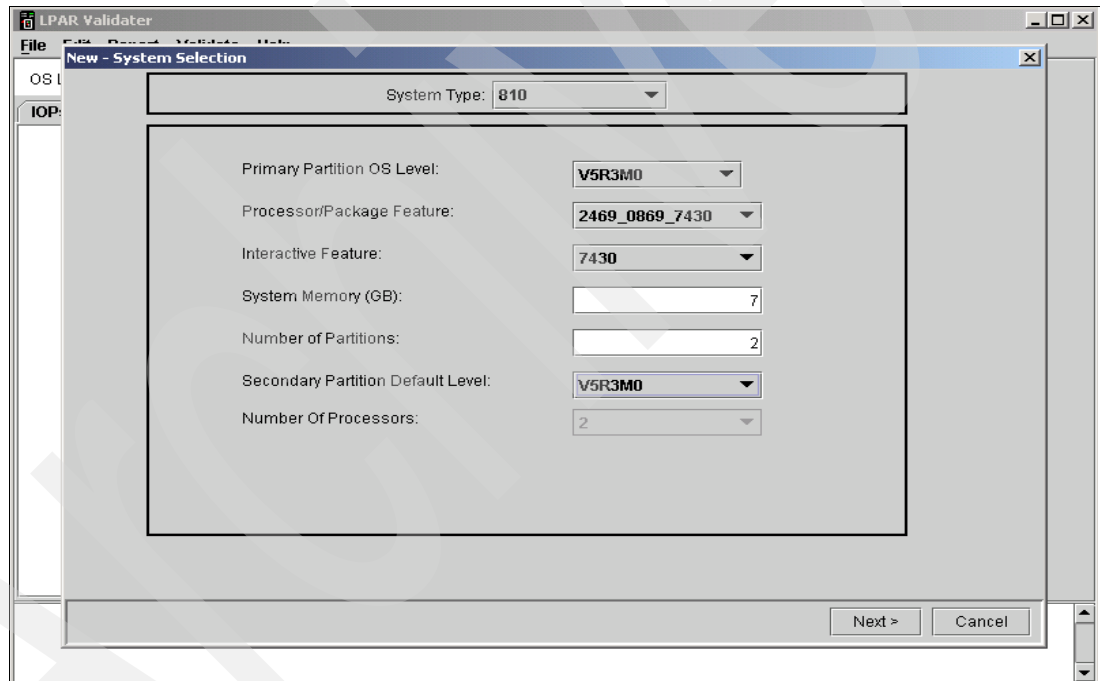


Figure 9-4 Example window System Selection options

- c. Complete the **Partition Specifications** window (each option described below) and click **Finish** (Figure 9-5).
 - i. **Primary Partition Console Type:** Select the console type that will be used by the Primary Partition (i.e. 9793 for Operations Console or 4746 for twinaxial, etc.)
 - ii. **Shared Processor Pool:** Enter amount of processor that will be shared among partitions.
 - iii. **Shared** checkbox: Check this box if the partition will be using shared processors
 - iv. **# Processors:** Indicate number of processors required by the partition. If Shared, the number can be decimal value. A minimum of 0.25 processor is recommended for the Primary partition. Secondary partitions must have at least 0.1 processors.
 - v. **Batch CPW:** It will be determined automatically based on the number of processors.
 - vi. **Memory (MB):** Enter amount of memory required by the partition in Mega Bytes.
 - vii. **Int% or Int CPW:** Enter the percentage of the processor CPW or amount of CPW that will be used for Interactive work.

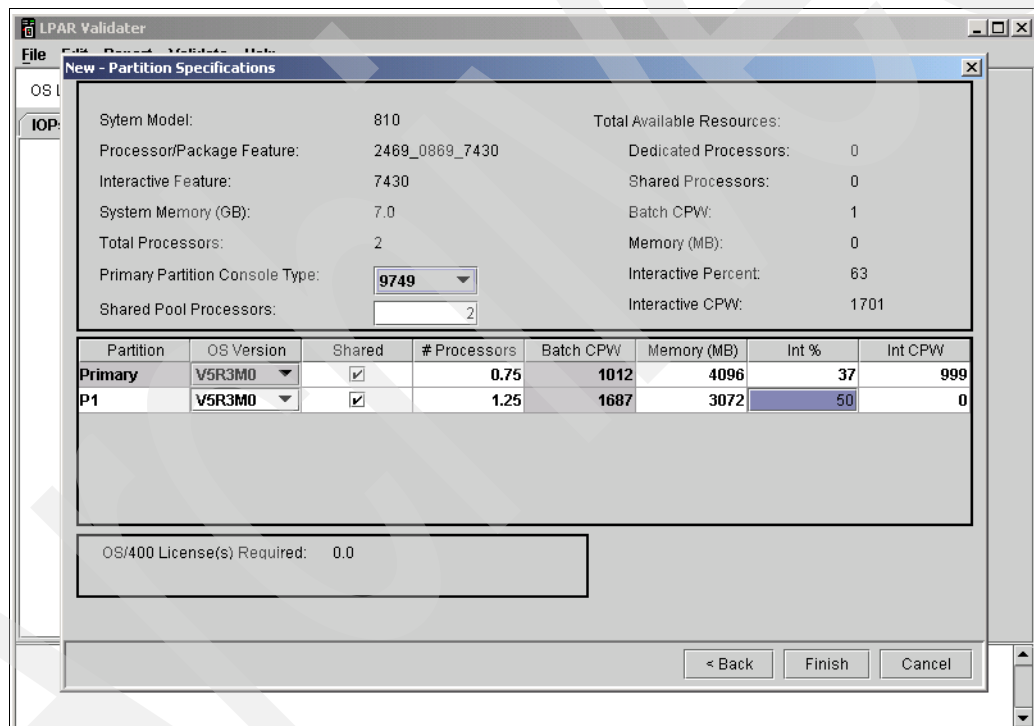


Figure 9-5 Example window Partition Specifications

- d. Select applicable features (IOPs, IOAs, Disk, Towers, etc.). To add a feature first select the **IOP**, **IOA**, **Drives** or **Linux** tabs. Then select the feature and click **Add** next to the desired slot or location (Figure 9-6).

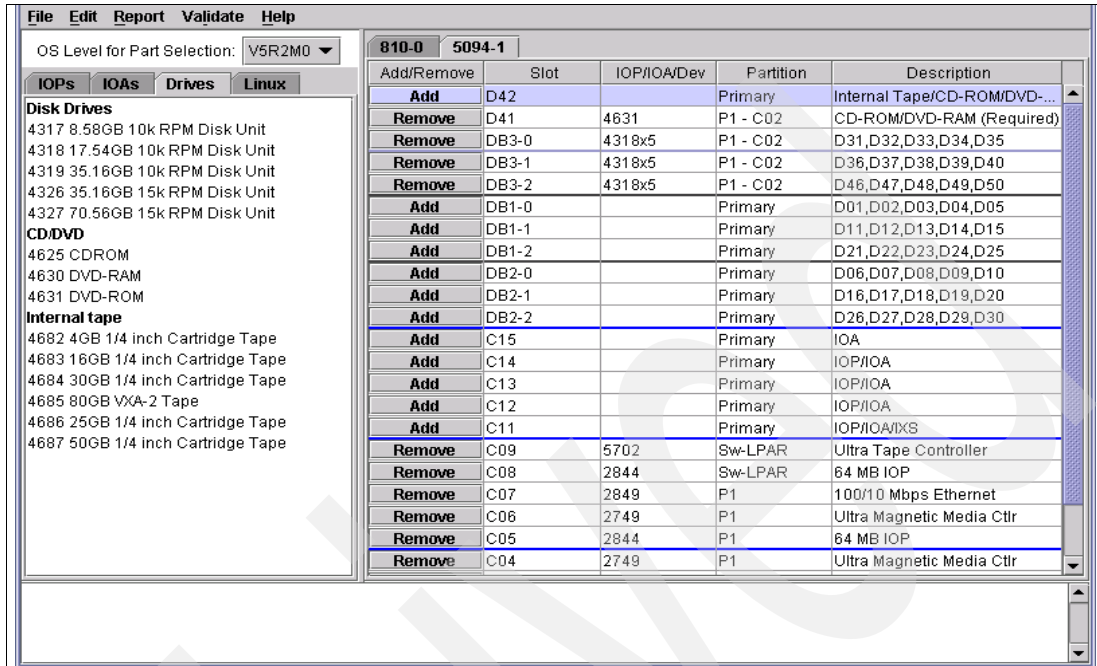


Figure 9-6 Example window features and towers

- e. Validate the LPAR configuration by selecting **Validate**. Correct any errors as they appear in the message window in red (Figure 9-7).

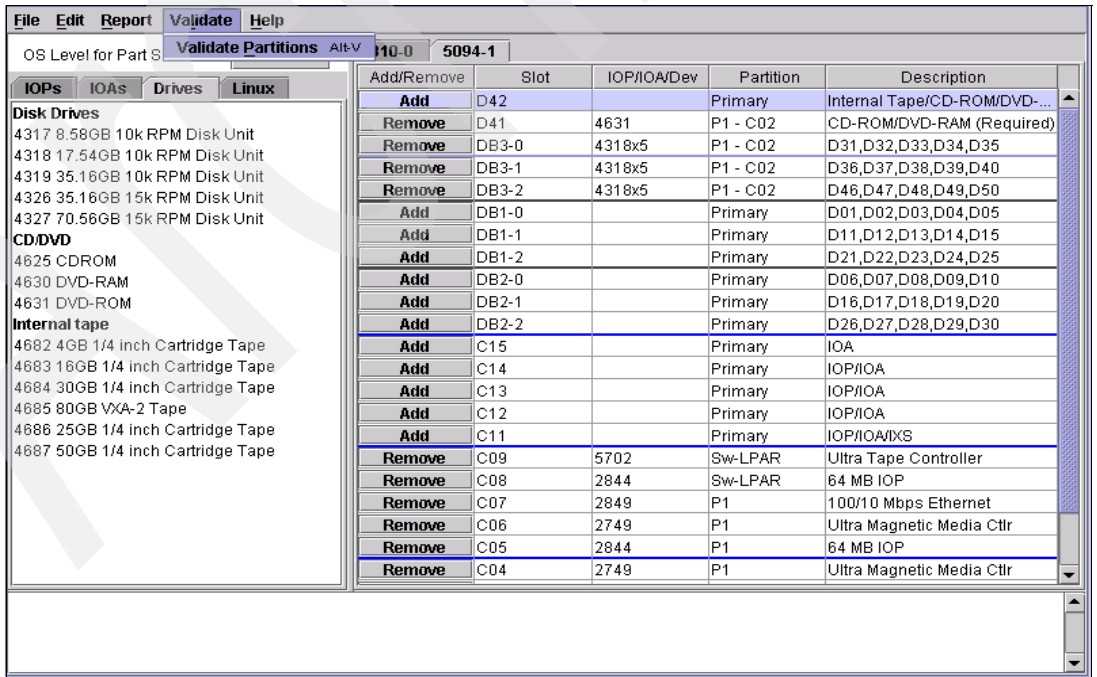


Figure 9-7 Example window Partition Validation

- f. The resulting report can be viewed or saved to disk by selecting Report. For more complete information, save both the **Detail** and **Summary** versions of the report using the **All** option (Figure 9-8).

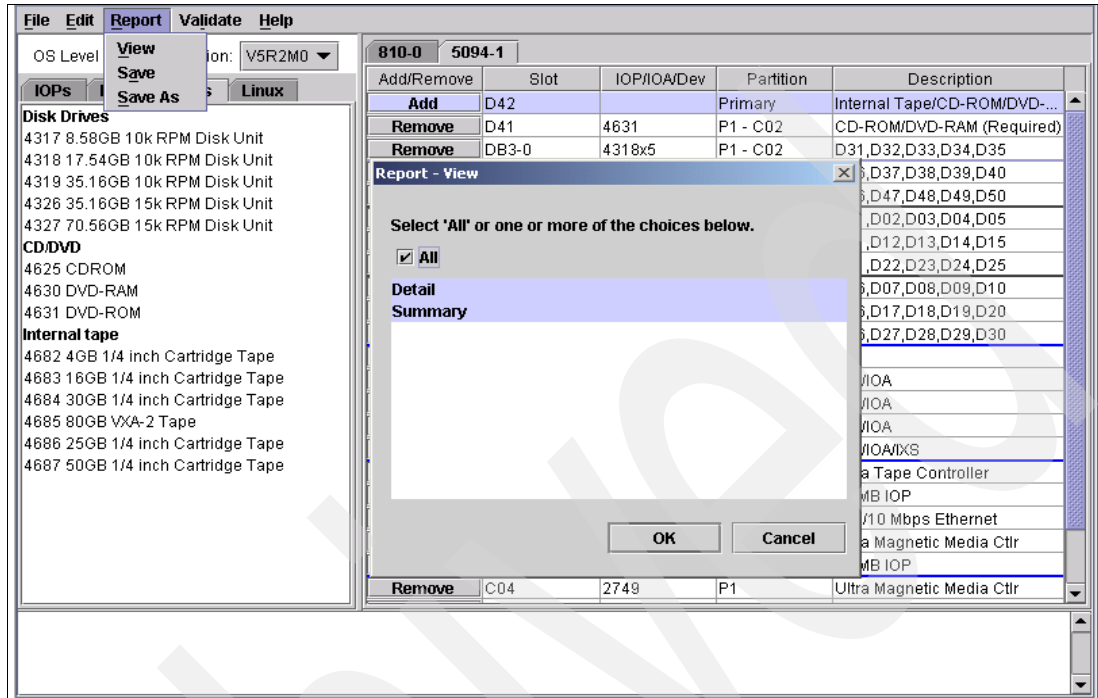


Figure 9-8 Example window view or save report

- g. The following images contain relevant portions of the **Detail** reports for the existing 810 and 825 in these scenarios (see Figure 9-9 here through Figure 9-16 on page 267).

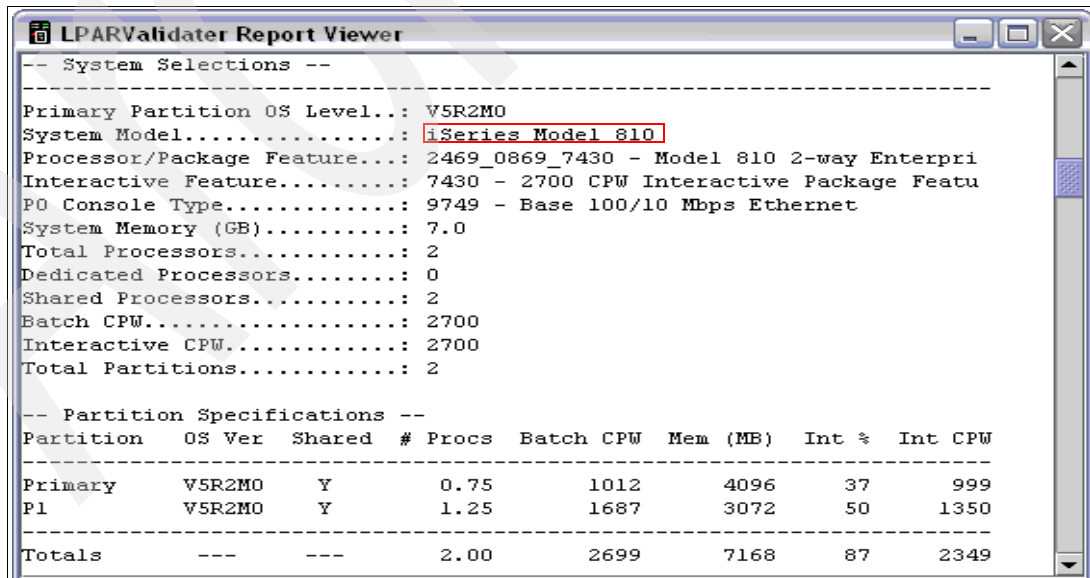


Figure 9-9 LVT report — current system first scenario page 1

LPARValidator Report Viewer

```
-- Tower Properties --
810-0 -- iSeries Model 810
```

Slot	Partition	Device	Description	
DB3	Primary -	4318x5	17.54GB 10k RPM Disk Unit	
DB3(a)	Primary -	None	Disk Drives	
DB3(b)	Primary -	None	Disk Drives	
DB2	Primary -	4318x5	17.54GB 10k RPM Disk Unit	
DB2(a)	Primary -	None	Disk Drives	
DB2(b)	Primary -	None	Disk Drives	
DB1	Primary -	4318x5	17.54GB 10k RPM Disk Unit	
DB1(a)	Primary -	None	Disk Drives	
DB1(b)	Primary -	None	Disk Drives	
D08	Primary -	None	Internal Tape/CD-ROM/DVD-RAM	
D07	Primary -	4530	DVD-RAM	

EMB	Primary	286x	Embedded IOP	B
C07	Primary	9793	Base 2-Line WAN w/Integrated Modem	U
C06	Primary	9749	Base 100/10 Mbps Ethernet	S
C05	Primary	4746	Twinaxial Workstation	#
C04	Primary	2844	64 MB IOP	1
C03	Primary	5701	1 Gbps Ethernet UTP	
C02	Primary	5702	Ultra Tape Controller	
C01	Primary	2757	RAID Disk Unit Ctrlr	

Figure 9-10 LVT report — current system first scenario page 2

LPARValidator Report Viewer

```
5094-1 -- Expansion Tower
```

Slot	Partition	Device	Description	

D42	Primary	None	Internal Tape/CD-ROM/DVD-RAM	
D41	P1 - C02	4631	CD-ROM/DVD-RAM (Required)	
DB3-0	P1 - C02	4318x5	D31,D32,D33,D34,D35	

DB3-1	P1 - C02	4318x5	D36,D37,D38,D39,D40	
DB3-2	P1 - C02	4318x5	D46,D47,D48,D49,D50	
DB1-0	Primary	None	D01,D02,D03,D04,D05	
DB1-1	Primary	None	D11,D12,D13,D14,D15	
DB1-2	Primary	None	D21,D22,D23,D24,D25	
DB2-0	Primary	None	D06,D07,D08,D09,D10	
DB2-1	Primary	None	D16,D17,D18,D19,D20	
DB2-2	Primary	None	D26,D27,D28,D29,D30	

C15	Primary	None	IOA	B
C14	Primary	None	IOP/IOA	U
C13	Primary	None	IOP/IOA	S
C12	Primary	None	IOP/IOA	#
C11	Primary	None	IOP/IOA/IXS	

C09	Sw-LPAR	5702	Ultra Tape Controller	B
C08	Sw-LPAR	2844	64 MB IOP	U
C07	P1	2849	100/10 Mbps Ethernet	S
C06	P1	2749	Ultra Magnetic Media Ctrlr	#
C05	P1	2844	64 MB IOP	

C04	P1	2749	Ultra Magnetic Media Ctrlr	B
C03	P1	4746	Twinaxial Workstation	U
C02	P1	4778	RAID Disk Unit Ctrlr	S
C01	P1	9844	Base 64 MB IOP	#

Figure 9-11 LVT report — current system first scenario page 3

```

LPARValidator Report Viewer
-- System Selections --
-----
Primary Partition OS Level..: V5R3M0
System Model.....: iSeries Model 825
Processor/Package Feature...: 2473_0873_7418 - Model 825 3/6-way Enterp
Interactive Feature.....: 7418 - 6600 CPW Interactive Package Featu
PO Console Type.....: 4746 - Twinaxial Workstation
System Memory (GB).....: 10.0
Total Processors.....: 6
Dedicated Processors.....: 3
Shared Processors.....: 3
Batch CPW.....: 6600
Interactive CPW.....: 6600
Total Partitions.....: 3

-- Partition Specifications --
-----
Partition  OS Ver  Shared  # Procs  Batch CPW  Mem (MB)  Int %  Int CPW
-----
Primary    V5R3M0    N        3.00    3300      6144      33      2178
P1         V5R3M0    Y        1.50    1650      2048      25      1650
P2         V5R3M0    Y        1.50    1650      2048      25      1650
-----
Totals     ---      ---      6.00    6600      10240     83      5478

```

Figure 9-12 LVT report — current system second scenario page 1

```

LPARValidator Report Viewer
-- Tower Properties --
825-0 -- iSeries Model 825
-----
Slot      Partition  Device      Description
-----
DB4       Primary -   4326x5      35.16GB 15k RPM Disk Unit
DB3       Primary -   4326x5      35.16GB 15k RPM Disk Unit
DB2       Primary -   4326x5      35.16GB 15k RPM Disk Unit
D18      Primary -   None        Internal Tape/CD-ROM/DVD-RAM
D17      Primary -   4631        DVD-ROM
-----
C12      Sw-LPAR    2844        64 MB IOP                                     B
C11      Sw-LPAR    5704        Fiber Channel Tape Ctrlr                     U
C10      Sw-LPAR    None        IOA                                           S
#
2
-----
C07      Primary    9844        Base 64 MB IOP                               B
EMB      Primary    288E        100/10 Mbps Ethernet                         U
C06      Primary    9793        Base 2-Line WAN w/Integrated Modem           S
C05      Primary    2757        RAID Disk Unit Ctrlr                         #
1
-----
C04      Primary    2844        64 MB IOP                                     B
C03      Primary    4746        Twinaxial Workstation                         U
C02      Primary    None        IOP/IOA                                       S
C01      Primary    None        IOA                                           #
3

```

Figure 9-13 LVT report — current system second scenario page 2

Slot	Partition	Device	Description	
5074-1 -- Expansion Tower				
DB1	P1 - C06	4318x10	17.54GB 10k RPM Disk Unit	
DB1(a)	P1 - C06	None	Disk Drives	
DB1(b)	P1 - C06	None	Disk Drives	
DB2	P1 - C12	4318x10	17.54GB 10k RPM Disk Unit	
DB2(a)	P1 - C12	None	Disk Drives	
DB2(b)	P1 - C12	None	Disk Drives	
DB3	P1 - C02	4318x10	17.54GB 10k RPM Disk Unit	
DB3(a)	P1 - C02	None	Disk Drives	
DB3(b)	P1 - C02	None	Disk Drives	
D42	P1 - C02	None	Internal Tape/CD-ROM/DVD-RAM	
D41	P1 - C02	4425	CDROM	
C15	P1	None	IOA	B
C14	P1	None	IOP/IOA	U
C13	P1	None	IOA	S
C12	P1	4778	RAID Disk Unit Ctlr	#
C11	P1	2843	64 MB IOP	
C10	P1	None	IOA	
C09	P1	4746	Twinaxial Workstation	
C07	P1	4746	Twinaxial Workstation	
C06	P1	4778	RAID Disk Unit Ctlr	
C05	P1	2843	64 MB IOP	
C04	P1	None	IOA	B
C03	P1	4838	100/10 Mbps Ethernet	U
C02	P1	4778	RAID Disk Unit Ctlr	S
C01	P1	9943	Base 64 MB IOP	#

Figure 9-14 LVT report — current system second scenario page 3

Slot	Partition	Device	Description	
5094-2 -- Expansion Tower				
D42	Primary	None	CD-ROM/DVD-RAM (Required)	
D41	P2 - C02	4630	CD-ROM/DVD-RAM (Required)	
DB3-0	P2 - C02	4326x5	D31,D32,D33,D34,D35	
DB3-1	P2 - C02	4326x5	D36,D37,D38,D39,D40	
DB3-2	P2 - C02	4326x5	D46,D47,D48,D49,D50	
DB1-0	P2 - C08	4326x5	D01,D02,D03,D04,D05	
DB1-1	P2 - C08	4326x5	D11,D12,D13,D14,D15	
DB1-2	P2 - C08	4326x5	D21,D22,D23,D24,D25	
DB2-0	P2	None	D06,D07,D08,D09,D10	
DB2-1	P2	None	D16,D17,D18,D19,D20	
DB2-2	P2	None	D26,D27,D28,D29,D30	
C15	P2	None	IOA	B
C14	P2	None	IOP/IOA	U
C13	P2	2742	2 Line WAN	S
C12	P2	2742	2 Line WAN	#
C11	P2	2844	64 MB IOP	
C09	P2	None	IOA	B
C08	P2	2757	RAID Disk Unit Ctlr	U
C07	P2	2742	2 Line WAN	S
C06	P2	2849	100/10 Mbps Ethernet	#
C05	P2	2844	64 MB IOP	
C04	P2	None	IOA	B
C03	P2	2849	100/10 Mbps Ethernet	U
C02	P2	2757	RAID Disk Unit Ctlr	S
C01	P2	9844	Base 64 MB IOP	#

Figure 9-15 LVT report — current system second scenario page 4

5088-3 -- Expansion Unit				
Slot	Partition	Device	Description	
C15	Primary	None	IOA	B
C14	Primary	None	IOP/IOA	U
C13	Primary	None	IOP/IOA	S
C12	Primary	None	IOP/IOA	#
C11	Primary	9710	Base Integrated xSeries Server (2.0 GHz)	
C09	Primary	5701	1 Gbps Ethernet UTP	B
C08	Primary	2844	64 MB IOP	U
C07	Primary	2849	100/10 Mbps Ethernet	S
C06	Primary	2849	100/10 Mbps Ethernet	#
C05	Primary	2844	64 MB IOP	
C04	Primary	None	IOA	B
C03	Primary	2742	2 Line WAN	U
C02	Primary	2742	2 Line WAN	S
C01	Primary	2844	64 MB IOP	#

Figure 9-16 LVT report — current system second scenario page 5

4. Using these reports, you can identify possible ways to simplify the migration. For example:

- In the first scenario, notice that there are currently (15) FC 4318 disks units in the 810 System Unit which belong to the Primary partition. The 520 System Unit can accommodate a maximum of 8 disk units (4 base and 4 optional). This means that there is no room to move over intact all of the Primary Partition disks. One option considers that there are enough PCI and Disks slots available in the 5094 to relocate all (15) disk units along with its disk and console IOAs from the System Unit to the 520 System i5 system during the migration. This option takes advantage of no longer requiring a Primary partition in System i5 systems System Unit, which allows all its resources to migrate intact (P0 will become P2 after the migration). Figure 9-17 illustrates this scenario.

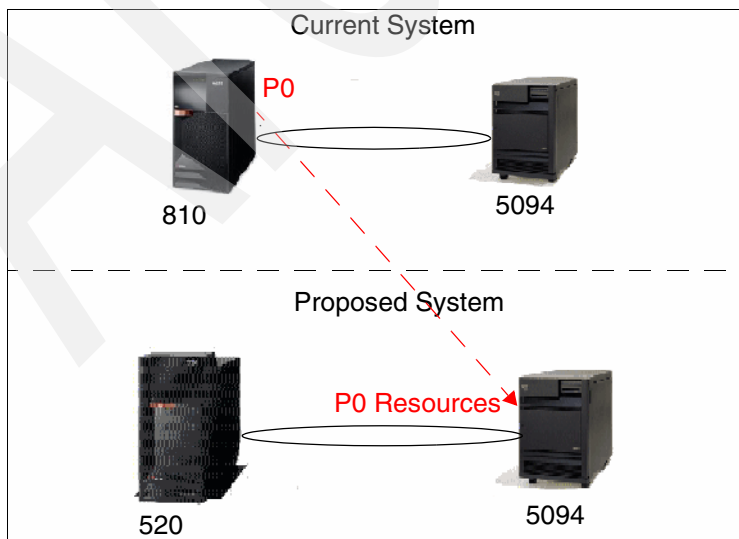


Figure 9-17 P0 migration first scenario

- In the second scenario, there are also enough available disk slots and PCI slots in the 5094 to relocate the (15) of the Primary partition during the migration. However, there are other considerations in this scenario. The 825 has three HSL Loops and the 570 System i5 only has one. It is important to consider the HSL cabling between different technology towers to avoid performance degradation. For instance, the 5074 uses HSL1, whereas the 5094 and 5088 uses more recent HSL2 technology. In the original 825 configuration they were placed in different loops to help with this consideration but now there is only one loop. The solution is to place the 5074 in between the 5094 and 5088, minimizing the effects of any contention caused by the slower 5074 hardware. See Figure 9-18.




Figure 9-18 HSL cabling second scenario

5. Steps 3a through 3f are repeated for the proposed 520 and 570 System i5 systems. This will illustrate the LPAR configuration and resource allocation as well as the configurations for the proposed systems to be ordered. The following images contain relevant portions of the *Detail* reports for both of the proposed systems. See Figure 9-19 on page 269 through Figure 9-26 on page 272.

LPARValidator Report Viewer

```
-- System Selections --
-----
Primary Partition OS Level...: V5R3M0
System Model.....: iSeries Model 520
Processor/Package Feature...: 8955_0905_7457 - Model 9406-520 SF2 2-way
Interactive Feature.....: 7457 - 6000 CPW Interactive Package Featu
PO Console Type.....: 5550 - Hardware Management Console
System Memory (GB).....: 8.0
Total Processors.....: 2
Dedicated Processors.....: 0
Shared Processors.....: 2
Batch CPW.....: 6000
Interactive CPW.....: 6000
Total Partitions.....: 2

-- Partition Specifications --
-----
Partition  OS Ver  Shared  # Procs  Batch CPW  Mem (MB)  Int %  Int CPW
-----
P1          V5R3M0    Y       1.10    3300      3072      55     3300
P2          V5R3M0    Y       0.90    2699      5120      40     2400
-----
Totals      ---      ---      2.00    5999      8192      95     5700
```

Figure 9-19 Detail LVT report — proposed system first scenario page 1

LPARValidator Report Viewer

```
9406_520-0 -- iSeries Model 520
-----
Slot      Partition  Device      Description
-----
P2        P1 - EMB     None        D01,D02,D03,D04
P3        P1 - EMB     None        D05,D06,D07,D08
P4D1     P1 - EMB     None        Internal Tape/CD-ROM/DVD-RAM
P4D2     P1 - EMB     None        DVD-ROM
P4D3     P1 - EMB     2640       IDE DVDROM
-----
EMBE      P1          EETH       Embedded 1 Gbps Ethernet (2 Ports)
-----
-----
EMB       P1          570B      Base DASD Feature
C06       P1          9844      Base 64 MB IOP
C03       P1          9793      Base 2-Line WAN w/Integrated Modem
C05       P1          None      IOA
-----
-----
C04       P2          4746      Twinaxial Workstation
C02       P2          5701      1 Gbps Ethernet UTP
C01       P2          2844      64 MB IOP
C01       P2          2844
-----
-----
```

Figure 9-20 Detail LVT report — proposed system first scenario page 2

LPARValidator Report Viewer

5094-1 -- Expansion Tower

Former Primary Disk and Disk/Console IOA

Slot	Partition	Device	Description	
D42	P1	None	CD-ROM/DVD-RAM (Required)	
D41	P1 - C02	4631	CD-ROM/DVD-RAM (Required)	
DB3-0	P1 - C02	4318x5	D31,D32,D33,D34,D35	
DB3-1	P1 - C02	4318x5	D36,D37,D38,D39,D40	
DB3-2	P1 - C02	4318x5	D46,D47,D48,D49,D50	
DB1-0	P2 - C15	4318x5	D01,D02,D03,D04,D05	
DB1-1	P2 - C15	4318x5	D11,D12,D13,D14,D15	
DB1-2	P2 - C15	4318x5	D21,D22,D23,D24,D25	
DB2-0	P1	None	D06,D07,D08,D09,D10	
DB2-1	P1	None	D16,D17,D18,D19,D20	
DB2-2	P1	None	D26,D27,D28,D29,D30	
C15	P2	2757	RAID Disk Unit Ctlr	B
C14	P2	5702	Ultra Tape Controller	U
C13	P2	None	IOP/IOA	S
C12	P2	2849	100/10 Mbps Ethernet	#
C11	P2	2844	64 MB IOP	
C09	P1	None	IOA	B
C08	P1	2849	100/10 Mbps Ethernet	U
C07	P1	4746	Twinaxial Workstation	S
C06	P1	2749	Ultra Magnetic Media Ctlr	#
C05	P1	2844	64 MB IOP	
C04	P1	2749	Ultra Magnetic Media Ctlr	B
C03	P1	4746	Twinaxial Workstation	U
C02	P1	4778	RAID Disk Unit Ctlr	S
C01	P1	9844	Base 64 MB IOP	#

Figure 9-21 Detail LVT report — proposed system first scenario page 3

LPARValidator Report Viewer

-- System Selections --

Primary Partition OS Level...: V5R3M0

System Model.....: iSeries Model 9406_570

Processor/Package Feature...: 8961_0920_7470 - Model 9406-570 2/4way 1.

Interactive Feature.....: 7470 - 11700 CPW Interactive Package Feat

P0 Console Type.....: 5550 - Hardware Management Console

System Memory (GB).....: 10.0

Total Processors.....: 4

Dedicated Processors.....: 2

Shared Processors.....: 2

Batch CPW.....: 11700

Interactive CPW.....: 11700

Total Partitions.....: 3

-- Partition Specifications --

Partition	OS Ver	Shared	#	Procs	Batch CPW	Mem (MB)	Int %	Int CPW
P1	V5R3M0	Y		1.25	3656	2048	31	3656
P2	V5R3M0	Y		0.75	2193	2048	18	2193
P3	V5R3M0	N		2.00	5850	6144	50	5850
Totals	---	---		4.00	11699	10240	99	11699

Figure 9-22 Detail LVT report — proposed system second scenario page 1

LPARValidator Report Viewer

9406_570-0 -- iSeries Model 9406_570

Slot	Partition	Device	Description	
P3	P3 - EMB	None	D01, D02, D03, D04, D05, D06	
P4	P3 - EMB	2640	IDE DVDROM	
C06	P1	None	IOA	B
C05	P1	None	IOP/IOA	U
C04	P1	None	IOA	S
C03	P1	None	IOP/IOA	#
				2
EMBE	P1	EETH	Embedded 1 Gbps Ethernet (2 Ports)	B
				U
				S
				#
				1
EMB	P3	5709	32MB Planar RAID Feature Card (DaughterBCard-YAWL)	
C02	P3	9793	Base 2-Line WAN w/Integrated Modem	U
C01	P3	9844	Base 64 MB IOP	S
C01	P3	9844		#

Figure 9-23 Detail LVT report — proposed system second scenario page 2

LPARValidator Report Viewer

5074-1 -- Expansion Tower

Slot	Partition	Device	Description	
DB1	P1 - C06	4318x10	17.54GB 10k RPM Disk Unit	
DB1(a)	P1 - C06	None	Disk Drives	
DB1(b)	P1 - C06	None	Disk Drives	
DB2	P1 - C12	4318x10	17.54GB 10k RPM Disk Unit	
DB2(a)	P1 - C12	None	Disk Drives	
DB2(b)	P1 - C12	None	Disk Drives	
DB3	P1 - C02	4318x10	17.54GB 10k RPM Disk Unit	
DB3(a)	P1 - C02	None	Disk Drives	
DB3(b)	P1 - C02	None	Disk Drives	
D42	P1 - C02	None	Internal Tape/CD-ROM/DVD-RAM	
D41	P1 - C02	4425	CDROM	
C15	P1	None	IOA	B
C14	P1	None	IOP/IOA	U
C13	P1	None	IOA	S
C12	P1	4778	RAID Disk Unit Ctlr	#
C11	P1	2843	64 MB IOP	
C10	P1	None	IOA	
C09	P1	4746	Twinaxial Workstation	
C07	P1	4746	Twinaxial Workstation	
C06	P1	4778	RAID Disk Unit Ctlr	
C05	P1	2843	64 MB IOP	
C04	P1	None	IOA	B
C03	P1	4838	100/10 Mbps Ethernet	U
C02	P1	4778	RAID Disk Unit Ctlr	S
C01	P1	9943	Base 64 MB IOP	#

Figure 9-24 Detail LVT report — proposed system second scenario page 3

LPARValidator Report Viewer

5094-2 -- Expansion Tower

Former Primary Disk and Disk/Console IOA

Slot	Partition	Device	Description	
D42	P1	None	CD-ROM/DVD-RAM (Required)	
D41	P1	None	CD-ROM/DVD-RAM (Required)	
DB3-0	P2 - C02	4326x5	D31,D32,D33,D34,D35	
DB3-1	P2 - C02	4326x5	D36,D37,D38,D39,D40	
DB3-2	P2 - C02	4326x5	D46,D47,D48,D49,D50	
DB1-0	P2 - C08	4326x5	D01,D02,D03,D04,D05	
DB1-1	P2 - C08	4326x5	D11,D12,D13,D14,D15	
DB1-2	P2 - C08	4326x5	D21,D22,D23,D24,D25	
DB2-0	P3 - C14	4326x5	D06,D07,D08,D09,D10	
DB2-1	P3 - C14	4326x5	D16,D17,D18,D19,D20	
DB2-2	P3 - C14	4326x5	D26,D27,D28,D29,D30	
C15	P3	4746	Twinaxial Workstation	B
C14	P3	2757	RAID Disk Unit Ctlr	U
C13	P3	2844	64 MB IOP	S
C12	P2	2742	2 Line WAN	#
C11	P2	2844	64 MB IOP	
C09	P2	2742	2 Line WAN	B
C08	P2	2757	RAID Disk Unit Ctlr	U
C07	P2	2742	2 Line WAN	S
C06	P2	2849	100/10 Mbps Ethernet	#
C05	P2	2844	64 MB IOP	
C04	P2	None	IOA	B
C03	P2	2849	100/10 Mbps Ethernet	U
C02	P2	2757	RAID Disk Unit Ctlr	S

Figure 9-25 Detail LVT report — proposed system second scenario page 4

LPARValidator Report Viewer

5088-3 -- Expansion Unit

Slot	Partition	Device	Description	
C15	Sw-LPAR	None	IOA	B
C14	Sw-LPAR	5704	Fiber Channel Tape Ctlr	U
C13	Sw-LPAR	2844	64 MB IOP	S
C12	P3	None	IOP/IOA	#
C11	P3	4710	Integrated xSeries Server (2.0 GHz)	
C09	P3	5701	1 Gbps Ethernet UTP	B
C08	P3	2844	64 MB IOP	U
C07	P3	2849	100/10 Mbps Ethernet	S
C06	P3	2849	100/10 Mbps Ethernet	#
C05	P3	2844	64 MB IOP	
C04	P3	None	IOA	B
C03	P3	2742	2 Line WAN	U
C02	P3	2742	2 Line WAN	S
C01	P3	2844	64 MB IOP	#

Figure 9-26 Detail LVT report — proposed system second scenario page 5

9.4.2 Post-sales customer tasks for both scenarios

At this point the current system information has been gathered and reviewed. This information was also used to plan the migration, and the necessary hardware has been purchased. Here are the remaining post-sales tasks to be completed by the customer:

1. Load V5R3 on all the logical partitions of the current system.
2. Do a complete system save of the current system.

3. Export the LPAR configuration. Perform the following steps:
 - a. Start an iSeries Navigator session and select the system that is partitioned (Figure 9-27).

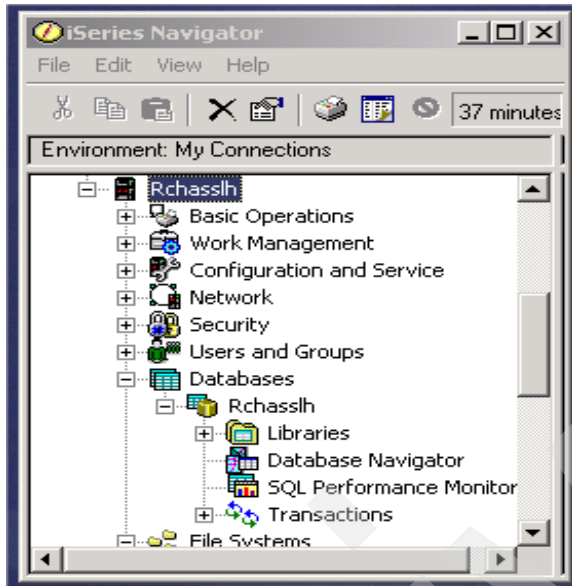


Figure 9-27 iSeries Navigator panel

- b. Select **Configuration and Service** (Figure 9-28).

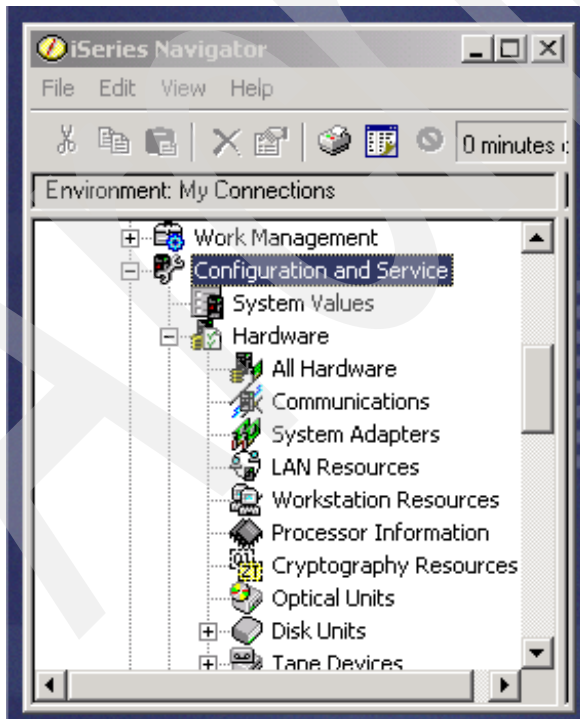


Figure 9-28 iSeries Navigator panel

- c. Select **Logical Partitions**, then right-click and select **Configure Partition** (Figure 9-29).

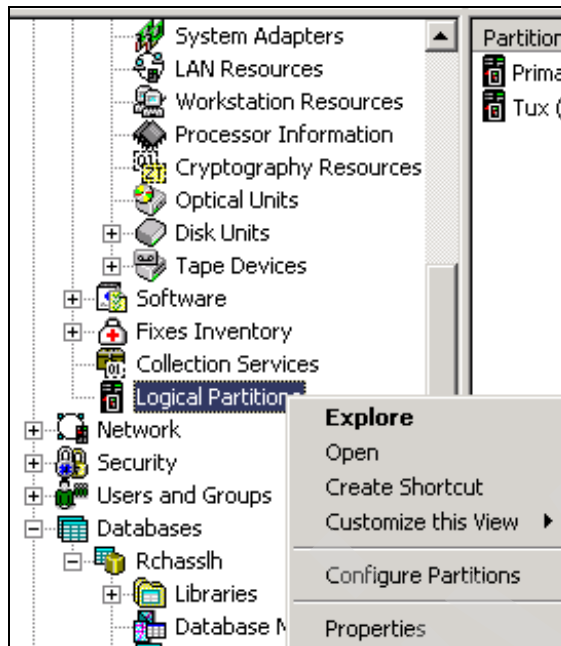


Figure 9-29 iSeries Navigator panel

- d. A list of partition configurations is displayed (Figure 9-30).

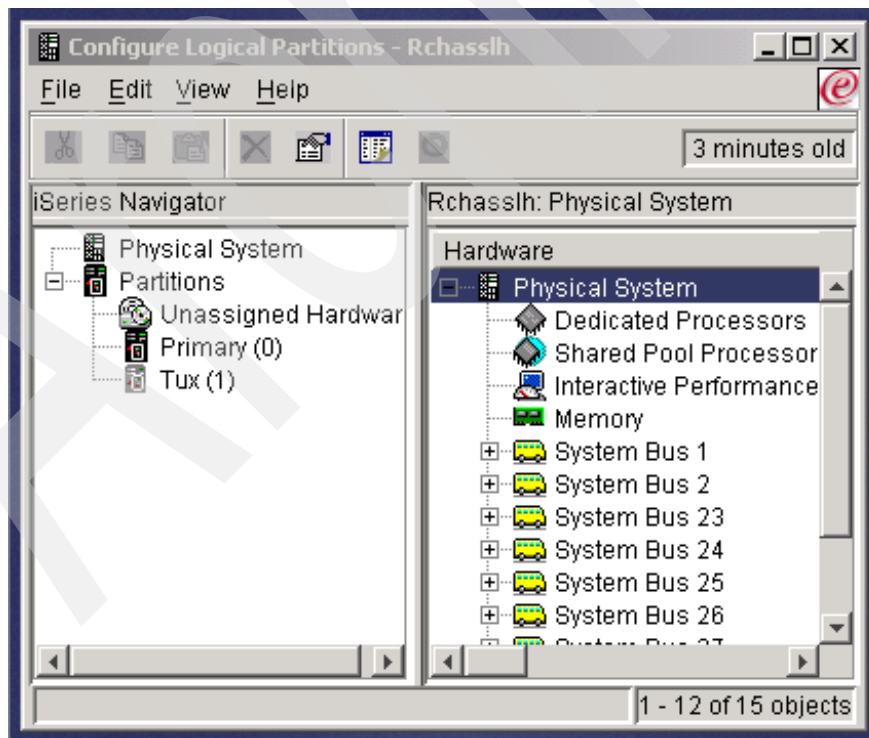


Figure 9-30 iSeries Navigator panel

- e. Right-click **Physical System** and select **Recovery**. Then select **Save All Configuration Data** (Figure 9-31).

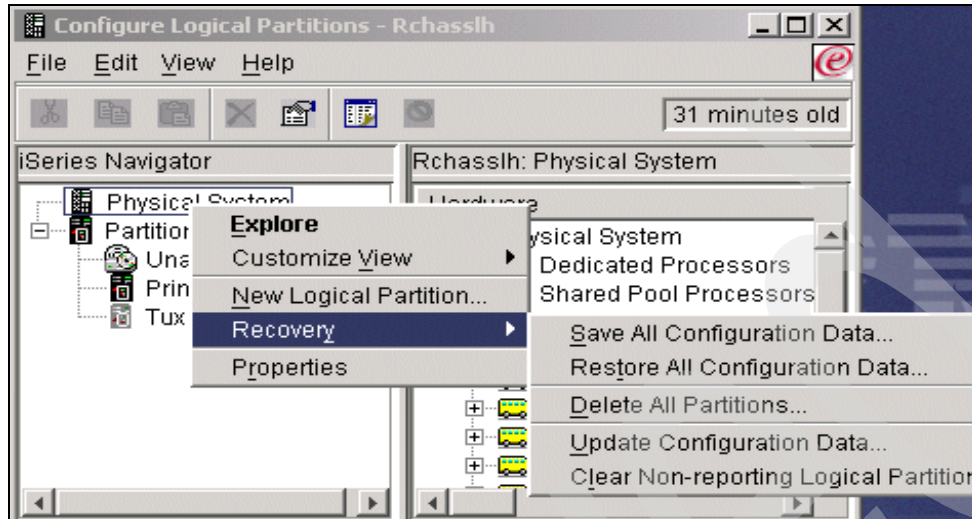


Figure 9-31 iSeries Navigator panel

- f. Enter a PC filename. The file should have been created prior to this step. You can browse for the filename as well. Click **OK**; the file is saved to the media of your choice. We recommend CD or diskette; the HMC uses either media (Figure 9-32).

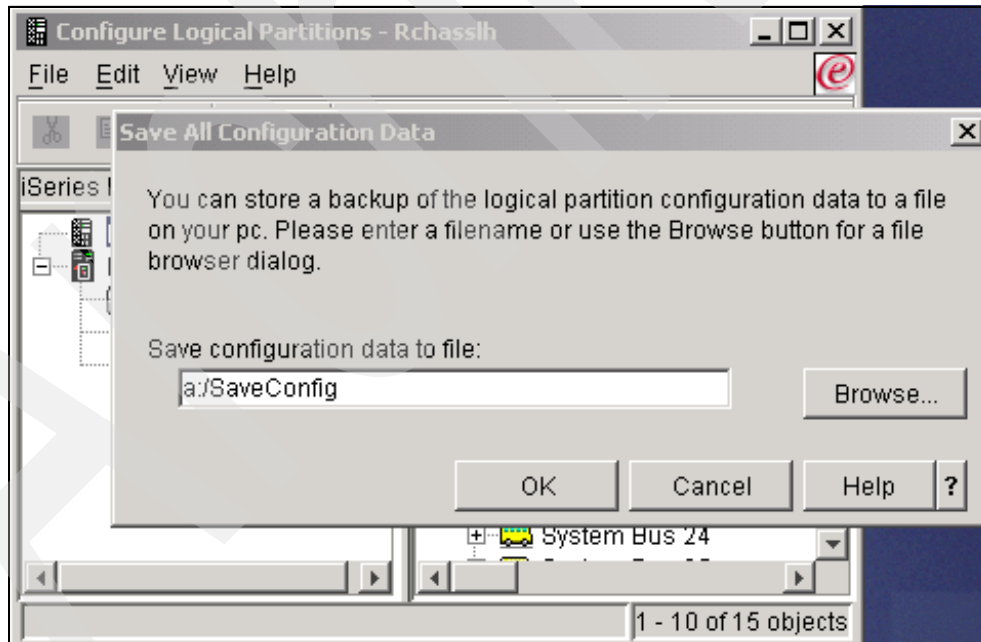


Figure 9-32 iSeries Navigator Save Configuration Data Panel

9.4.3 Post sales tasks: IBM

Note: Depending on the model System i5 system being install, some or all of the HMC setup is done by the customer. The smaller System i5 systems are *Customer Set Up (CSU)*. Refer to 3.1.3, “Initial setup of the HMC” on page 48 for more details.

The IBM Customer Service Representative (CSR) performs the following tasks:

1. Sets up the new hardware and connects the Hardware Manager Console (HMC) to the new server:
 - The setup consists of connecting the HMC to the System i5 system. A UTP CAT 5 cable will be connected from an ethernet port on the back of the HMC to the port labeled HMC1 on the back of the System i5 system unit (a crossover cable is not necessary). See the example in Figure 9-33.

Important: Do not connect the System i5 system to utility power until the HMC is attached. Refer to the HMC Overview Chapter for details.




Figure 9-33 Example of HMC attached to 520 System i5 system

- The IBM CSR would initially only attach the existing towers to the System i5 system unit via HSL. Any additional towers would be added later. This is to avoid any bus re-numbering. The server is then powered on and IPL to the Hypervisor level. At this time, the System i5 system unit has been assigned new bus numbers by the *bus master* firmware. The Hypervisor *bus master* tracks and assigns all bus numbers. The existing towers retain their bus numbers.
2. The server is then powered down and any new towers are attached by the CSR. The server is once again to the Hypervisor level, and any new towers are assigned new bus numbers. The System i5 system unit and original towers retain their bus numbers. See the example in Figure 9-34.
 3. The CSR configures and test the ECS connection and turns the system over to the customer.




Figure 9-34 Example of HMC attached to 520 System i5 system now with ECS connection

9.4.4 Customer migration tasks

The customer must perform the following tasks:

1. Complete the setup of the HMC:
 - The new HMC will include a setup document and will come pre-loaded with the *Hardware Information Center* which contains additional information on many topics related to the HMC. Review 3.1.3, “Initial setup of the HMC” on page 48 for an overview of the HMC setup.

Note: To print topics from the *Hardware Information Center*, consider accessing HMC via the Web-based System Management remote client, which allows you to print from a PC browser. Review 6.6.1, “Web SM” on page 208 for details on Web-based Systems Management Remote Client.

2. The customer imports/migrates their LPAR configurations to the HMC:
 - During the migration of the LPAR configurations, partition P0/Primary is reassigned the next available partition number, Pn+1. The Primary partition, as you know it today, no longer exists. The HMC is now used to manage the partitions.
 - To start the import process, a Linux command is executed via the command line from an `rshe11` session at the HMC.
3. Perform the following steps to migrate the configuration data to the HMC:
 - a. Minimize your HMC window. You should have a blank screen.

- b. Click the right mouse key on the blank screen. The box in Figure 9-35 is displayed.




Figure 9-35 HMC window

- c. Select **Terminals** and then **rshterm**. A term session is started (Figure 9-37).



Figure 9-36 HMC window

- d. From the xterminal session (Figure 9-37) you can enter the commands to start the migration process.




Figure 9-37 xTerminal session window

- e. Load your diskette or CD that contains your configuration data into your drive.
f. Now enter the following command: `migrcfg -t 1 -m [system-name] -f [filename]` (see Figure 9-38).

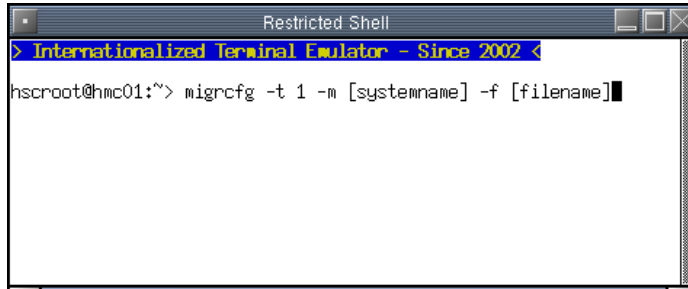


Figure 9-38 xTerminal session window

4. After the LPAR configurations are migrated (Figure 9-39), correct any resource reallocations resulting from P0/Primary being reassigned. Allocate new hardware resources as required. Validate the new Pn+1 partition (former primary) against the configuration and resource allocation documentation gathered in early steps. Refer to Chapter 5, “Partition creation using the HMC” on page 139, for detailed instructions for allocating resources and creating new partitions.

Important: Before you can activate the former Primary partition in both scenarios, it will be necessary identify the new locations of all resources that were in the original System Unit in the *Tagged I/O* tab of the HMC P2 Profile (Load source disk, IOA with other disk, and console IOA, etc.). For more information, refer to the HMC Profile and Partition Profile creation in Chapter 5, “Partition creation using the HMC” on page 139.




Figure 9-39 520 with migrated partitions

5. After all allocations have been made and partition profiles verified, the partitions can be activated. In both of the scenarios listed earlier, all partitions moved over intact.

9.5 Backing up HMC profiles

As with any new install, migration, or upgrade, getting a backup of the new information is critical. In the following sections, we describe how to obtain the backup.

9.5.1 Backing up Critical Console Data

Using your HMC you can back up the following data:

- ▶ User-preference files
- ▶ User information
- ▶ HMC platform-configuration files
- ▶ HMC log files

The *Back up Critical Console Data* function saves the HMC data stored on the HMC hard disk to the DVD-RAM and is critical to support HMC operations. Back up the HMC after you have made changes to the HMC or to the information associated with partitions.

To back up the HMC, you must be a member of one of the following roles:

- ▶ System Administrator
- ▶ Operator
- ▶ Service Representative

To back up the HMC, do the following steps:

1. In the Navigation area, open the Licensed Internal Code Maintenance icon (Figure 9-40).

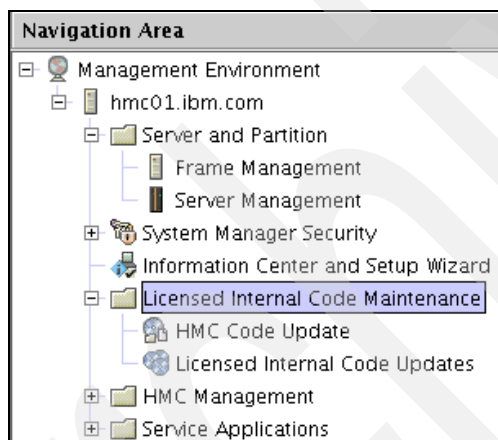


Figure 9-40 Navigation Area Window, Licensed Internal Code Maintenance

2. In the Navigation area, double-click the HMC Code Update icon (Figure 9-41).

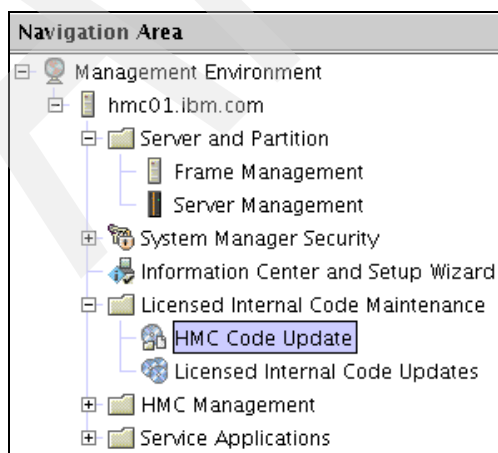


Figure 9-41 Navigation Area window, HMC Code Update

3. In the Contents area, select **Back up Critical Console Data** (Figure 9-42).

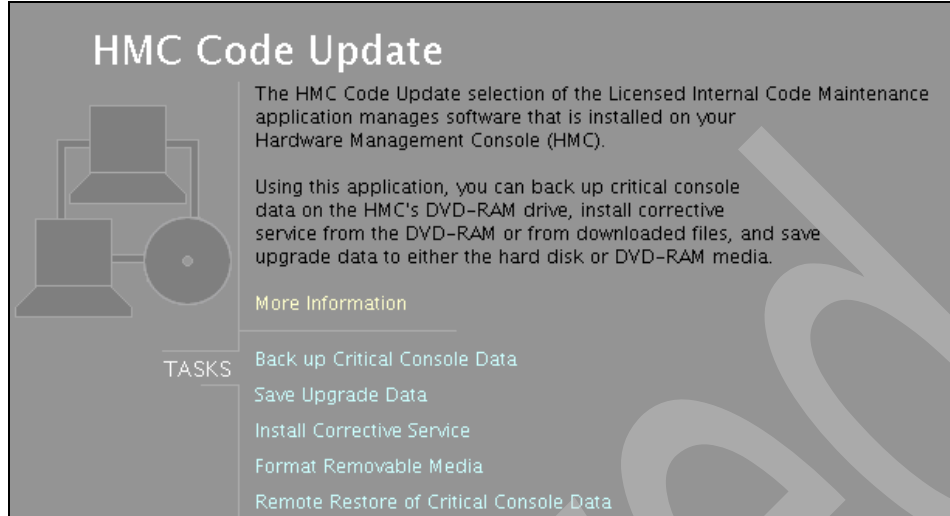


Figure 9-42 Back up Critical Console Data window

4. Insert a formatted DVD-RAM media into the drive.
5. Select **Backup to DVD in Local System** to save your critical console data on the HMC DVD-RAM and click **Next** (Figure 9-43).

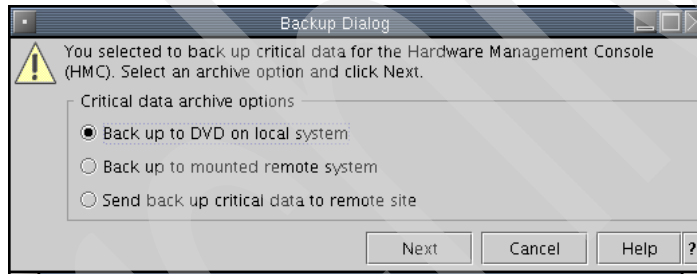


Figure 9-43 Backup dialog window

Note: This backup could take a significant amount of time to complete, perhaps 1.5 hours if it were over 1 GB in size.

9.5.2 Scheduling and reviewing scheduled HMC backups

You can schedule a backup to DVD to occur once, or you can set up a repeated schedule. You must provide the time and date that you want the operation to occur. If the operation is scheduled to repeat, you must select how you want this backup to repeat (daily, weekly, or monthly).

Note: Only the most-recent backup image is stored at any time on the DVD.

To schedule a backup operation, do the following steps:

1. In the Navigation area, open the **HMC Management** folder.
2. In the Navigation area, double-click the **HMC Configuration** folder.
3. In the Contents area, click **Schedule Operations**.
4. From the list, select the HMC you want to back up and click **OK**.
5. Select **Options > New**.
6. In the Add a Scheduled Operation window, select **Backup Profile Data** and click **OK**.
7. In the appropriate fields, enter the time and date that you want this backup to occur.
8. If you want this scheduled operation to repeat, click the **Repeat** tab and enter the intervals at which you want the backup to repeat. You can schedule backup operations to repeat at monthly, weekly, daily, or hourly intervals.
9. In the Set Up Scheduled Operations window, click the **Options** tab and name this file.
10. When you are finished setting the backup time and date, click **Save**.

After you have defined a scheduled backup operation, a description of the operation displays in the Scheduled Operations window.



Using the Advanced Systems Management Interface

This chapter describes the setup and use of the Advanced Systems Management Interface (ASMI). The Advanced Systems Management Interface provides a terminal interface via a standard Web browser to the service processor that allows you to perform general and administrator level service tasks. An example of these tasks would include reading service processor error logs, reading vital product data, setting up the service processor, and controlling the system power.

The ASMI may also be referred to as the service processor menus.

10.1 ASMI introduction

All System i5 systems would use an ASMI to communicate with the service processor. The ASMI function provides much of the same function that had been provided in OS/400 DST/SST in all previous releases before i5OS V5R3.

Most System i5 systems would also typically be controlled using the Hardware Management Console (HMC) introduced in Chapter 3. Any server that would be divided into a multi-partitioned environment would require the HMC to create and maintain the LPAR environment. So you would most likely continue to use the HMC for ongoing service processor management tasks.

However, not all System i5 systems are required to have an HMC for daily operation and control. One example would be a system with only a single system image (a non-LPAR'd system). This might be the case with a smaller, less powerful standalone or non-rack mounted model of the System i5 family. In this situation you could control the management of the system solely with the ASMI function.

However, using the HMC as the primary console for controlling the system does not necessarily exclude using the ASMI as a secondary console. The two can be used on the same system concurrently. When used in this manner, the HMC will have control over most system management functions, and the ASMI will control a reduced set of functions.

10.2 ASMI concepts and initial setup

The main functions of the ASMI are to perform service functions and various system management functions. Some of the servicing functions that were previously in the iSeries service tools will now be performed by the ASMI.

10.2.1 Types of ASMI

There are three methods to provide the access to the ASMI:

- ▶ Access via a Web browser
- ▶ Access via an ASCII console
- ▶ Access via the HMC

In this book we have chosen the method of using a Web browser to access ASMI.

10.2.2 Initial setup of terminal interface to ASMI

For details on the setup of the terminal interface to ASMI, please refer to the book, *Accessing the Advanced System Management Interface*. There you will be provided details of the hardware and software requirements for your terminal or personal computer as well as the actual setup of the three access methods:

http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/info/iphby/iphby.pdf

Tip: If you connected your PC or mobile computer directly to the first service processor port HMC1, you can access ASMI with your Web browser by typing <https://192.168.2.147>. If you are connected to service processor port HMC2, you can use <https://192.168.3.147>. These are the factory default settings and could, however, have been changed on your system.

10.3 Initial tour of the interface

When you connect to the server using the correct secure link and IP address, you will see an initial Welcome login panel similar to the one shown in Figure 10-1. You need to sign on using the Administrator profile (or a profile with Administrator authority levels) in order to see and execute most of the functions described in the remainder of this chapter.

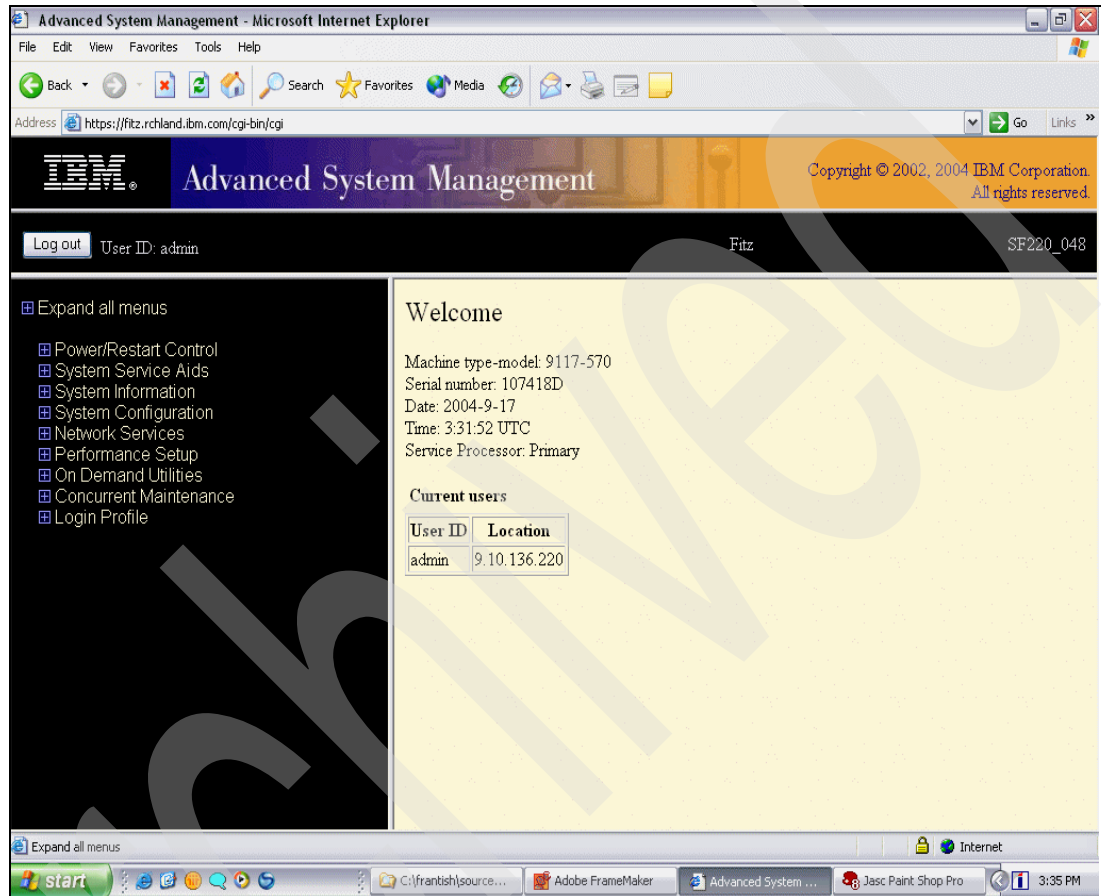


Figure 10-1 ASMI Welcome login panel

Enter a valid User ID and Password, and select the Language you want to use. If this is the first time that ASMI has been accessed, you may be able to use one of the factory default User IDs and Passwords shown in Table 10-1.

Table 10-1 Default

User ID	Password
admin	admin
general	general

Note: You should sign on initially with a User ID having administrative authority. It is also recommended that you change the default passwords to protect the security of your system. The procedure for changing of passwords is described later, near the end of this chapter.

When you first log in to the server, you will see the panel shown in Figure 10-2. You can choose to expand one or more of the service menus, or you can choose to expand all service menus to begin.

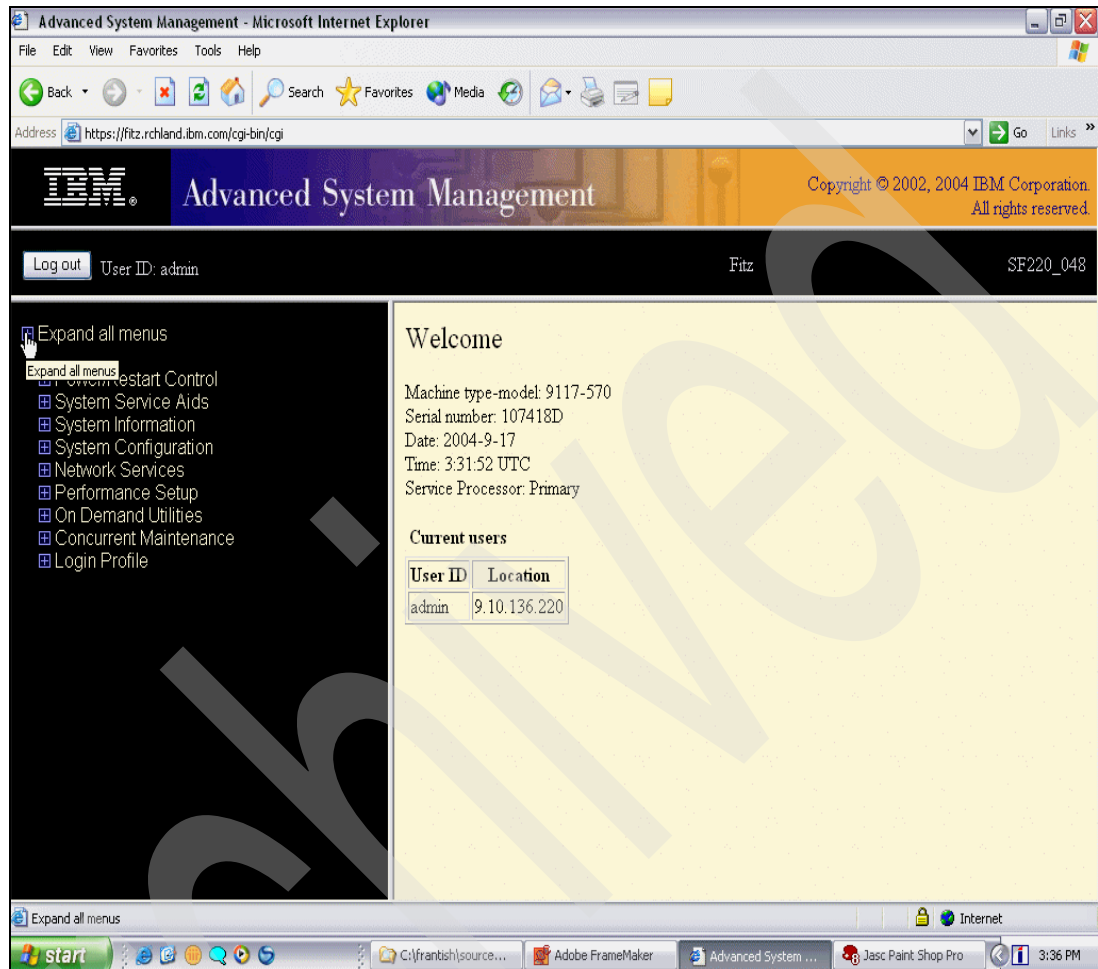


Figure 10-2 First ASMI panel after login

More than one user can be signed on to ASMI at the same time. You will see one or more *Current users* on the panel for each User ID that is signed on. You will also see the IP address that the user is working from.

If you are using ASMI and your session remains inactive for any period of time, you will lose your session. This is a normal security measure.

In the remainder of this chapter we will discuss each of the functions that can be performed using ASMI, which is sometimes also called the service processor menu.

10.3.1 Power/restart control

Figure 10-3 shows the expanded Power/Restart Control menu. Using this menu, you can:

- ▶ Power the system on or off.
- ▶ Set the function to allow an auto-power restart of the system if the system has experienced a power interruption.
- ▶ Perform an immediate power off.
- ▶ Perform a system reboot.
- ▶ Set the function to allow the system to be powered on remotely through a local area network (LAN) connection.



Figure 10-3 Power/Restart Control menu

Power on / power off system

You can shut down the system and change one or more of the power restart/IPL options, or you can restart the system. In earlier releases of OS/400 this was done using the physical control panel of the system or the DST/SST service functions. With ASMI, you can do this using the options shown in Figure 10-4.

If the system is currently in a powered on state, you will see an option to save settings and power off. If the system is currently powered off, you will instead see an option to power on the system.

Power On/Off System

Current system power state: On
 Current firmware boot side: Temporary
 Current system server firmware state: Running

System boot speed: Fast ?

Firmware boot side for the next boot: Temporary ?

System operating mode: Manual ?

Boot to system server firmware: Standby ?

System power off policy: Automatic ?

Save settings ?

Save settings and power off ?

Figure 10-4 Powering off the system

Firmware boot side for the next boot

Select the side from which the firmware will boot: permanent or temporary. Typically, firmware updates should be tested on the temporary side before being applied to the permanent side. This selection is analogous to the previous OS/400 concept of starting the system using micro-code from the “A” side (with permanent PTFs applied only) or using micro-code from the “B” side (using both temporary and permanent PTs applied).

System operating mode

Select the system operating mode: manual or normal. Manual mode is intended to be used when the system hardware is being serviced. When the system is in manual mode, various automatic power-on functions such as auto-power restart and timed power-on are disabled. This selection is analogous to the previous OS/400 concept of starting the system from the control panel or DST/SST in manual mode or normal mode.

AIX/Linux partition mode boot

Select the boot type for an AIX/Linux partition. This option is enabled only when the system is not HMC-managed. The choices are:

1. Continue to the operating system (the partition boots to the operating system without stopping).
2. Boot to the SMS menu (the partition stops at the System Management Services (SMS) menu).
3. Perform service mode boot from a saved list (the system boots from the saved service mode boot list).
4. Perform service mode boot from a default list (the system boots from the default boot list).
5. Boot to open firmware prompt (the system stops at the open firmware prompt).

The “service mode boot from default list” option is the preferred way to run stand-alone diagnostics from a CD-ROM drive. The “service mode boot from saved list” option is typically a shorthand method of executing the partition's diagnostics. To make this option useful, the partition operating system must support diagnostic boot, and diagnostics must be loaded on the partition's disk drive.

Boot to system server firmware

Select the state for the system server firmware: standby or running.

System power off policy

Select the system power off policy. The system power off policy flag is a system parameter that controls the system's behavior 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:

- 1) Power off. (When the last partition is shut down, the system will power down)
- 2) Stay on. (When the last partition powers off, the system will stay powered up.)
- 3) Automatic. (This is the default setting. When the system is not partitioned, the behavior is the same as "power off". When the system is partitioned, the behavior of the system is the same as "stay on".)

i5/OS partition boot mode

Select the i5/OS partition mode for next system boot. This option is available only when the system is not HMC-managed.

Save settings

Click this button to save the selected options. The power state will not change.

Save settings and power on/off

Click this button to save the selected options and change the system power state. The power-on option is available only when the system is powered off, and vice versa.

Save settings and continue system server firmware boot

Click this button to save the selected options and change the system server firmware state. This option is available only when the system server firmware is in standby mode.

Figure 10-5 shows an example of the confirmation you will receive when you successfully power off the system.

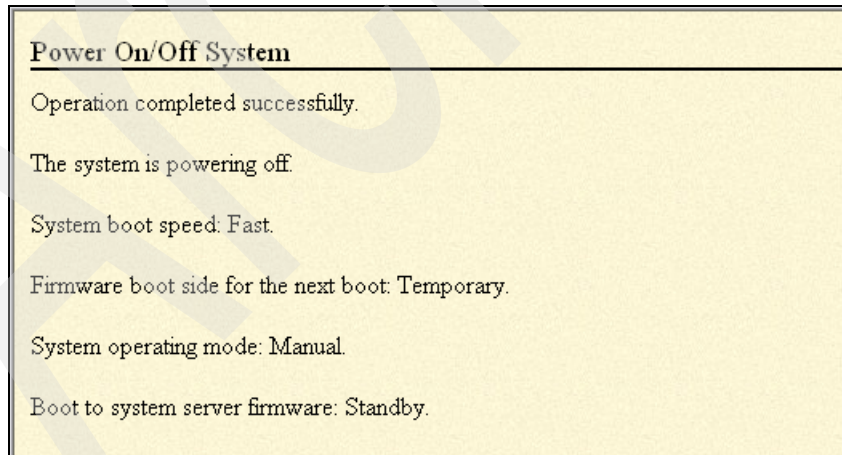


Figure 10-5 Example: Power off confirm

Figure 10-6 shows an example of the confirmation you will receive when you successfully power on the system.

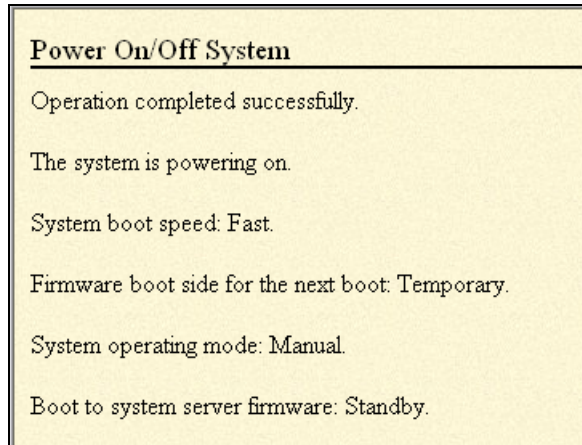


Figure 10-6 Example: Power on confirm

Auto power restart

You can set your system to automatically restart. This function is useful when power has been restored after an unexpected power line disturbance had caused the system to be shut down unexpectedly. Select either *Enable* or *Disable* from the example shown in Figure 10-7). 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 has been resolved.



Figure 10-7 Auto power restart

Immediate power off

You can power off the system quickly using the immediate power off function shown in Figure 10-8. 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 prior to performing an immediate power off.

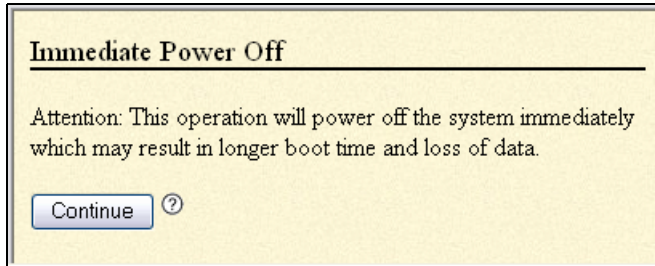


Figure 10-8 Immediate power off

System reboot

You can reboot the system quickly using the reboot function shown in Figure 10-9. The operating system is not notified before the system is rebooted.

Attention: Rebooting the system will immediately shut down all partitions. To avoid experiencing data loss and a longer IPL the next time the system or logical partitions are booted, shut down the operating system prior to performing a reboot.

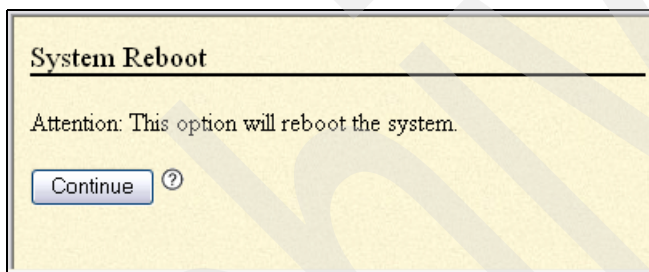


Figure 10-9 System reboot

Wake on LAN

You can power on your system remotely through a local area network (LAN) connection as shown in Figure 10-10. You can choose to Enable or Disable the Wake on LAN® function.

Note: Wake on LAN is supported on ethernet port 0. It is not supported on ethernet port 1.




Figure 10-10 Wake on LAN

10.3.2 System service aids

Figure 10-11 shows the expanded System Service Aids menu. Using this menu, you can:

- ▶ Display system error logs.
- ▶ Set the function to allow a serial port snoop.
- ▶ Initiate a system dump.
- ▶ Initiate a service processor dump.
- ▶ Initiate a Partition Dump
- ▶ Set up a serial port for the call-home and call-in function.
- ▶ Configure the modem connected to the service processor's serial ports.
- ▶ Set up the call-in and call-home policy.
- ▶ Test the call-home function.
- ▶ Reset the service processor.
- ▶ Reset your system to the factory-shipped configuration settings.



Figure 10-11 System Service Aids menu

Error/event logs

From the System Service Aids menu, you can select the option to display the system 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 problems. You will see a selection panel similar to the one shown in Figure 10-12.

Error/Event Logs					
Error logs					
√	Log ID	Time	Failing subsystem	Severity	SRC
<input checked="" type="checkbox"/>	50014C99	2004-09-13 21:47:34	System Hypervisor Firmware	Unrecoverable Error	B70047CC
<input type="checkbox"/>	50012B05	2004-09-01 00:02:50	System Hypervisor Firmware	Unrecoverable Error	B2018105
<input type="checkbox"/>	50010DFD	2004-08-30 22:19:01	System Hypervisor Firmware	Unrecoverable Error	B2018105
<input type="checkbox"/>	500105F8	2004-08-30 02:41:27	System Hypervisor Firmware	Unrecoverable Error	B2018105
<input type="checkbox"/>	500105B7	2004-08-30 02:38:18	System Hypervisor Firmware	Unrecoverable Error	B2018105
<input type="checkbox"/>	5000F765	2004-08-30 00:51:27	System Hypervisor Firmware	Unrecoverable Error	B2018105
<input type="checkbox"/>	5000EFF0	2004-08-29 22:43:15	System Hypervisor Firmware	Unrecoverable Error	B2018105
<input type="checkbox"/>	5000AA54	1970-01-14 00:08:47	System Hypervisor Firmware	Unrecoverable Error	A7004730
<input type="checkbox"/>	5000A817	1970-01-13 23:09:04	System Hypervisor Firmware	Unrecoverable Error	A7004730

Figure 10-12 Error/Event logs

You will first see a listing of all the Error Logs present, then you will see a listing of Informational Logs. From these panels you may select to display the details of one or more events, or to clear all error log entries. To select the detail, you would pre-select the event(s) and click the **Show Details** button shown in Figure 10-13. To clear all event(s), you would click the **Clear all error/event log entries** button, also shown in Figure 10-13.



Figure 10-13 Show detail of Error/Event logs

A panel similar to Figure 10-14 will be displayed for each of the events that you may have selected. You would then be able to use the information when working with your hardware service provider.

Error/Event Logs

Log ID: 50014C99

Platform Event Log

Created at	: 09/13/2004 21:46:32			
Failing Subsystem	: System Hypervisor Firmware			
Event Severity	: Unrecoverable Error			
Event Type	: Not Applicable			
Error Actions Flags	: Report Externally			
	: Service Action Required			
	: Call Home Required			

Primary System Reference Code

Reference Code	: B70047CC			
User Data 2 - 5	: 00000062	: 00060000	: 0000026A	: 00000000
User Data 6 - 9	: 00000000	: 00000000	: 00000000	: 00000000

Symbolic FRU

Priority	: Medium Priority			
Location Code	:			
Part Number	: NEXTLVL			

Log Hex Dump

00000000	50480030	01004346	20040913	21463235	PH.O..CF ...!F25
00000010	20040913	21473402	48000104	00000000	...!G4.H.....

Figure 10-14 Detail of Error/Event logs

Serial port snoop

You can disable or enable serial port snoop on a serial service port. When enabled, data received on the selected serial port is examined, or *snooped*, as it arrives. You can also specify the snoop string, a particular sequence of bytes that resets the service processor if detected. In this manner the serial port S1 serves as a “catch-all” reset device for service personnel.

This function is not available if your system is managed by an HMC. If your system is managed by an HMC, you will see a panel as shown in Figure 10-15.

Serial Port Snoop

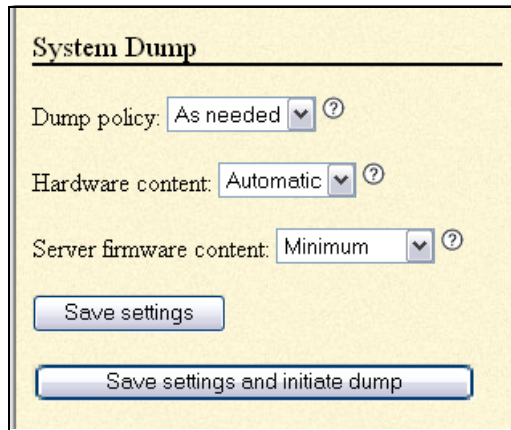
This feature is not available when the system is managed by an HMC.

Figure 10-15 Serial port snoop not available

System dump

Use this procedure only under the direction of your service provider. You can initiate a system dump in order 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 problem. A system dump may also be automatically initiated after a system malfunction, such as a checkstop or hang. It is typically 34 MB.

Figure 10-16 shows the options you can choose relating to a system dump. You can choose to change one or more of the system dump options, or you can choose to change one or more of the options and initiate a system dump.



The screenshot shows a window titled "System Dump" with a yellow background. It contains three dropdown menus: "Dump policy" set to "As needed", "Hardware content" set to "Automatic", and "Server firmware content" set to "Minimum". Each dropdown menu has a question mark icon to its right. Below the dropdowns are two buttons: "Save settings" and "Save settings and initiate dump".

Figure 10-16 Initiating a system dump

Dump policy

Select the policy to determine when system dump data is collected. The **As needed** policy instructs the service processor (SP) to collect system dump data only when the SP determines it is necessary — typically only when a specific failure has not been identified. The **Always** policy instructs the SP to collect the dump data after any hardware or firmware error resulting in a system crash. The default policy is **As needed**.

Hardware content

Select the policy to determine how much hardware data is collected for a system dump. The **Automatic** policy instructs the service processor (SP) to collect the hardware data that it determines is necessary, depending on the particular failure. **Automatic** is the default policy.

Note that collection of hardware data can be time consuming. In some cases, the user may wish to override the default policy. The **Minimum** policy instructs the SP to collect the minimum amount of hardware data. This selection allows the user to minimize the hardware data portion of the system dump.

Note that if this option is selected, the debug data collected for some errors may be insufficient. The **Medium** policy instructs the SP to collect a moderate amount of hardware data. This option strikes a balance between dump collection time and amount of data available for problem debug.

The **Maximum** policy instructs the SP to collect the maximum amount of hardware data. Note that if this selection is chosen, the collection of hardware data can be quite time consuming, especially for systems with a large number of processors.

Service processor dump

Use this procedure only under the direction of your service provider. You can initiate a service processor dump in order to capture and preserve error data after a service processor application failure, external reset, or user request for a service processor dump. This information can be used to resolve a hardware or server firmware problem. The existing service processor dump is considered valid if neither the server firmware nor the HMC has collected the previous failure data. Figure 10-17 shows that you can select the option to enable or disable a system processor dump, or initiate a service processor dump.

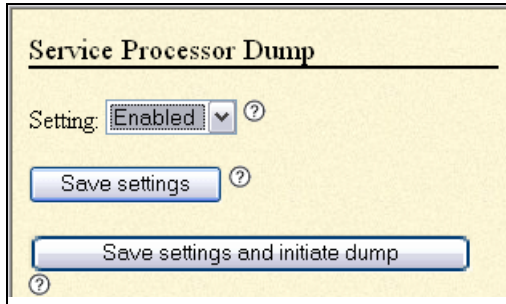


Figure 10-17 Service processor dump

Setting

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

Save settings

Click this button to save the setting for service processor dump.

Save settings and initiate dump

This button is visible only when a service processor (SP) dump is allowed; that is, when SP dumps are enabled and the previous SP dump data has been retrieved. Click this button to initiate an SP dump. Unsaved changes to the enable/disable setting are ignored.

Partition Dump

This feature is available only on non-HMC managed systems, and after the system server firmware is in running state.

Use this procedure only under the direction of your hardware service provider. By initiating a partition dump, you can preserve error data that can be used to diagnose server firmware or operating system problems. The state of the operating system is saved on the hard disk and the partition restarts. This function can be used when the operating system is in an abnormal wait state, or endless loop, and the retry partition dump function is not available. The retry partition dump feature is present only on iSeries systems. See Figure 10-18.

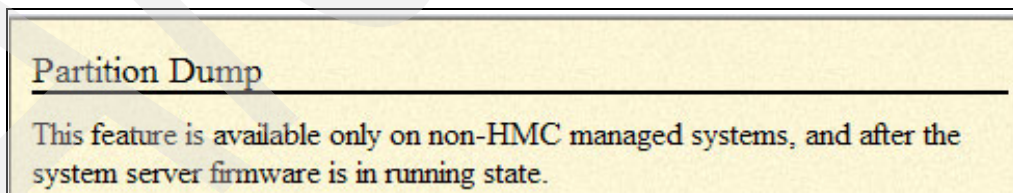


Figure 10-18 Partition Dump feature

Attention: You may experience data loss when using this operation. This feature is only available on non-HMC managed systems that have the system server firmware in the Running state.

To perform this operation, your authority level must be one of the following possibilities:

- ▶ Administrator
- ▶ Authorized service provider

Serial port setup

You can configure the serial ports used with the call-home and call-in features with this option. You can also set the baud rate for the serial ports.

This function is not available if your system is managed by an HMC. If your system is managed by an HMC, you will see a panel as shown in Figure 10-19.




Figure 10-19 Serial port setup not available

Modem configuration

You can configure the modem connected to the serial ports used for the server processor related to the call-home and call-in features with this option.

This function is not available if your system is managed by an HMC. If your system is managed by an HMC, you will see a panel as shown in Figure 10-20.




Figure 10-20 Modem configuration not available

Call-Home/Call-In Setup

You can select which serial port and modem is used, set various telephone numbers, and add customer and contact formation related to the call-home and call-in features with this option.

This function is *not* available if your system is managed by an HMC. If your system is managed by an HMC, you will see a panel as shown in Figure 10-21.



Figure 10-21 Call-Home setup not available

Call-Home Test

You can test the call-home configurations and settings after the modem is installed and set up correctly.

This function is not available if your system is managed by an HMC. If your system is managed by an HMC, you will see a panel as shown in Figure 10-22.




Figure 10-22 Call-home test not available

Reset service processor

Use this procedure only under the direction of your service provider. In critical systems situations, such as during system hangs, you can reset or “reboot” the service processor.

This function is not available if your system is powered on. If your system is powered on, you will see a panel as shown in Figure 10-23.




Figure 10-23 Reset service processor not available

If your system is powered off, you will see a panel as shown in Figure 10-24 asking you to confirm and continue with the reset of the service processor.

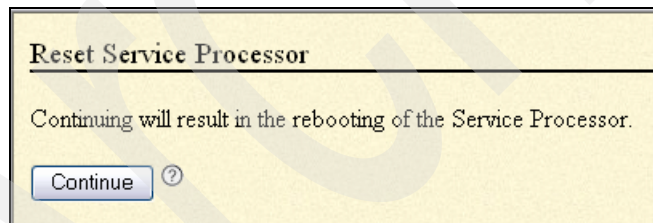


Figure 10-24 Confirm reset of service processor

Clicking **Continue** will cause the service processor to reboot. You will receive a second confirmation panel as shown in Figure 10-25, and the service processor will reboot. As the service processor reboots, your ASMI session will be dropped and you will have to reconnect your session to continue.

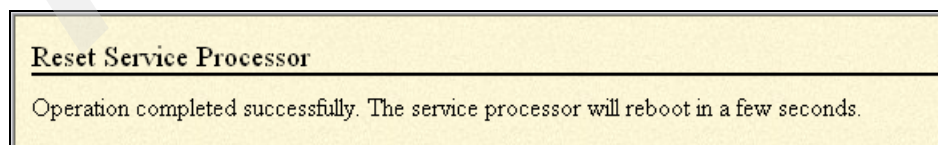


Figure 10-25 Reset service processor success

Factory configuration reset

Use this procedure only under the direction of your service provider. In critical systems situations, you can restore your system to the factory default settings.

Continuing will result 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 once again through the service processor interfaces. Also, you will lose the system error logs and partition-related information.

Attention: Before continuing with this operation, make sure you have manually recorded all settings that need to be preserved.

This function is not available if your system is powered on, if the firmware levels of the permanent and temporary sides are different, or if there is an active HMC session. If any of these conditions apply, you will see a panel as shown in Figure 10-26.

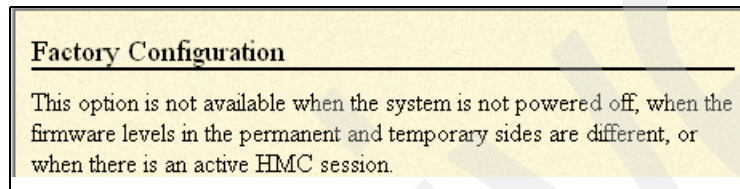


Figure 10-26 Factory configuration reset not available

10.3.3 System information

Figure 10-27 shows the expanded System Information menu. Using this menu, you can:

- ▶ Display vital product data.
- ▶ Perform an 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

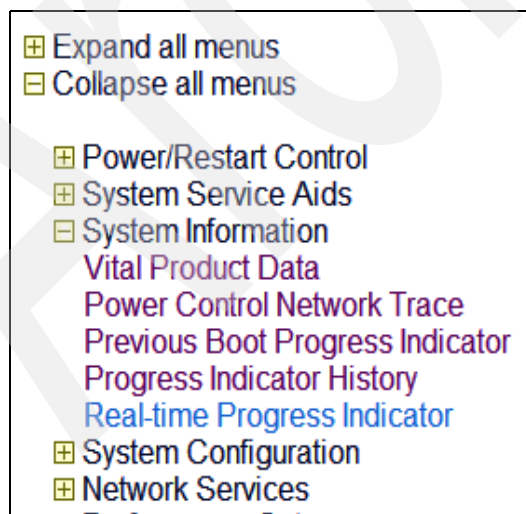


Figure 10-27 System Information menu

Vital product data

From the System Information menu, you can select the option to display the system Vital Product Data. This is the manufacturer's data that defines the system. This data was stored from the system boot prior to the one in progress now. You will see a selection panel similar to the one shown in Figure 10-28.

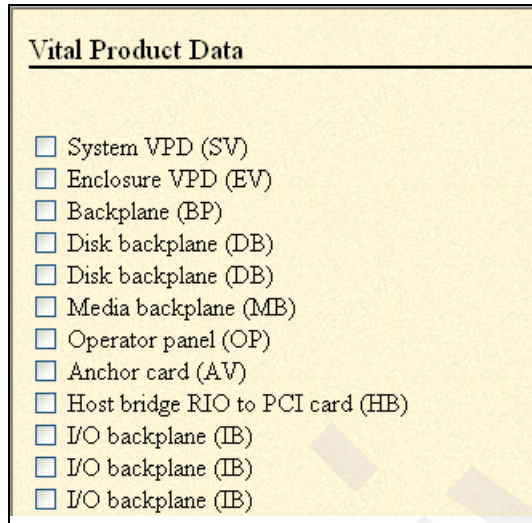


Figure 10-28 Vital product data selection

From these panels, you may select to display the details of one or more hardware components, or to display the details of all the hardware associated with your system. As shown in Figure 10-29, to select specific details, you would pre-select the component(s) and click the **Display details** button. To select all details, you would click the **Display all details** button.




Figure 10-29 Show details of vital product data

A panel similar to Figure 10-30 will be displayed for each of the vital product detail entries that you may have selected. You would then be able to use the information when working with your hardware service provider.

Vital Product Data					
Build name: fips220/b0830a_0434.220					
System brand: P0					
System serial number: 107418D					
Machine type-model: 9117-570					
FRU ID	Part number	Serial number	FRU number	CCIN RID	Location code
SV	80P5110	YL1004000473	80P5110	528E 0x1c00	U9117.570.107418D
EV	97P6644	YL1094180056	97P6644	28DA 0x1e00	U7879.001.DQD04GG
BP	97P6644	YL1094180056	97P6644	28DA 0x800	U7879.001.DQD04GG-P1
DB	80P4812	YL1324178613	80P4812	28DB 0x2400	U7879.001.DQD04GG-P3
DB	80P4812	YL1324178613	80P4812	28DB 0x2401	U7879.001.DQD04GG-P3
MB	80P4556	YL1224120680	80P4556	28DC 0x2500	U7879.001.DQD04GG-P4
OP	97P4940	YL142417715E	97P4940	28D4 0x300	U7879.001.DQD04GG-D1
AV	80P5110	YL1004000473	80P5110	528E 0x500	U7879.001.DQD04GG-P1-C10
HB	97P6644	YL1094180056	97P6644	28DA 0x2100	U7879.001.DQD04GG-P1
IB	97P6644	YL1094180056	97P6644	28DA 0x2000	U7879.001.DQD04GG-P1

Figure 10-30 Details of vital product data

Power control network trace

You can perform a system power control network (SPCN) trace and display the results. This information is gathered to provide additional debug information when working with your hardware service provider.

Note: Producing a trace can take an extended amount of time based upon your system type and configuration. This is a normal delay due to the amount of time the system requires to query the data.

After several minutes, you will see a panel similar to Figure 10-31. Your service provider can make use of this data if requested.

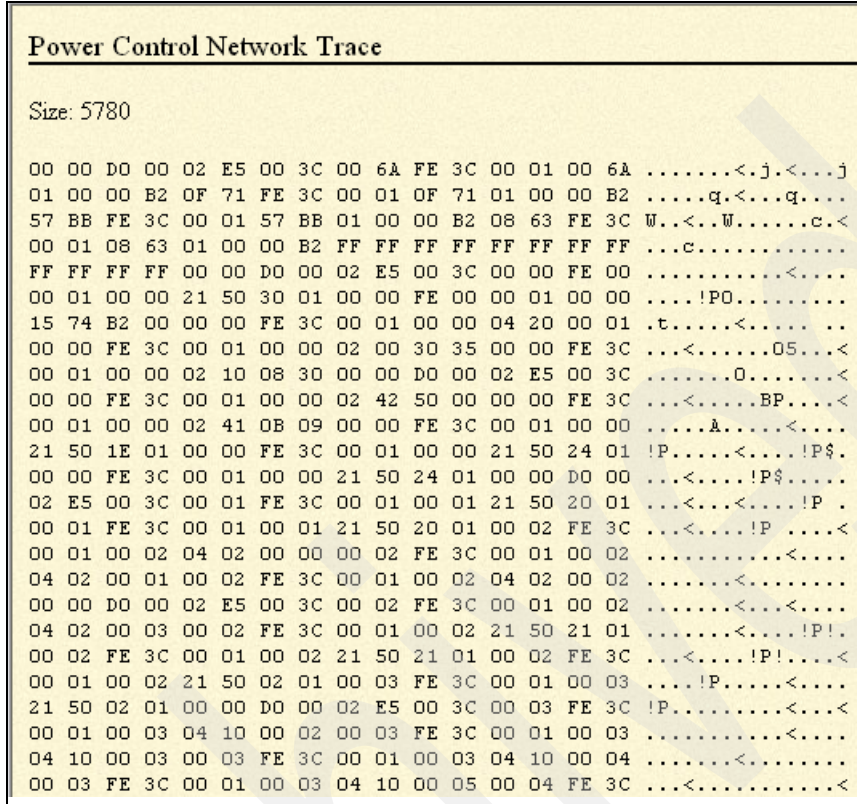


Figure 10-31 Power control network trace

Previous boot progress indicator

You can view the progress indicator that displayed in the control panel during the previous boot, if the previous boot had failed. During a successful boot, the previous progress indicator is cleared. If this option is selected after a successful boot, you will see no indicator. You will see a display similar to Figure 10-32. This information is gathered to provide additional debug information when working with your hardware service provider after a boot has failed.



Figure 10-32 Previous boot progress indicator

Progress indicator history

With this option you can review the progress of codes that appeared in the control panel during the previous boot. The codes display in reverse chronological order. (The first entry seen is the most recent entry.) This information is gathered to provide additional debug information when working with your hardware service provider.

Figure 10-33 shows an example of the progress indicator history selection panel. You can select one or more codes, as directed by your hardware service provider, and click the **Show details button** as shown in Figure 10-34.

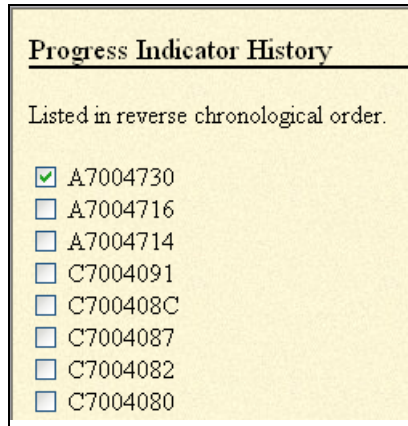


Figure 10-33 Progress indicator history selection




Figure 10-34 Displaying the selected entries

A display similar to Figure 10-35 will be shown for the entries that you had selected. The details can be interpreted by your hardware service provider.

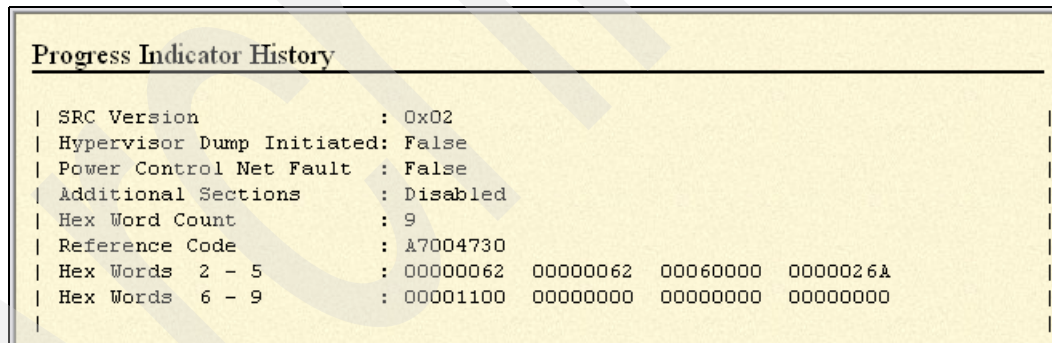


Figure 10-35 Details of progress indicator history

Real-Time Progress indicator

You can view the progress and error codes that currently display on the control panel. Viewing progress and error codes is useful when diagnosing boot-related issues. To perform this operation, your authority level must be one of the following possibilities:

- ▶ General
- ▶ Administrator
- ▶ Authorized service provider

If you have the required authority level, you will see a panel similar to Figure 10-36.

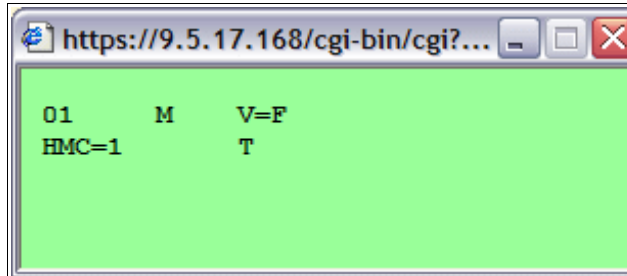


Figure 10-36 Real-Time Progress Indicator

System configuration

Figure 10-37 shows the expanded System Configuration menu. Using this menu, you can:

- ▶ Change the system name.
- ▶ Display the processing unit identifier.
- ▶ 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
- ▶ Change the HSL Opticonnect Connection
- ▶ Enable I/O Adapter Enlarged Capacity
- ▶ View Hardware management Consoles connection



Figure 10-37 System configuration

System name

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 used to identify the system or server. The system name may not be blank and may not be longer than 31 characters. To change the system name, enter a new value and click the store settings button. See Figure 10-38 for an example of the panel you would use to change the system name.

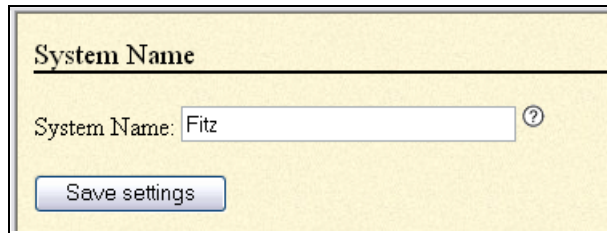


Figure 10-38 System Name

The system is shipped with the default system name initialized to a 31 character value as follows: Server-tttt-mmm-SN0000000:

- ▶ tttt = Machine type
- ▶ mmm = Model number
- ▶ 0000000 = Serial number

You can change the system name to any valid 31 character (or less) string. It does not have to follow the factory set format.

Processing unit identifier

You can change the processing unit identifier, also referred to as the processing unit system power control network (SPCN) ID. The processing unit SPCN ID is used by the SPCN firmware to identify the system type. It is also used to identify the primary service processor if there are two service processors in the system.

This function is not available if your system is powered on. If your system is powered on, you will see a panel as shown in Figure 10-39.

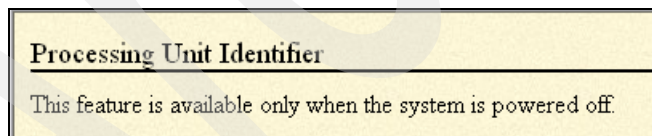


Figure 10-39 Processing unit identifier not available

If your system is powered off, you will see a panel as shown in Figure 10-40 allowing you to enter a new setting value for the processing unit identifier. You can enter a new value and click *Save settings* to reset the value.

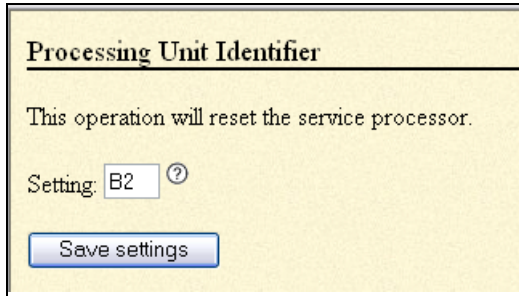


Figure 10-40 Reset the processing unit identifier

Processing unit identifier values

The power control network identifier is intended to uniquely identify each enclosure on the power control network. Typically, these identifiers are automatically assigned by firmware. In some cases, a user may wish to assign specific identifiers to specific drawers. This value is 2 hexadecimal digits. Supported processing unit identifiers are shown in Table 10-2.

Table 10-2 Processing unit identifiers:

Model or expansion unit	Processing unit identifiers
520	B4
550	B5
570	B2
570 (with one or more secondary units)	B3
5074 and 5079	81
5088 and 0588	89
5094 and 5294	8A
0595 and 5095	8B
D10	88
D11	B6
D20	8C

Configure I/O enclosures

This function would normally be used by your hardware service provider. After the server firmware has reached 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.

Refer to Figure 10-41 to see how you can modify the following options for each enclosure that you select.

	Status	Rack address	Unit address	Power Control Network Identifier	Type - Model	Serial number	Location code
<input type="radio"/>	Active	0x3C00	0x1	0xB2	7879-001	DQD04GG	U7879.001.DQD04GG

Identify enclosure ?

Turn off indicator ?

Change settings ?

Clear inactive enclosures ?

Figure 10-41 Configure I/O enclosures

Next we provide a description of these options.

Identify enclosure

Click this button to turn on the indicator on the selected enclosure. You can then visually inspect the enclosure to see the indicator turned on.

Turn off indicator

Click this button to turn off the indicator on the selected enclosure. You can then visually inspect the enclosure to see the indicator turned off.

Change settings

Click this button to change the settings for the selected enclosure. The next page displays options for changing configuration ID, machine type-model, and serial numbers for the enclosure.

► Clear inactive enclosures:

Click this button to clear the rack and unit addresses of all inactive enclosures. You might want to do this if you have removed an enclosure from the system.

► Power control network indicator:

Enter a hexadecimal number for the power control network identifier.

► Type-Model:

Enter the machine type and model in the form TTTT-MMM, where TTTT are the four characters of the enclosure machine type, and MMM are 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

Time of day

You can display and change the system's current date and time. This function is not available if your system is powered on. If your system is powered on, you will see a panel as shown in Figure 10-42.

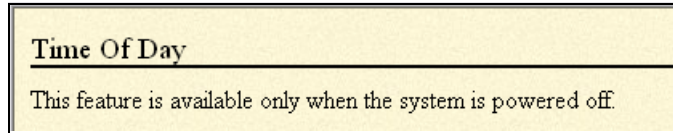


Figure 10-42 Time of day not available

If your system is powered off, you will see a panel as shown in Figure 10-43 allowing you to make changes to the system date or time.

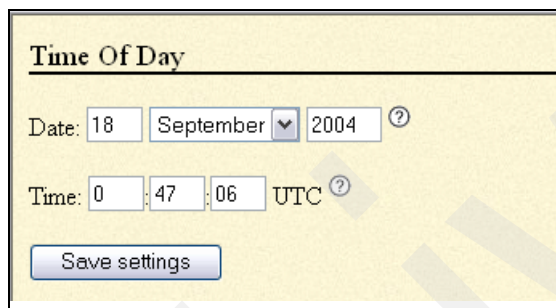


Figure 10-43 Reset date or time of day

Use the following information to help you change the date or time. Click the **Save settings** button to initiate the changes.

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 UTC (Coordinated Universal Time) format. UTC is equivalent to mean solar time at the prime meridian (0 degrees longitude), formerly expressed as Greenwich Mean Time (GMT).

Firmware update policy

Using this option, you can specify the source for firmware updates to your system. Firmware updates can be received from one of two sources:

- ▶ From the Hardware Management Console (HMC):

If an HMC is attached to the system, you can receive firmware updates directly to the HMC via the Web. The HMC can then be used to initiate the firmware updates to the system.

- ▶ From the operating system:

Firmware updates are also available from the operating system (for example, i5/OS) and are delivered by the Program Temporary Fix (PTF) process. So if you have i5/OS installed in one or more partitions of your server you could update the firmware via the operating system. (For the i5/OS operating system, these PTFs are delivered as a new type of PTF with the format MHXXXX. Note that in previous releases of OS/400 the PTF process provided only for MFXXXX and SFXXXX PTFs.)

For example, if you choose the Hardware Management Console (HMC) as the source for a firmware update, the HMC must be used to perform the update. You would use the panel shown in Figure 10-44 to select your source for firmware updates.




Figure 10-44 Selecting the firmware update policy

The default setting of this policy is *not* to allow firmware updates via the operating system. Note that this policy only takes effect when a system is HMC managed. When a system is not HMC managed, firmware updates can only be made 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 powered off. When a system is powered off, firmware updates can be performed from an HMC, regardless of the setting of this policy. However, care should be taken when updating firmware from both an HMC and the operating system. For example, i5/OS load source synchronization may cause firmware updated from an HMC to be overlaid with older firmware from the load source.

PCI error injection policy

The policy for PCI error injection can be set using the display shown in Figure 10-45. You can enable or disable the injection of errors onto the PCI bus. (The PCI bus is one of the hardware busses that the PCI IOP and IOA cards are plugged into. An IOP card is defined as Input Output Processor card and an IOA card is defined as an Input Output Adapter card.)




Figure 10-45 PCI error injection policy

Enabling or disabling the PCI error injection policy

This option controls the PCI error injection policy. If enabled, utilities on the host operating system can inject PCI errors onto the bus. One use of this, for example, is where independent software developers who develop hardware device drivers can inject errors to test the error handling code in the device driver.

We recommend that you set the policy to **Disabled** and only use Enabled under those circumstances where you might want to test a newly developed error handling function.

Interposer Plug Count

We can track the number of times that a multiple chip module (MCM) has been replaced or re-seated on a given interposer. This interposer plug count provides you with information needed to prevent field problems due to damaged or overused interposers. Whenever a service action is performed on a system that requires the replacement or re-seating of an MCM, service personnel are responsible for updating the plug count for that interposer (Figure 10-46).

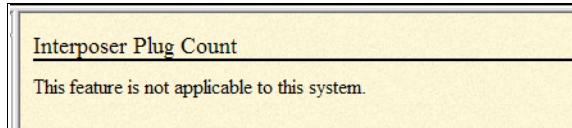


Figure 10-46 Interposer Plug Count

Note: The Interposer Plug Count option is supported only on certain system types and models. If your server does not support this option and you select this option from the menu, the firmware returns a message indicating that this option is not applicable to your system.

HSL Opticonnect Connections

You can view and change the maximum number of Highspeed Link (HSL) Opticonnect Connections allowed for your system (Figure 10-47).

There are two options available to select:

- ▶ Automatic
- ▶ Custom




Figure 10-47 HSL Opticonnect connections

Select **Automatic** to allow the system to select the maximum number of HSL Opticonnect Connections allowed for the system.

Select **Custom** to enter a value for the maximum number of HSL Opticonnect Connections allowed for the system. The value must be between 0 and 4096.

Once the option is selected, click **Save settings** to save the settings. The following screen appears after the save operation (Figure 10-48).

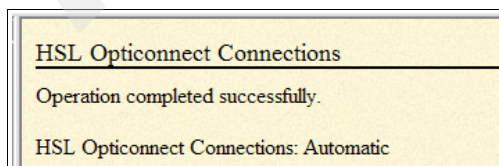


Figure 10-48 HSL Opticonnect Connection save operation

I/O Adapter Enlarged Capacity

You can increase the amount of I/O adapter memory for specified PCI slots.

This option controls the size of PCI memory space allocated to each PCI slots. When enabled, selected PCI slots, including those in external I/O subsystems, receive the larger DMA and memory mapped address space. Some PCI adapters may require this additional DMA or memory space, per the adapter specification.

This option increases system mainstore allocation to these selected PCI slots, and enabling this option may result in some PCI host bridges and slots not being configured because the installed mainstore is insufficient to configure all installed PCI slots (Figure 10-49).




Figure 10-49 I/O Adapter Enlarged Capacity

To change the option, change setting **Disabled** to **Enabled** and click **Save Settings**. Figure 10-50 appears after completion of the save operation.

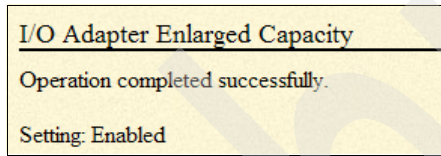


Figure 10-50 I/O Adapter Enlarged capacity saved operation

Hardware Management Consoles

You can see the current HMC connections, as shown in Figure 10-51.

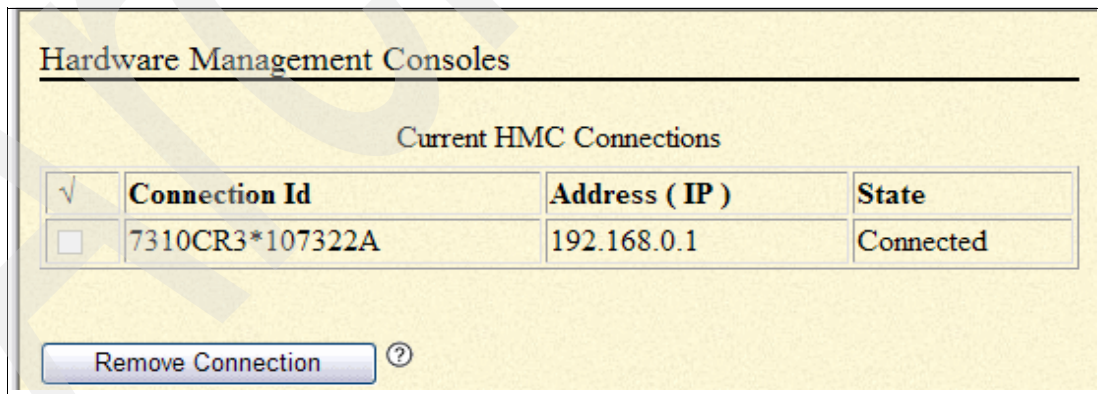


Figure 10-51 HMC Consoles

Hardware deconfiguration policies

You can set various policies to deconfigure processors and memory in certain situations. 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. This can be automated to some degree though the use of *policies*.

You can enable policies that will deconfigure the processor when failures occur, such as a predictive failure (for example, correctable errors generated by a processor exceeding the threshold level for errors), floating point failure, functional failure of the processor, or system bus failure.

Figure 10-52 shows the expanded Hardware Deconfiguration menu, which is a sub-menu under the main System Configuration menu. Using this menu, you can:

- ▶ Set general hardware deconfiguration policies.
- ▶ Set processor deconfiguration policies.
- ▶ Set memory deconfiguration policies.

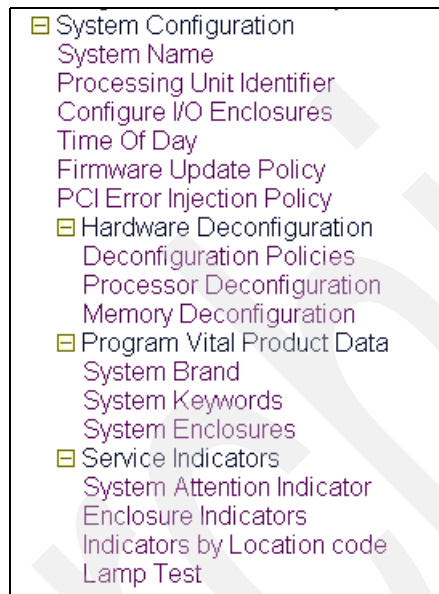


Figure 10-52 Expanded Hardware Deconfiguration menu

General hardware deconfiguration policies

You can use the panel shown in Figure 10-53 to set policies for much of the general system hardware. You would select the policy that you want to enable or disable, and then click the **Save settings** button to make the change.

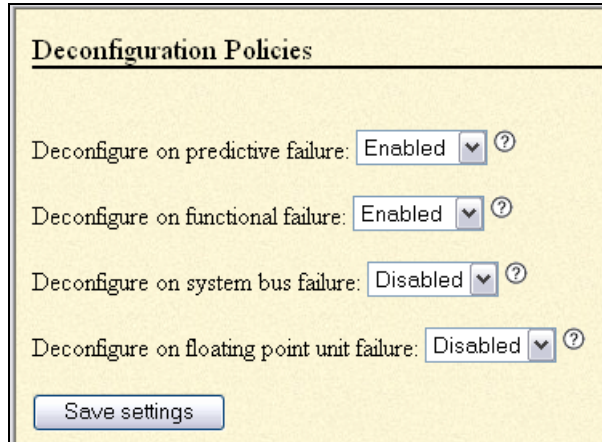


Figure 10-53 General hardware deconfiguration policies

Deconfigure on predictive failure

Select the policy for deconfigure on predictive failures. This 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 affected by the failure will be deconfigured.

Deconfigure on functional failure

Select the policy for deconfigure on functional failures. This applies to run time or persistent boot time deconfiguration of processing unit resources or functions with functional failures, such as checkstop errors or uncorrectable errors. If enabled, the particular resource or function affected by the failure will be 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 affected by the failure will be deconfigured.

Deconfigure on floating point unit failure

Select the policy for deconfigure on floating point unit test failures. If enabled, the periodic floating point unit test is executed, and if it detects an error, then it deconfigures the resource.

Processor deconfiguration

Most System i5 systems will have more than one processor. 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 shown in Figure 10-54 below to start the process of removing processors that might have failed or are beginning to generate serious failure conditions. You can also see processors that may have become deconfigured due to 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 will be 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 run-time 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 run-time 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, you will use a panel similar to the one shown in Figure 10-54. Select the processing unit you want to work with (one or more processing units may be shown) and click the **Continue** button. You should get a panel similar to Figure 10-55.

Processor Deconfiguration

Total system processors: 4

Total system configured processors: 4

Total system deconfigured processors: 0

	Processing unit	Total processors	Configured	Deconfigured
<input type="radio"/>	0	4	4	0

Figure 10-54 Processor deconfiguration, processing unit selection

Processor Deconfiguration

Processing unit: 0

Processor ID	State	Error type	Change settings
0	Configured	0	Configured <input type="button" value="v"/> ?
1	Configured	0	Configured <input type="button" value="v"/> ?
8	Configured	0	Configured <input type="button" value="v"/> ?
9	Configured	0	Configured <input type="button" value="v"/> ?

Figure 10-55 Processor deconfiguration, processor selection

Processor deconfiguration

Refer to Figure 10-55 to see how a processor might be deconfigured. You would select whether each processor should remain configured or become deconfigured, and click the **Save settings** button. State changes take effect on the next platform boot.

Memory deconfiguration

Most System i5 systems will have several megabytes (MB) of memory. Each memory bank contains two DIMMs (dual inline memory module). If the firmware detects a failure, or predictive failure, of a DIMM, it deconfigures the DIMM with the failure, as well as 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* to prevent them from being configured on subsequent boots. Memory banks marked as deconfigured remain offline and will be 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 run-time 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 run-time 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 the Memory Deconfiguration option from the Hardware Configuration sub-menu, you will see a panel similar to the one shown in Figure 10-56 allowing you to view the total memory installed on your system. From this panel you can select the Processing Unit (one or more processing units may be shown) and click **Continue** to advance to the next panel.

A new panel similar to the one shown in Figure 10-57 is then displayed. You can then see any Memory Banks that may have become deconfigured due to some error condition that the system was able to detect and isolate, and the operational state of the memory blocks as either configured or deconfigured.

Memory Deconfiguration

Total system memory: 20480 MB

Total system configured memory: 20480 MB

Total system deconfigured memory: 0 MB

	Processing unit	Total memory	Configured	Deconfigured
<input type="radio"/>	0	20480 MB	20480 MB	0 MB

Figure 10-56 Memory deconfiguration, processing unit selection

Memory Deconfiguration

Processing unit: 0

Memory bank	Size	State	Error type	Change settings
0	1024 MB	Configured	0	Configured <input type="button" value="v"/> ?
1	2048 MB	Configured	0	Configured <input type="button" value="v"/> ?
2	2048 MB	Configured	0	Configured <input type="button" value="v"/> ?
3	1024 MB	Configured	0	Configured <input type="button" value="v"/> ?
4	2048 MB	Configured	0	Configured <input type="button" value="v"/> ?
5	2048 MB	Configured	0	Configured <input type="button" value="v"/> ?
6	1024 MB	Configured	0	Configured <input type="button" value="v"/> ?
7	2048 MB	Configured	0	Configured <input type="button" value="v"/> ?
8	2048 MB	Configured	0	Configured <input type="button" value="v"/> ?
9	1024 MB	Configured	0	Configured <input type="button" value="v"/> ?
10	2048 MB	Configured	0	Configured <input type="button" value="v"/> ?
11	2048 MB	Configured	0	Configured <input type="button" value="v"/> ?

Figure 10-57 Memory deconfiguration, memory bank selection

Memory deconfiguration

Refer to Figure 10-57 to see how a memory bank block might be deconfigured. You would select whether each memory bank should remain configured or become deconfigured, and click the **Save settings** button. State changes take effect on the next platform boot.

Program vital product data

Figure 10-58 shows the expanded Program Vital Product Data menu, which is a sub-menu under the main System Configuration menu. Using this menu, you can:

- ▶ Display system brand.
- ▶ Display system keywords.
- ▶ Display system enclosures.

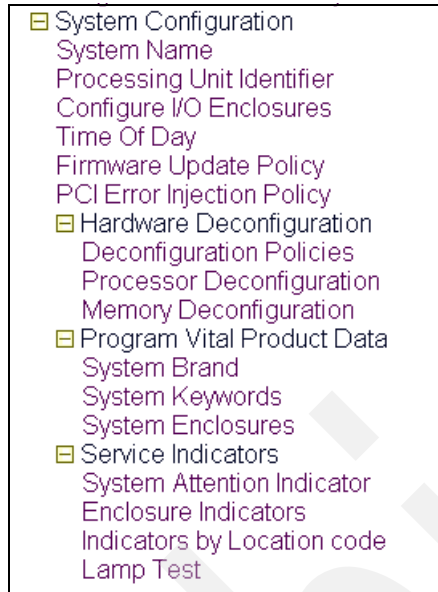


Figure 10-58 Expanded Program Vital Product Data menu

System brand

This menu (Figure 10-59) is available only when the system is powered off.




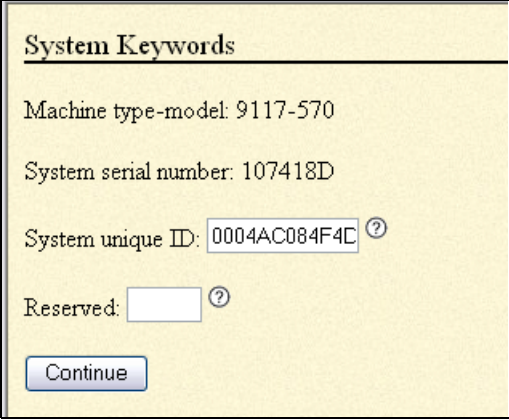
Figure 10-59 System brand

System brand

Enter a 2-character brand type. The first character must be one of the following: D (IBM Storage), I (IBM iSeries), N (OEM iSeries only), O (OEM pSeries only), or P (IBM pSeries), and the second character is reserved. A value of zero means that there is no specific information associated with it. This entry is write once only, except in the case where it is all blanks, or when changing from a pSeries 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

This menu is available only when the system is powered off (Figure 10-60).



The screenshot shows a window titled "System Keywords" with a yellow background. It contains the following text and input fields:

- Machine type-model: 9117-570
- System serial number: 107418D
- System unique ID: 0004AC084F4C (with a question mark icon)
- Reserved: (empty field) (with a question mark icon)
- A "Continue" button at the bottom left.

Figure 10-60 System keywords

System unique ID

Enter a system-unique serial number as 12 hexadecimal digits. The value should be unique to a given system anywhere in the world. A valid value is required for the machine to boot.

Storage facility system 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.

Storage facility system serial number

Enter a system serial number in the form XYYYYY, where XX is the code for the plant of manufacture, and YYYYYY is the unit sequence number. Valid characters are '0' to '9' and 'A' to 'Z'. A valid value is required for the machine to boot. This entry is write once only.

Storage facility machine type-model

Enter a storage facility 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. This entry is write once only.

Storage facility serial number

Enter a storage facility serial number in the form XYYYYY, where XX is the code for the plant of manufacture, and YYYYYY is the unit sequence number. Valid characters plant of manufacture are '0' to '9' and 'A' to 'Z'. The sequence number must be in the form DDDD0, or ADDD0, or AADD0, or AAAD0, or AAAA0, where D is a digit '0' to '9' and A is an alphabetic character 'A' to 'Z', excluding 'E', 'I', 'J', 'O', 'Q', 'S', 'U'. The rightmost character must always be zero. 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 ID

Enter a 2-character system ID. This value is used to identify each system in a storage facility. A valid value is required for the machine to boot. Valid values are '00' and '01'. This entry is write once only.

Storage facility system unique ID

Enter a system-unique serial number as 12 hexadecimal digits. The value should be unique to a given storage facility anywhere in the world. 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.

Storage facility manufacturing ID

Enter a storage facility manufacturing ID in the form JJJYYYY, where JJJ is the Julian date and YYYY is the year. A valid value is required for the machine to boot. This entry is write once only.

Node ID

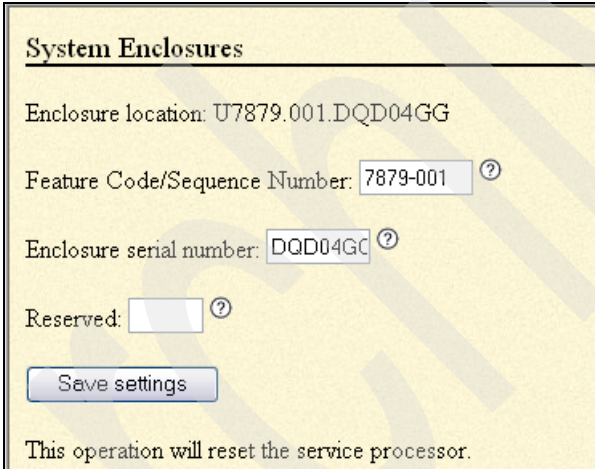
Enter a 16-digit hexadecimal number for the worldwide node name. This value is an IEEE-assigned 64-bit identifier for the storage facility. A valid value is required for the machine to boot. This entry is write once only.

Reserved

Reserved — this should be set to blanks unless directed by Level 4 Service.

System enclosures

This menu is available only when the system is powered off (Figure 10-61).



System Enclosures

Enclosure location: U7879.001.DQD04GG

Feature Code/Sequence Number: 7879-001 ?

Enclosure serial number: DQD04GG ?

Reserved: ?

Save settings

This operation will reset the service processor.

Figure 10-61 System enclosures

Feature code/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 attached to the system. A valid value is required for the machine to boot. When this value is changed, the service processor will reboot so that the location codes can be updated accordingly.

Enclosure serial number

Enter an enclosure serial number in the form XXYYYYY, where XX is the code for the plant of manufacture, and YYYYYY is the unit sequence number. Valid characters are '0' to '9' and 'A' to '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 will reboot so that the location codes can be updated accordingly.

Reserved

Reserved — this should be set to blanks unless directed by Level 4 Service.

Service indicators

Figure 10-62 shows the expanded Service Indicators menu, which is a sub-menu under the main System Configuration menu. Using this menu, you can:

- ▶ Display the system attention indicator
- ▶ Display the enclosure indicators
- ▶ Display indicators by location code
- ▶ Perform a lamp test

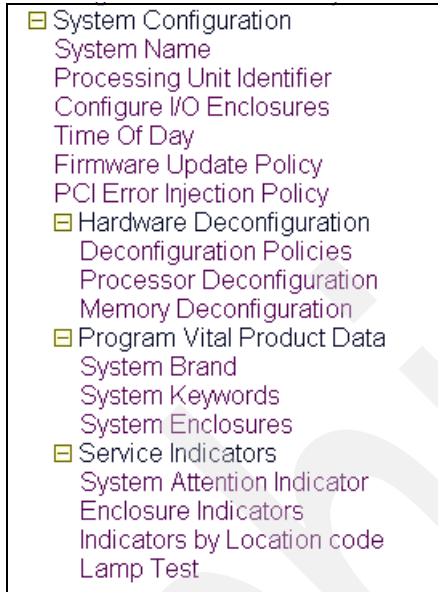


Figure 10-62 Expanded Service Indicators menu

System attention indicator

Figure 10-63 shows the system attention indicator.

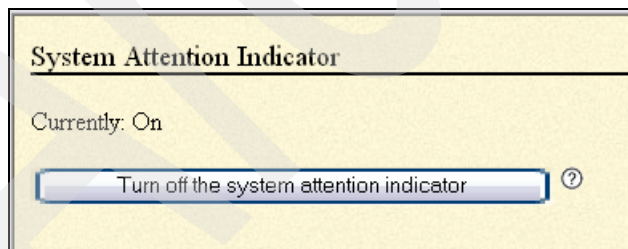


Figure 10-63 System attention indicator

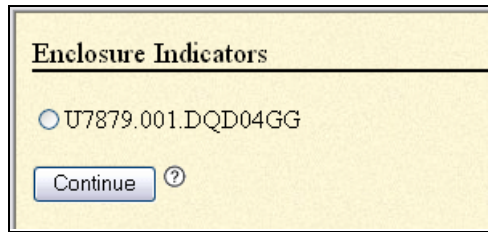
System attention indicator

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 once again.

Enclosure indicators

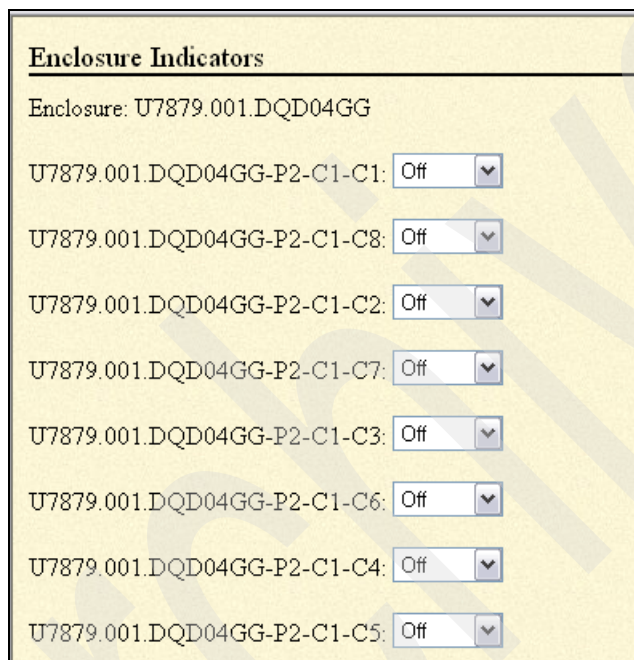
Figure 10-64 shows the enclosure indicators.



The screenshot shows a window titled "Enclosure Indicators" with a yellow background. It contains a radio button next to the enclosure ID "U7879.001.DQD04GG". Below the radio button is a "Continue" button with a question mark icon.

Figure 10-64 Select enclosure indicators

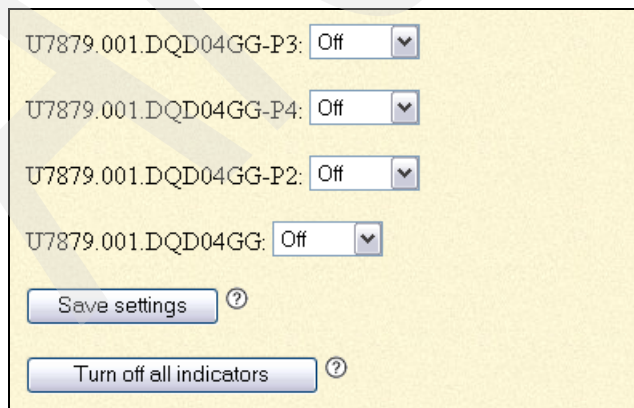
Select enclosure and continue (Figure 10-65 and Figure 10-66).



The screenshot shows a window titled "Enclosure Indicators" with a yellow background. The enclosure ID "U7879.001.DQD04GG" is displayed at the top. Below it, there are eight rows, each with a label and a dropdown menu set to "Off":

- U7879.001.DQD04GG-P2-C1-C1: Off
- U7879.001.DQD04GG-P2-C1-C8: Off
- U7879.001.DQD04GG-P2-C1-C2: Off
- U7879.001.DQD04GG-P2-C1-C7: Off
- U7879.001.DQD04GG-P2-C1-C3: Off
- U7879.001.DQD04GG-P2-C1-C6: Off
- U7879.001.DQD04GG-P2-C1-C4: Off
- U7879.001.DQD04GG-P2-C1-C5: Off

Figure 10-65 Enclosure indicators, part 1 of 2



The screenshot shows a window titled "Enclosure Indicators" with a yellow background. It contains four rows, each with a label and a dropdown menu set to "Off":

- U7879.001.DQD04GG-P3: Off
- U7879.001.DQD04GG-P4: Off
- U7879.001.DQD04GG-P2: Off
- U7879.001.DQD04GG: Off

At the bottom, there are two buttons: "Save settings" with a question mark icon, and "Turn off all indicators" with a question mark icon.

Figure 10-66 Enclosure indicators, part 2 of 2

Off/Identify = the two options.

Continue

Click this button to display another page of indicators for the selected enclosure.

Save settings

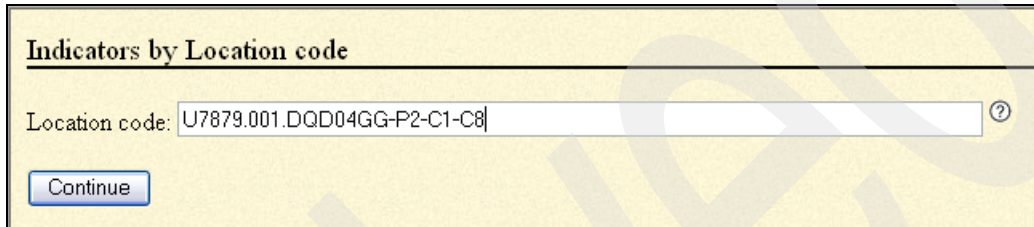
Click this button to update the state of all the indicators for this enclosure.

Turn off all indicators

Click this button to turn off all the indicators for this enclosure.

Indicators by location code

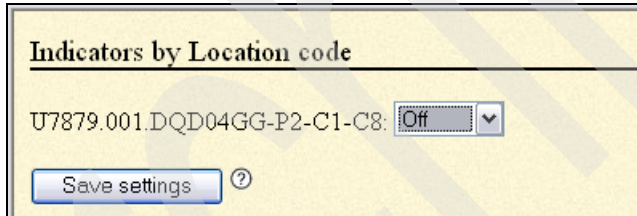
This is the same as the previous section (enclosure indicators) if you already know the location code (Figure 10-67).



The screenshot shows a panel titled "Indicators by Location code". Below the title is a text input field with the value "U7879.001.DQD04GG-P2-C1-C8" and a help icon. Below the input field is a "Continue" button.

Figure 10-67 Indicators by location code

U7879.001.DQD04GG-P2-C1-C8 is an example of a Location code for Enclosure: U7879.001.DQD04GG. In this example we enter that code and select **Continue**. You will see a panel similar to Figure 10-68.



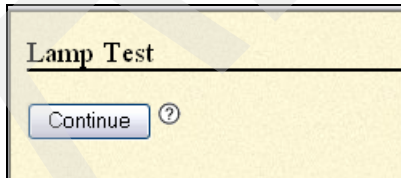
The screenshot shows the same panel as Figure 10-67, but now the location code is "U7879.001.DQD04GG-P2-C1-C8:" followed by a dropdown menu showing "Off". Below this is a "Save settings" button with a help icon.

Figure 10-68 Changing indicators by location code

Off/Identify = the two options.

Lamp test

Figure 10-69 shows the lamp test.



The screenshot shows a panel titled "Lamp Test". Below the title is a "Continue" button with a help icon.

Figure 10-69 Lamp test

Lamp test

Click this button to perform the lamp test. The test changes all indicators to the “on” state for about four minutes.

10.3.4 Network services

Figure 10-70 shows the expanded Network Services menu. Using this menu, you can:

- ▶ Display or change the ethernet port network configurations for the service processor.
- ▶ Display or change the IP addresses that are allowed access to the service processor ethernet ports.

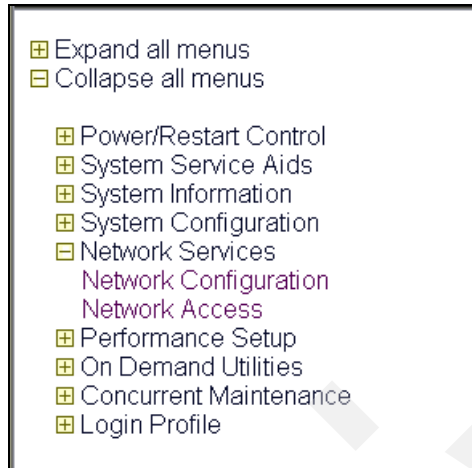


Figure 10-70 Network Services menu

Network configuration

Using this option you can display, or display and change, the system's ethernet network interfaces to the service processor. You can change the network configuration only when the system is powered off. If your system is powered on, you will see panels as shown in Figure 10-71 for ethernet service port 0 (eth0), and Figure 10-72 for ethernet service port 1 (eth1).

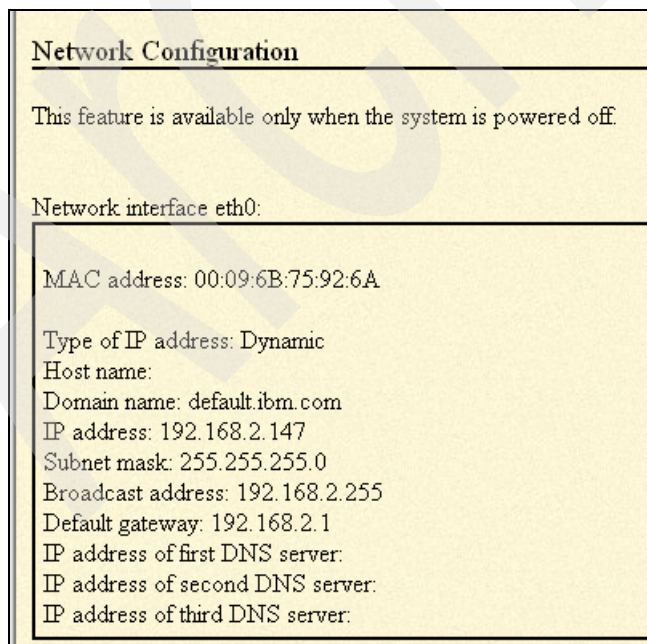


Figure 10-71 Display of network interface port 0

Network interface eth1:

MAC address: 00:09:6B:75:92:6B

Type of IP address: Static

Host name: fitz

Domain name: rchland.ibm.com

IP address: 9.5.17.204

Subnet mask: 255.255.255.128

Broadcast address: 9.5.17.255

Default gateway: 9.5.17.129

IP address of first DNS server: 9.10.244.100

IP address of second DNS server: 9.10.244.200

IP address of third DNS server:

Figure 10-72 Display of network interface port 1

When the system is powered off, you will be able to see the current network settings and you will also be able to make changes to the network configuration. You can select **Configure this interface** (for eth0, eth1, or both) and then click the **Continue** button. In Figure 10-73 we will continue with the example of selecting to change the configuration for ethernet service port 1 (eth1).

Network interface eth1:

Configure this interface? ?

MAC address: 00:09:6B:75:92:6B

Type of IP address: Static ?

Host name: fitz ?

Domain name: rchland.ibm.com ?

IP address: 9.5.17.204 ?

Subnet mask: 255.255.255.128 ?

Default gateway: 9.5.17.129 ?

IP address of first DNS server: 9.10.244.100 ?

IP address of second DNS server: 9.10.244.200 ?

IP address of third DNS server: ? ?

Continue

Figure 10-73 Changing network configuration for port eth1

Make any changes that you want based upon the following descriptions, and once again, select **Continue**. You will see a confirmation panel appear, similar to Figure 10-74.

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.
- ▶ If Static is selected, the IP address, subnet mask, broadcast address, default gateway, and first DNS server address must be entered. The second and third DNS server addresses are optional.

Host name

Enter a new value for the hostname. The valid characters are: hyphen and period [- .]; upper and lower case alphabetic [A - Z] and [a - z]; 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 hostname 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 '-', '_', and '.' 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, if a second DNS server is being used.

IP address of third DNS server

Enter a new value for the third DNS server, if a third DNS server is being used (Figure 10-74).

Eth0:
Type of IP address: Dynamic Host name: Domain name: default.ibm.com
Eth1:
Type of IP address: Static Host name: fitz Domain name: rchland.ibm.com IP address: 9.5.17.204 Subnet mask: 255.255.255.128 Broadcast address: 9.5.17.255 Default gateway: 9.5.17.129 IP address of first DNS server: 9.10.244.100 IP address of second DNS server: 9.10.244.200 IP address of third DNS server:
Attention: This operation will reset the service processor. Current network sessions, including hardware management console (HMC) connections, will be terminated. The new settings must be used to re-establish network connections.
<input type="button" value="Save settings"/>

Figure 10-74 Confirm network changes

Selecting **Save Settings** will then cause the network configuration changes to be made and the service processor to be rebooted. As the service processor reboots, your ASMI session will be dropped and you will have to reconnect your session to continue. When you reconnect you will be using the new settings.

Attention: If incorrect network configuration information is entered, you may not be able to use the ASMI after the service processor reboots. To remedy this situation, your hardware service provider may have to reset the service processor to the default settings. This procedure will also reset all user IDs and passwords to their default values.

Network access

Using this option, you can display or change the IP addresses that are allowed access to the service processor ethernet ports. You can specify a list of allowed IP addresses and a list of denied IP addresses. An example is shown in Figure 10-75.

Allowed IP addresses ?	Denied IP addresses ?
1. 9.10.136.221	1. 9.10.136.224
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"/>

Save settings

Figure 10-75 Change Network access IP addresses allowed and denied

Make any changes that you want based upon the following descriptions and select **Save Settings** from the bottom of the panel shown in Figure 10-75. In this example we will be allowing access for one IP address and denying access for one IP address. You will see a confirmation panel appear, similar to Figure 10-76.

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. “ALL” is a valid IP address that can be entered. An empty allowed list is ignored and access is granted from any IP address.

Tip: The IP address of the browser you are currently using to connect to ASMI is shown in the Network Access panel. In our example, Figure 10-75 shows our IP address as: 9.10.136.220

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 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 may be permanently denied. In this case, reset the network parameters by using the network reset parameters switch on the service processor card. Note that 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. See Figure 10-76.

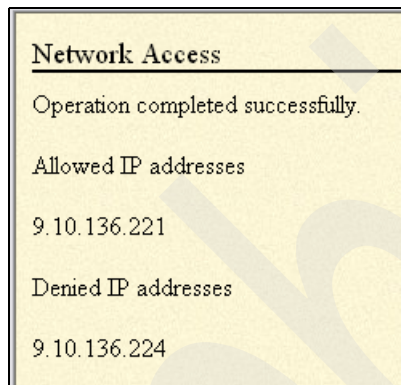


Figure 10-76 Confirm network IP addresses allowed and denied

10.3.5 Performance setup

Figure 10-77 shows the expanded Performance Setup menu. Using this menu, you can select the logical memory block size for the system.

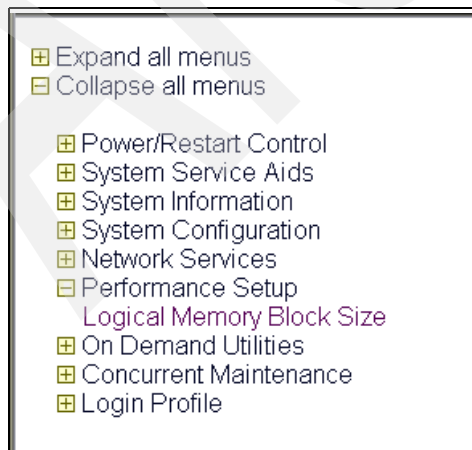


Figure 10-77 Performance Setup menu

Logical memory block size

Using this option, you can display or change the Logical Memory Block (LMB) size used by your system. To display or change the memory block size currently used, you would select the Logical Memory Block Size option from the Performance Setup menu. You will be presented with a panel similar to the one in Figure 10-78.

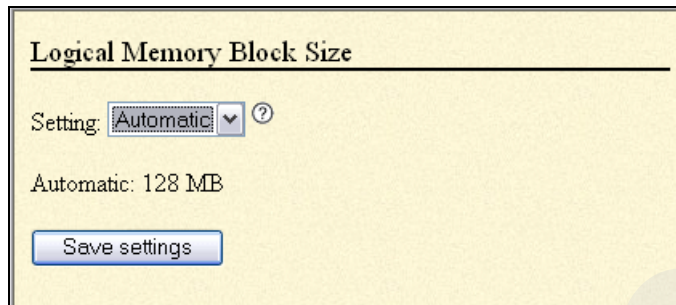


Figure 10-78 Logical memory block size selection

The designation of the logical memory block size you intend to use is a system wide parameter, so that everything on a given system will be using the same size memory blocks. In the past we have been able to assign or move memory from a partition at the single MB level. We could move as little as 1 MB into or out of a partition. Memory on the System i5 machines cannot be manipulated that finely. Instead it must be assigned or moved in blocks. The LMBs are also sometimes referred to, more accurately, as Memory Regions. This is similar to the method which the current Model 690 pSeries servers handle memory.

Care should be taken when manually specifying a memory block size, however, because memory block size selection has an effect on:

- ▶ Initial Program Load (IPL) times
- ▶ Memory granularity

For this reason we recommend that you allow the system to remain at the default of Automatic.

How logical memory block size affects IPL times

There is some increased IPL times associated with using small LMBs (16MB, or 32MB) when compared to using the larger sizes of LMBs (64MB, 128MB, or 256MB). So choosing a small LMB size will increase your IPL time, sometimes dramatically, if this is a very large system.

This is a consideration for all System i5 systems whether the server is being used as a single server image (a single partition) or as a LPAR machine with several images. If the server was being used only as a single image you might want to have a larger LMB size that would produce a shortened IPL time.

How logical memory block size affects memory granularity

We had stated earlier that, in the past, we have been able to assign or move memory from a partition at the single MB level. We could move as little as 1 MB into or out of a partition. Memory on the System i5 machines cannot be manipulated that finely. Instead it must be assigned or moved in blocks. So rather than a memory granularity of 1MB, the System i5 memory granularity will be between 16MB and 256MB. This then becomes the smallest unit of memory that can be manipulated. This becomes the smallest unit of memory that can be taken from an LPAR, inserted into an LPAR, or moved between LPARs.

For System i5 systems running as a single server image this has much less impact. You will seldom need to be concerned with granularity of the memory because you have *all* of the system's memory assigned to a single partition.

Note: All System i5 systems require that some amount of system memory must be allocated for the controlling Hypervisor. Selecting a larger MLB size may have an effect on the amount of memory the system will require to be assigned to the Hypervisor. This may slightly reduce the amount of memory remaining to be used for system or LPAR assignment.

Changing logical memory block size

Using the panel shown in Figure 10-78 you can select a new **Setting** value for the memory block size for your system and click the **Save** settings button. Allowable values are 16 MB, 32 MB, 64 MB, 128 MB, 256 MB, and Automatic.

The *Automatic* setting instructs the service processor to compute a value based on the total amount of available physical memory installed on your system. (Even if your system is not fully using the CoD portion of the installed memory, the calculation is based upon the total amount of physical memory installed.)

Updates to the logical memory block size do not take effect until the next system reboot. For most systems we recommend that this parameter be set to Automatic.

10.3.6 On demand utilities

Capacity on Demand (CoD) allows you to permanently activate inactive processors or inactive system memory without requiring you to restart your server or interrupt your business. You must first have sufficient inactive processor or inactive memory capacity physically installed on your system. Then you will order and receive activation keys allowing you to enable this capacity for your use.

Figure 10-79 shows the expanded On Demand Utilities menu. Using this menu, you can:

- ▶ Display CoD ordering information for your system (required for placing an order for adding capacity).
- ▶ Enter CoD activation keys.
- ▶ Allow new processors or new memory to become allocated once CoD activation keys have been entered.
- ▶ Enter special CoD commands that may be given to you by your hardware service provider.

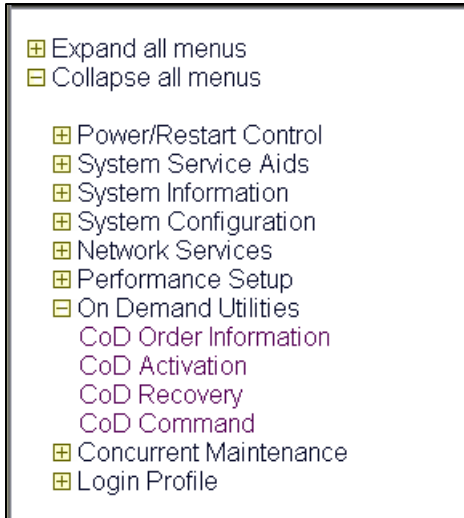


Figure 10-79 On Demand Utilities menu

CoD order information

You can use this option to generate the system information required when you need to order additional processor or memory activation features from IBM or your business partner. When you place an order with IBM or your business partner, you will receive activation codes (keys) that must be entered into the system.

Note: This feature is not available to be displayed prior to the system server firmware reaching the state of standby mode.

Figure 10-80 shows an example of the system information displayed. The information will be different for your system. You must print or record this information so that it will later be available for your order.




Figure 10-80 Example: Displaying CoD order information

Important: If your system has an HMC available for managing the system, we recommend that you use the HMC to perform the ordering and activation tasks for increasing Capacity on Demand.

CoD activation keys

You will use the display shown in Figure 10-81 to enter the processor and memory activation keys provided to you by IBM or your business partner. You may have more than one key to enter.

Note: This feature is not available prior to the system server firmware reaching the state of standby mode.

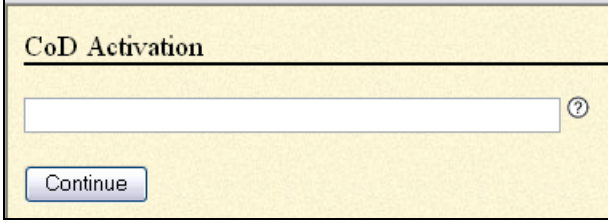


Figure 10-81 CoD activation

Enter the CoD activation key and click **Continue**. You may need to enter one or more keys. To receive the key, call the hardware provider with the information displayed on the CoD Order Information display.

Entering the key or keys does not make the processor or memory immediately available to be used. You must first complete the action to allow CoD Recovery in order for the keys to be recognized by the firmware and the hardware to be made available.

Enabling activation keys using CoD recovery

You will use the display shown in Figure 10-82 to enable the processor and memory activation keys. By continuing with CoD recovery, the firmware causes the CoD key(s) to become recognized and the additional hardware becomes allocated and available for use.

Note: This feature is not available prior to the system server firmware reaching the state of standby mode.



Figure 10-82 CoD Recovery

Click the **Continue** button after all CoD activation keys have been entered.

CoD commands

There may be situations where your service provider may need to have CoD commands entered into the system firmware. The service provider will specify the command, which may then be entered onto the panel shown in Figure 10-83.

Note: This feature is not available prior to the system server firmware reaching the state of standby mode.

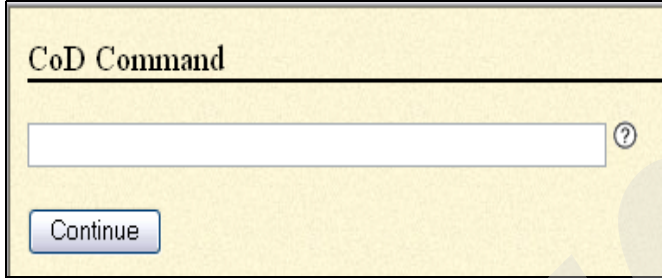


Figure 10-83 Enter CoD command (optional)

Enter a CoD command. If needed, the command is supplied by hardware provider. You will need to use the command response to complete the CoD recovery process.

10.3.7 Concurrent maintenance

Using ASMI, your hardware service provider can provide concurrent maintenance for some of the system hardware. Figure 10-84 shows the expanded Concurrent Maintenance menu.

Using this menu, you can:

- ▶ Remove and repair a control panel that exists in a processor frame, or add a new control panel.
- ▶ Turn on or off power to an integrated drive electronics (IDE) device for removal or repair. Examples of IDE devices are CD-ROM drives or a DVD-R/W drive.

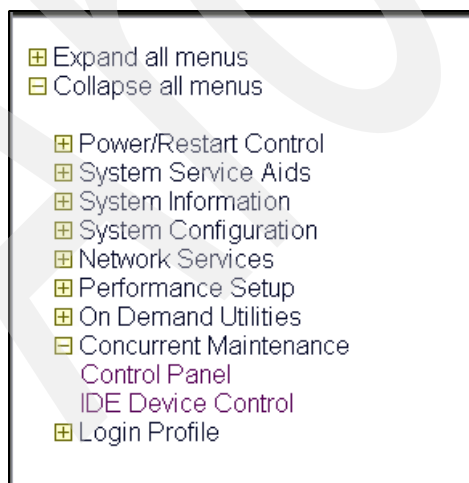


Figure 10-84 Concurrent Maintenance menu

Control panel

Selecting Control Panel from the Concurrent Maintenance menu will show the display in Figure 10-85. From this display you can click **Continue** to remove and replace an existing control panel, or to add a new control panel. You could, for example, remove an existing control panel that has become inoperative and replace it with a new control panel.

This option prepares the control panel for concurrent maintenance by *logically* isolating the control panel. As a result, the firmware does not recognize the control panel as being active and you can remove it without powering off the frame.



Figure 10-85 Control Panel

Select an action, either Install or Remove, and click the **Continue** button. The next page displays either a location code for the control panel, or a list of location codes for all control panels present. Most system will only have a single control panel. Our example in Figure 10-86 shows that we are going to remove a control panel with concurrent maintenance.

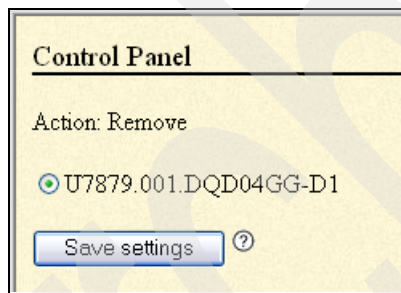


Figure 10-86 Example: Control panel removal using concurrent maintenance

Attention: If you *remove* the control panel, you should return later to *install* the control panel after the hardware has been replaced.

Choose a control panel for the selected action and click the **Save** settings button. You will be presented with a confirmation panel as shown in Figure 10-87. It is now safe to have the service representative do a physical removal, or removal and replacement of the control panel.

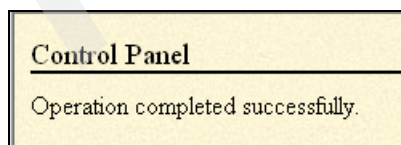


Figure 10-87 Confirm control panel removal

With the control panel now replaced, you can now return and use the Install action shown in Figure 10-88 to activate the new control panel, completing the concurrent maintenance procedure.

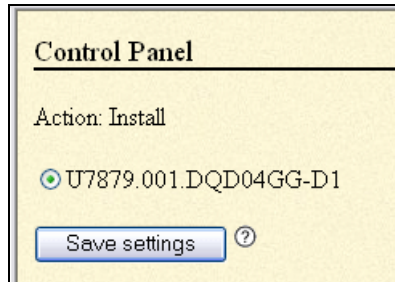


Figure 10-88 Example: Control panel install using concurrent maintenance

IDE Device Control

An IDE device can be either a CD-ROM drive or a DVD-R/W drive. When you select IDE Device Control from the Concurrent Maintenance menu you will see a panel similar to Figure 10-89. From this panel, your hardware service provider can perform concurrent maintenance on any failing IDE device.

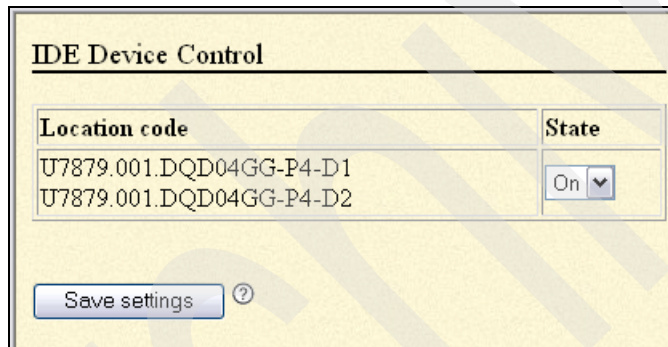


Figure 10-89 IDE Device Control selection

At a minimum, your system will have one IDE device. Most systems, however, will have several IDE drives installed. This is especially true for large systems, which may have several drives installed to support several LPARs. Your system may have a different number of drives than we have shown in our example.

A system processor frame can contain up to two IDE devices per frame. In addition, a system expansion frame might contain up to four IDE devices per frame. When four devices are installed in an expansion frame they would typically be installed with a “pair” of devices in the top half of the frame and a second “pair” of devices in the bottom half of the frame. Power to the IDE devices, a pair at a time, can be turned off and on without affecting power to the rest of the devices in the system.

To perform concurrent maintenance, you would identify, by location code, the failing IDE device you want to repair. You would then change the state of the “pair” of devices as shown in Figure 10-90 by selecting a state of **Off** and clicking the **Save settings** button.

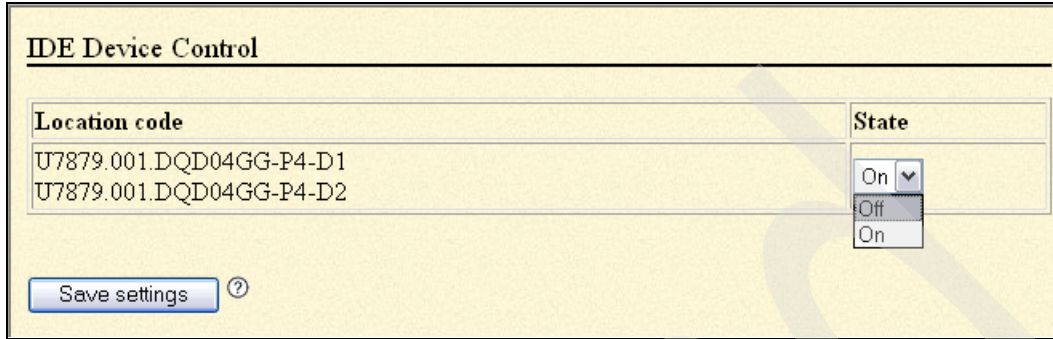


Figure 10-90 Example: IDE device power off using concurrent maintenance

You will next see a confirmation panel as shown in Figure 10-91. It is now safe to have the service representative do a physical removal, or removal and replacement of the failing IDE device.

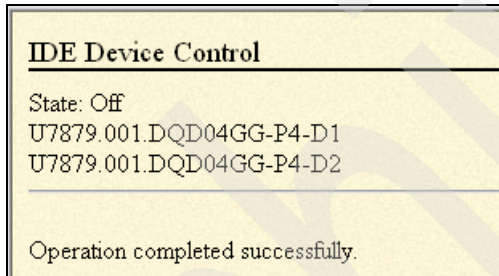


Figure 10-91 Confirm IDE device power off

Attention: If you power *Off* a pair of IDE devices, you should return later to power *On* the IDE devices after the failing hardware has been replaced.

With the IDE device now replaced, you can now return and use the action shown previously in Figure 10-89 on page 334 to power *On* the pair containing the new IDE device, completing the concurrent maintenance procedure.

10.3.8 Login Profile

You must use a Login Profile each time you access the ASMI menus. This Login Profile consists of a User ID and Password set to a specific authority level. Figure 10-92 shows the expanded Login Profile menu. Using this menu, you can:

- ▶ Change the password for a user.
- ▶ Display the successful or failed login attempts to the service processor.
- ▶ Select or change the default language used when accessing the ASMI Welcome panel.
- ▶ Select the language to be installed on the next firmware update.

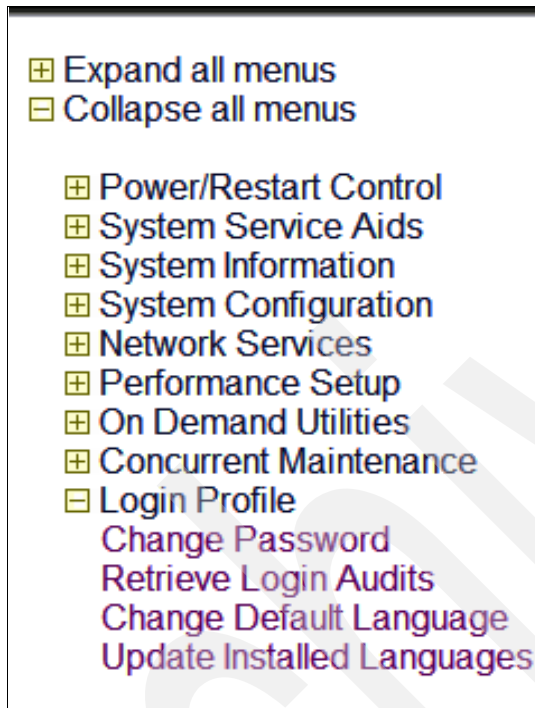


Figure 10-92 Login Profile menu

Change password

Using this menu option you can change the password for a ASMI user. You can only change the password for the User IDs of *general*, *admin*, or *HMC*. Figure 10-93 shows the panel you would use to change passwords.

If you are a general user, you can only change your own password. If you are an administrator, you can change your password and the passwords for the general user. If you are an authorized service provider, you can change your password, the passwords for general and administrator IDs, and the HMC access password. After you make the User ID selection and enter the Current password and New password (twice), you can click the **Continue** button to make the password change.

Figure 10-93 Change password

User ID to change

Select the user ID of the user whose password you wish to change. Choices you can make are general, admin, or HMC.

Current password for current user

As a security measure, the current password must be supplied. The initial factory default passwords are set to:

- ▶ **genera1** for the User ID general
- ▶ **admin** for the User ID admin
- ▶ **abc123** for the HMC user ID of **hscroot** (may have been changed during HMC guided setup).

New password for user

Enter the new password for the user whose password you wish to change. Passwords can be any combination of up to 64 alphanumeric characters.

New password again

Enter the new password for the user again for verification.

Click the **Continue** button to make the password change.

Retrieve login audits

Using this menu option you can display an audit log of both the successful ASMI logins to your system as well as the failed login attempts. An example of the successful login audit log is shown in Figure 10-94. An example of the failed login audit log is shown in Figure 10-95.

You will see up to the last 20 successful logins and up to the last 20 failed login attempts. You can only display the information, you cannot change or delete the information. You may see User IDs of **dev** and **ce1ogin** which are special User IDs available to authorized service personnel only. As a user you cannot use or change these User IDs. They are highly secure due to the fact that the service processor uses a function to regenerate and change the passwords on a daily basis and only the IBM authorized service personnel can have access to and use these special passwords.

In the unsuccessful login audit log you may also see an entry of *unknown*. This occurs when an attempt is made to login with any other User ID other than the ones supplied by IBM that have been described in this section and the Change password section.

Retrieve Login Audits		
Successful logins		
User ID	Date	Location
dev	Tue Jan 13 22:53:38 2026	9.10.127.17
dev	Tue Jan 13 23:35:29 2026	9.10.127.17
dev	Wed Jan 14 04:04:44 2026	9.53.179.169
dev	Wed Jan 14 22:20:07 2026	9.10.227.24
dev	Wed Jan 14 22:32:18 2026	9.10.227.24
dev	Wed Jan 14 22:35:52 2026	9.10.227.24
dev	Thu Jan 15 05:37:01 2026	9.10.227.24
admin	Fri Jan 16 05:16:20 2026	9.10.136.27
admin	Fri Jan 16 05:17:36 2026	9.41.179.222
dev	Fri Jan 16 05:18:53 2026	9.41.179.222

Figure 10-94 Show successful login attempts

Failed logins		
User ID	Date	Location
unknown	Wed Jan 14 22:19:53 2026	9.10.227.24
dev	Wed Jan 14 22:31:58 2026	9.10.227.24
celogin	Fri Jan 16 05:18:31 2026	9.41.179.222
celogin	Fri Jan 16 05:18:42 2026	9.41.179.222
dev	Fri Jan 16 05:20:27 2026	9.10.136.27
celogin	Fri Jan 16 05:23:01 2026	9.41.179.222
admin	Fri Jan 16 05:23:19 2026	9.41.179.222
dev	Mon Sep 13 23:36:17 2004	9.10.136.220

Figure 10-95 Show failed login attempts

Change default language

Using the display shown in Figure 10-96, you can change the default language for ASMI users. This controls the language that is displayed on the ASMI Welcome panel prior to login, and during your ASMI session if you do not choose an alternative language at the time of login.

Change Default Language	
Setting:	English ?
<input type="button" value="Save"/>	<ul style="list-style-type: none"> English Français Deutsch Italiano Español

Figure 10-96 Changing the default language

From the pull down menu, select the default language to use for the ASMI Welcome panel and click the **Save settings** button. You will receive a confirmation panel similar to Figure 10-97. The change will take place in just a few minutes with no restart of the service processor firmware required.

Note that a user can override the default at ASMI login time. If no override is selected at login, then the default language is used for that session.

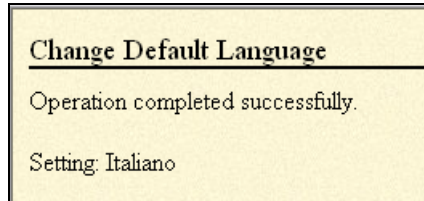


Figure 10-97 Confirm change of default language

Update Installed languages

Using the display shown in Figure 10-98, you can change the installed language.

Select which languages to install on next firmware update. English is always installed as part of a firmware update. A maximum of five languages may reside on the service processor; choose up to four languages.

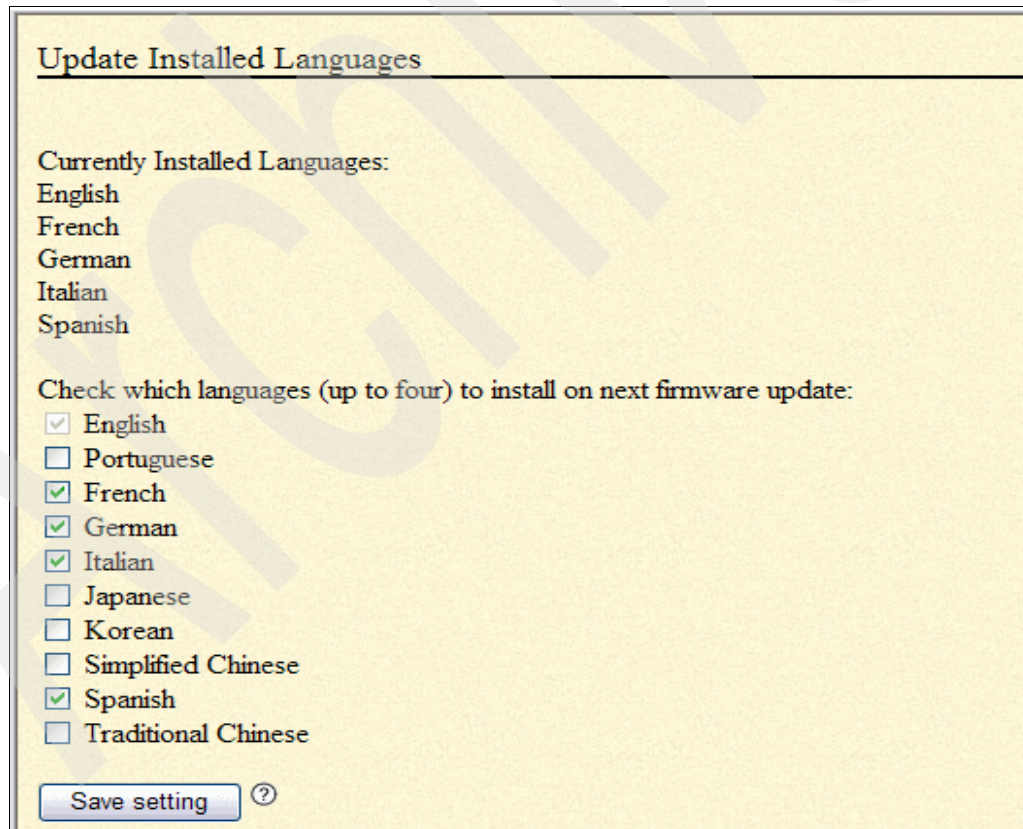


Figure 10-98 Changing the installed language

Archived

OpenSSH

The licence program 5733-SC1 contains the OpenSSH (Secure SHell), OpenSSL, and zlib open source packages ported to i5/OS using the i5/OS PASE runtime environment. The SC1 licensed program requires i5/OS V5R3 or later and also requires that i5/OS Option 33 (i5/OS PASE - Portable Solutions Application Environment) be installed.

TCP/IP connectivity applications such as telnet and ftp transmit data and passwords over the network in plain text. This means that the data and passwords can be intercepted and read by other users on the network. The SSH protocol suite is a software solution that provides secure alternatives for telnet and ftp. SSH verifies the authenticity of both the client and server, and all of the data (including userids and passwords) is encrypted as it travels on the network. This encryption is done transparently to the end user.

OpenSSH is the open source implementation of the SSH protocol suite. OpenSSH is widely available for use on many other platforms including Linux, AIX, and z/OS.

For more details about OpenSSH, refer to the Web site:

<http://www.openssh.org/>

For more details about i5/OS PASE, refer to the Web site:

<http://publib.boulder.ibm.com/infocenter/iseres/v5r3/ic2924/info/rzalf/rzalfwhatispase.htm>

11.1 Utilities available in Open SSH

The following utilities are available in Open SSH:

1. **ssh** — A secure telnet replacement that allows an i5/OS user to connect as a client to a server running the **sshd** daemon. An **ssh** client can also be used to connect to the HMC on the IBM @server@ 5xx iSeries models.
2. **sftp** — A secure ftp replacement. As with all implementations of **sftp** on other platforms, **sftp** can only transfer data in binary format. Note that **sftp** also does not provide the enhanced functions available in the i5/OS ftp utility when transferring files in the QSYS.LIB file system, nor does it provide the CCSID data conversion options available in the i5/OS ftp utility.
3. **scp** — A secure file copy program — basically an alternative to **sftp** for copying a single file in the integrated file system (IFS).
4. **ssh-keygen** — A public/private key generation and management tool. SSH allows users to authenticate using these public and private keys as an alternative to using their OS signon password.
5. **ssh-agent** — An authentication agent that can store private keys. The **ssh-agent** allows a user to load their public/private key pass phrase into memory to avoid needing to retype the pass phrase each time an SSH connection is started.
6. **sshd** — The daemon that handles incoming **ssh** connections. The **sshd** daemon utility allows users to connect to i5/OS via an **ssh** client. In contrast to connecting to i5/OS via telnet and being presented with a 5250 screen interface, users that connect via **ssh** to an i5/OS system running the **sshd** daemon will be presented with a character interface and an i5/OS PASE command line.

More details on these utilities can be found at:

<http://www.openssh.org/manual.html>

11.2 Installing the OpenSSH licence program in i5/OS

Here is the procedure for you to follow:

- ▶ Install the licence program 5722-SS1 Option 33 - **Portable Solutions Application Environment**. You can see the installed licence program shown in Figure 11-1.
- ▶ Install the licence program **IBM Portable Utilities for i5/OS** using the command:
`RSTLICPGM LICPGM(5733SC1) DEV(OPTxx) OPTION(*BASE) RSTOBJ(*ALL) LNG(2924)`
- ▶ Install the licence program 5733SC1 Option 1 **OpenSSH, OpenSSL, zlib** using the command:
`RSTLICPGM LICPGM(5733SC1) DEV(OPTxx) OPTION(1) RSTOBJ(*PGM)`

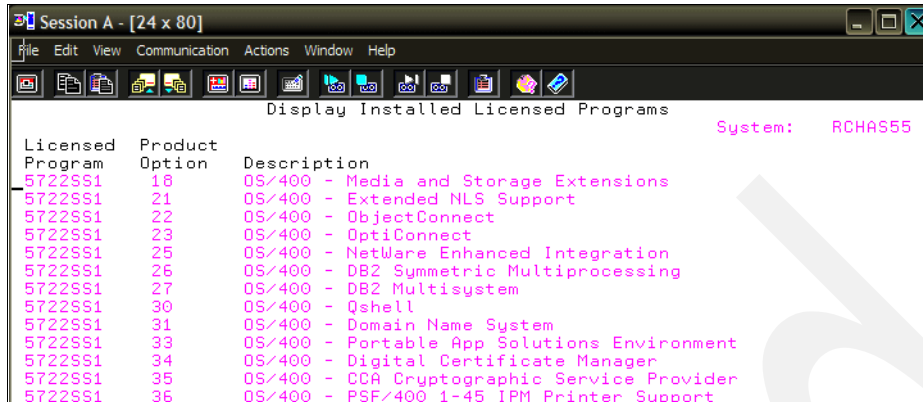


Figure 11-1 Installed licence program -5722SS1-33

The installed licence program 5733SC1 is shown in Figure 11-2.

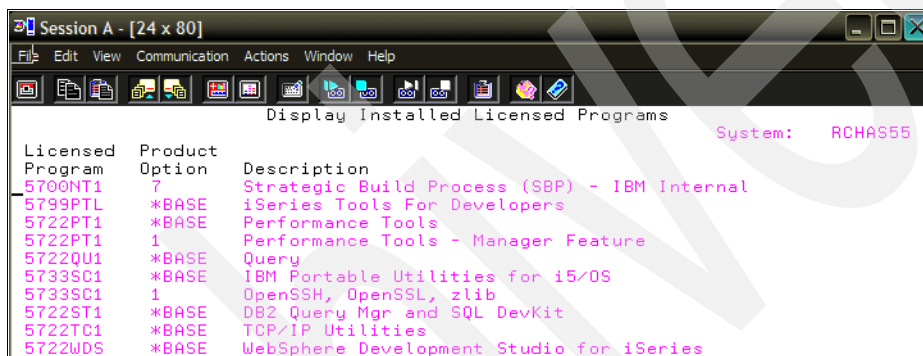


Figure 11-2 Installed licence program -5733SC1

11.3 Using the HMC from i5/OS with OpenSSH

Although the i5/OS, with the OpenSSH license program installed, supports the `ssh` server and client, here we are using OpenSSH as a client. Basically, to work with OpenSSH in i5/OS, you need to create the users on HMC and i5/OS with the same name (Figure 11-3).

Also, an `ssh` authentication key needs to be created on i5/OS and the same key needs to be updated on HMC. The following step-by-step-procedure will guide you to complete these tasks:

- ▶ Logon to the HMC using the `hscroot` userid.
- ▶ Click **HMC Management** from the navigation area.
- ▶ Click **HMC Users**.
- ▶ Click **Manage HMC Users and Access** from the content area.

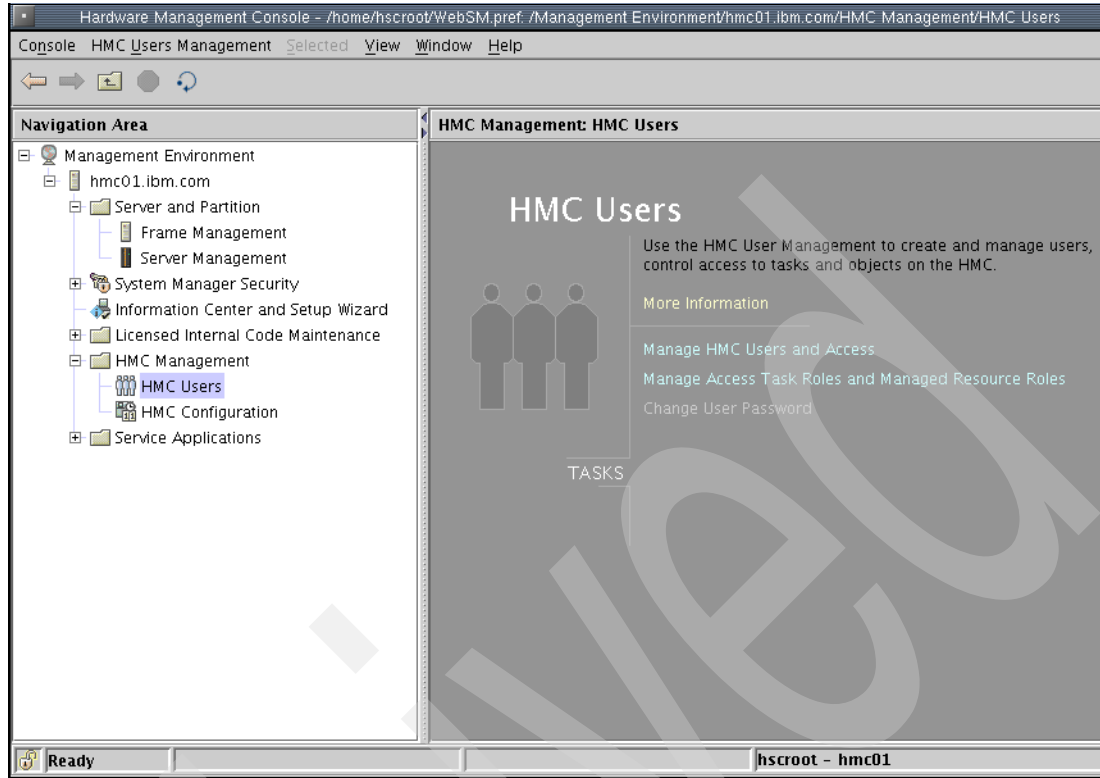


Figure 11-3 HMC Users

- ▶ Click the **User** tab from the User profiles window.
- ▶ Click **Add** to add the user (Figure 11-4).




Figure 11-4 HMC User Add

- ▶ Fill the details of the user you want to create in this **Add User** window, and also select the role of the user from the **Task Roles**. For example, create the user **fred** and select the **Task Role** as **HMC superadmin**.

- ▶ Click **OK** to continue. See Figure 11-5.

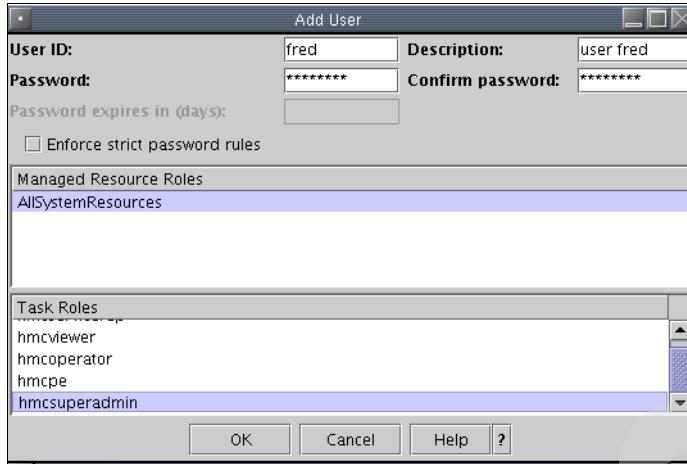


Figure 11-5 Add User Details

- ▶ As shown in Figure 11-6, displaying User profiles will now show Fred in the listing.




Figure 11-6 User (fred) added.

- ▶ Create a user profile in i5/OS. For example, create the user called **fred** as shown in Figure 11-7.

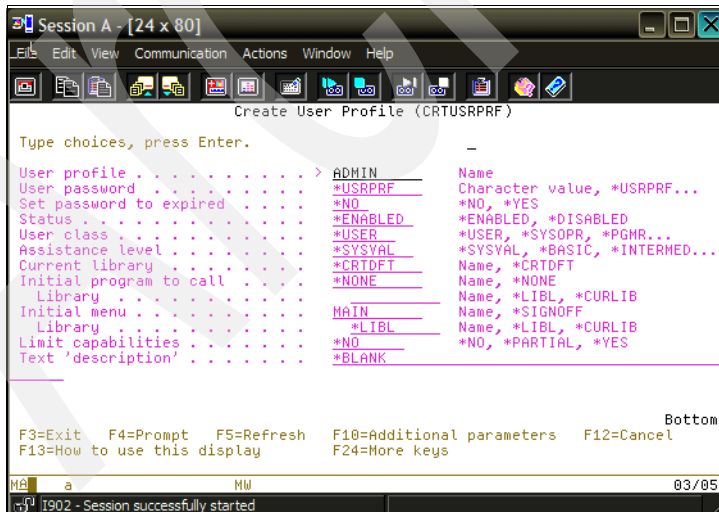


Figure 11-7 User Profile on i5/OS

- ▶ Logon as *fred* from the i5/OS and run the command `call qp2term` (Figure 11-8).




Figure 11-8 Qp2term

- ▶ Create the directory called `fred` under `/home`.
- ▶ Change the ownership of the directory `fred` using the command `chown fred fred`.
- ▶ Go to the directory `cd /home/fred`.

Note: The `qp2term` shell environment is not a true TTY device and this can cause problems when trying to use `ssh`, `sftp`, or `scp` within one of these sessions. Use the `-T` option to not allocate a TTY when connecting.

- ▶ Generate the `ssh` key (Figure 11-9) by using the command `ssh-keygen`.

`ssh-keygen -t rsa`

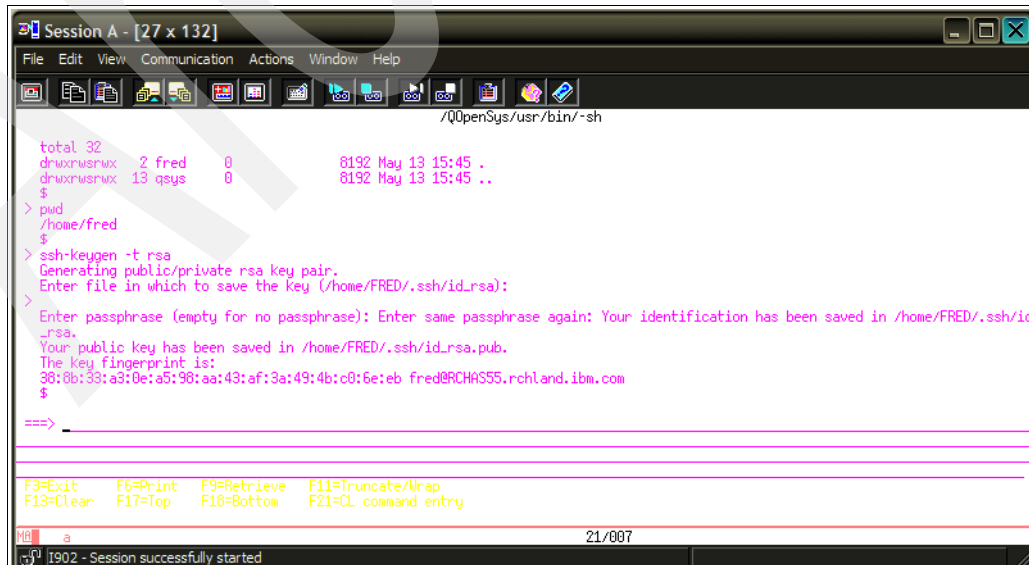


Figure 11-9 SSH-Key generation

The following directory and files will be created under the directory /home/fred:

- /home/fred/.ssh
- /home/fred/.ssh/id_rsa (private key)
- /home/fred/.ssh/id_rsa.pub (public key)

► Go to the directory `cd /home/fred/.ssh` (Figure 11-10).

Note: The write bits for both group and other are turned off for these `ssh` key files. Ensure that the private key has a permission of 600.




Figure 11-10 `ssh` key directory

► Run a command `cat id_rsa.pub` (Figure 11-11).




Figure 11-11 `SSH` Key

- ▶ Copy the key from 5250 emulator screen as shown in Figure 11-12.




Figure 11-12 SSH Key content

- ▶ Establish the connection to HMC from qp2term shell (Figure 11-13) using the command: `ssh -T 9.5.92.92` (For example, 9.5.92.92 is an IP address of the HMC).

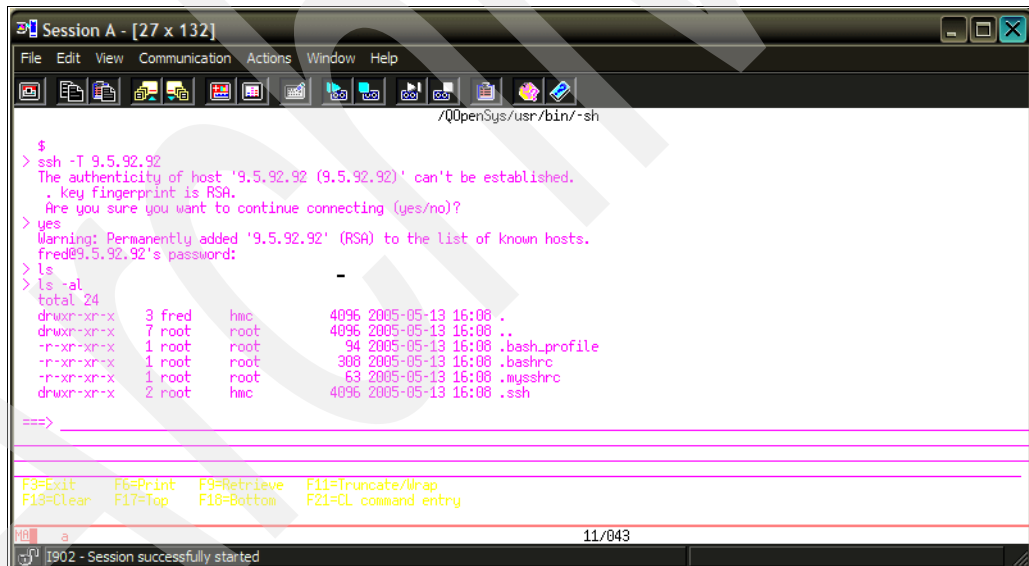
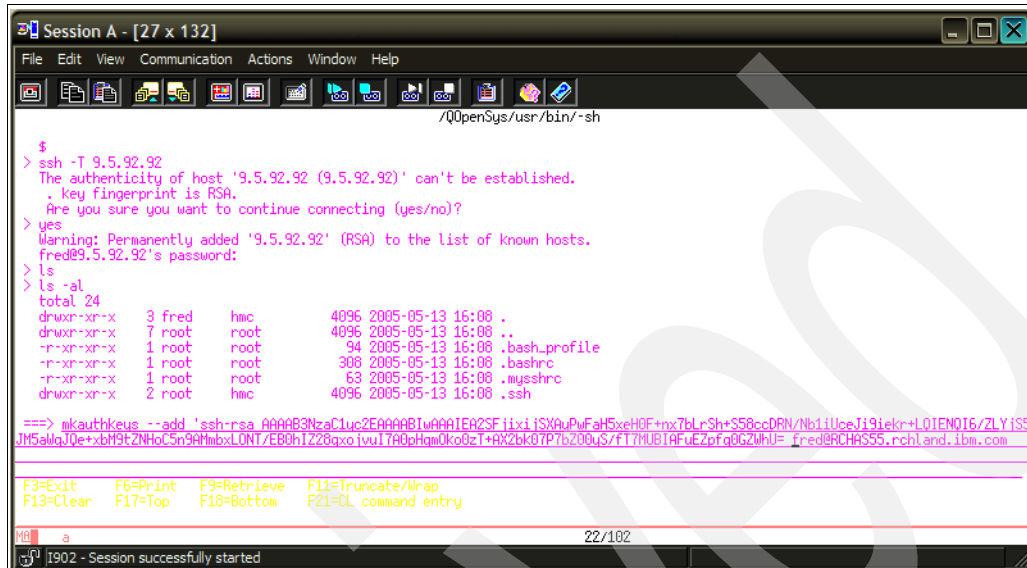


Figure 11-13 SSH to HMC

Follow the instructions shown in Figure 11-13 above to logon into the HMC.

- ▶ Once you logon to the HMC, run the command **mkauthkeys** to authenticate the key which we have generated, and paste the key here, as shown in Figure 11-14.



```
Session A - [27 x 132]
File Edit View Communication Actions Window Help
~/QOpenSys/usr/bin/-sh

$
> ssh -T 9.5.92.92
The authenticity of host '9.5.92.92 (9.5.92.92)' can't be established.
: key fingerprint is RSA.
Are you sure you want to continue connecting (yes/no)?
> yes
Warning: Permanently added '9.5.92.92' (RSA) to the list of known hosts.
fred@9.5.92.92's password:
> ls
> ls -al
total 24
drwxr-xr-x  3 fred  hmc      4096 2005-05-13 16:08 .
drwxr-xr-x  7 root   root      4096 2005-05-13 16:08 ..
-r-xr-xr-x  1 root   root         94 2005-05-13 16:08 .bash_profile
-r-xr-xr-x  1 root   root      308 2005-05-13 16:08 .bashrc
-r-xr-xr-x  1 root   root        63 2005-05-13 16:08 .msshrc
drwxr-xr-x  2 root   hmc      4096 2005-05-13 16:08 .ssh

====> mkauthkeys --add 'ssh-rsa AAAAB3NzaC1uc2EAAAABIwAAAIEA2SFjixiISXAUvFah5xeHOF+nx7blrSh+S58ccDRN/Nb1UceJi9iekr+L0IEN0I6/ZLYiSS
JH5aMqJ0e+xbH9tZNHoC5n9AMbXLONT/EB0hIz28qxoIvuI7A0pHqm0ko0zT+AX2bk07P7b200uS/ft7MUBIAFuEzpf0GZWhUf_fred@RCHASS5_rchland.ibm.com

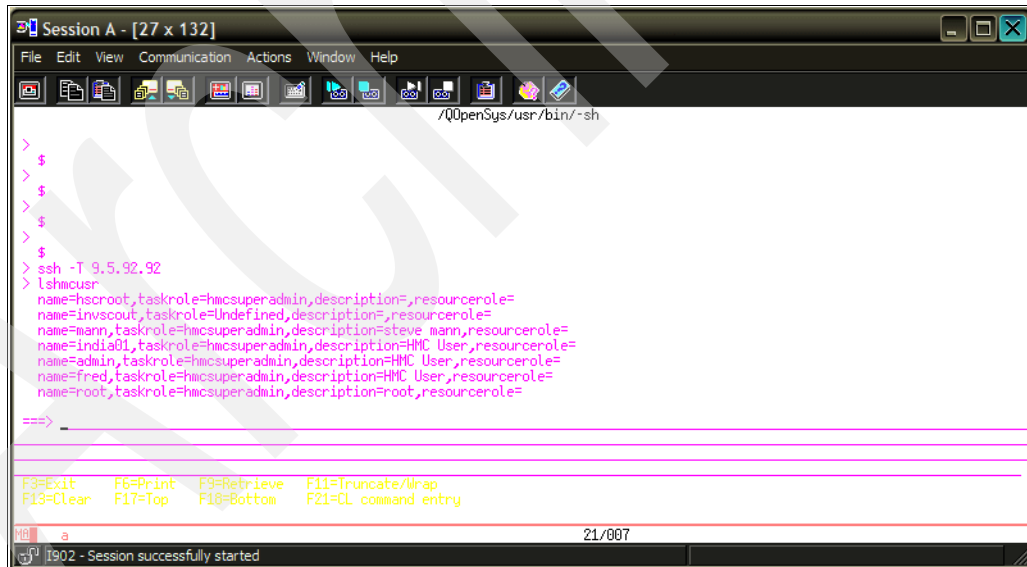
F3=Exit      F6=Print    F9=Retrieve  F11=Truncate/Wrap
F13=Clear    F17=Top     F18=Bottom   F21=CL command entry

MB a 22/102
I902 - Session successfully started
```

Figure 11-14 *Mkauthkeys*

Once the key authentication is done, you can logon to the HMC without userid and password.

- ▶ Run a command **ssh -T 9.5.92.92** to logon to the HMC.



```
Session A - [27 x 132]
File Edit View Communication Actions Window Help
~/QOpenSys/usr/bin/-sh

>
> $
> $
> $
> $
> ssh -T 9.5.92.92
> lshmcusr
name=hmcroot,taskrole=hmcsuperadmin,description=,resource=
name=invscout,taskrole=Undefined,description=,resource=
name=swann,taskrole=hmcsuperadmin,description=steve swann,resource=
name=india01,taskrole=hmcsuperadmin,description=HMC User,resource=
name=admin,taskrole=hmcsuperadmin,description=HMC User,resource=
name=fred,taskrole=hmcsuperadmin,description=HMC User,resource=
name=root,taskrole=hmcsuperadmin,description=Root,resource=

====>

F3=Exit      F6=Print    F9=Retrieve  F11=Truncate/Wrap
F13=Clear    F17=Top     F18=Bottom   F21=CL command entry

MB a 21/007
I902 - Session successfully started
```

Figure 11-15 *HMC logon without password*

11.4 Running DLPAR scripts from i5/OS

Once you logon to the HMC from i5/OS, all the HMC commands are available to use. You can use these commands to do DLPAR functions. Also, you can write the scripts to run the specific tasks and schedule the script to run from i5/OS. Scripts can be written from the desktop PC, and using the operation navigator, drag and drop them to the desired IFS directory (for example, /home/fred) in i5/OS.

To schedule the scripts from the i5/OS, you need to run the commands from the QSHELL instead of QP2TERM.

The QSHELL environment does not have the PATH for **qsh**. You need to add the **qsh** command path to the QSHELL environment. The following steps show how to execute the HMC commands from QSHELL.

- ▶ Logon as *fred* from i5/OS and run the command **strqsh** (Figure 11-17).

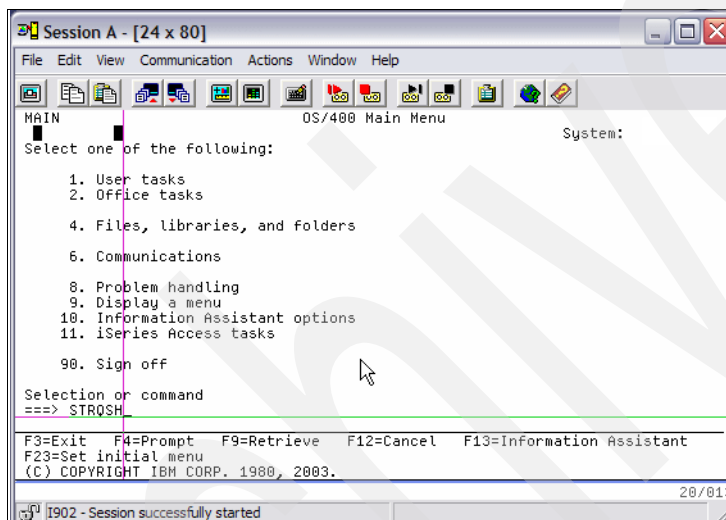


Figure 11-16 *strqsh*

To run the **ssh** command, a path needs to be added. You can see that the following script contains the path for the **ssh**. You can write the script from a Windows workstation and then copy it to the IFS folder (for example, /home/fred), or, from the i5/OS, use the command **edtf**.

You can run the script **systemname** as shown in Example 11-1 and Figure 11-17.

Example 11-1 *Systemname*

```
PATH=$PATH:/QOpenSys/usr/bin:/usr/ccs/bin:/QOpenSys/usr/bin/X11:/usr/sbin:./usr/bin
ssh -T 9.5.92.92 lssyscfg -r sys -F name
```

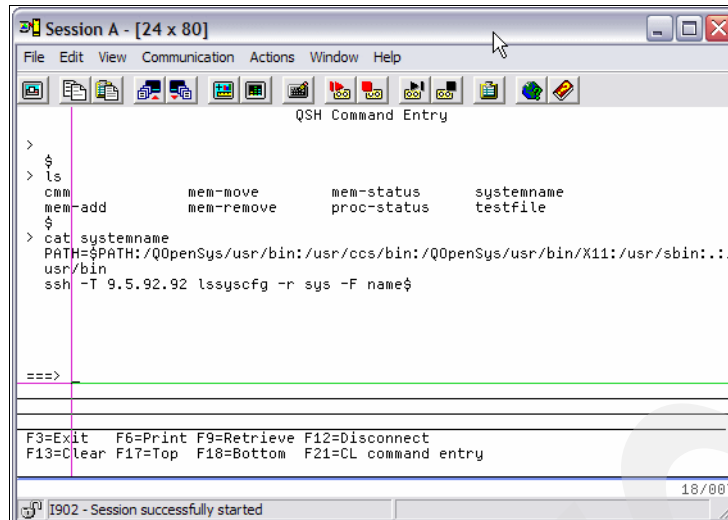
In this script example, the first line will add the **ssh** path and the second line will execute the **ssh** command and logon to the HMC (ip address 9.5.92.92) and then execute the command:

```
lssyscfg -r sys -F name
```

Note: This path need to be added to the scripts to execute the **ssh** command:

```
PATH=$PATH:/QOpenSys/usr/bin:/usr/ccs/bin:/QOpenSys/usr/bin/X11:/usr/sbin:./usr/bin
```

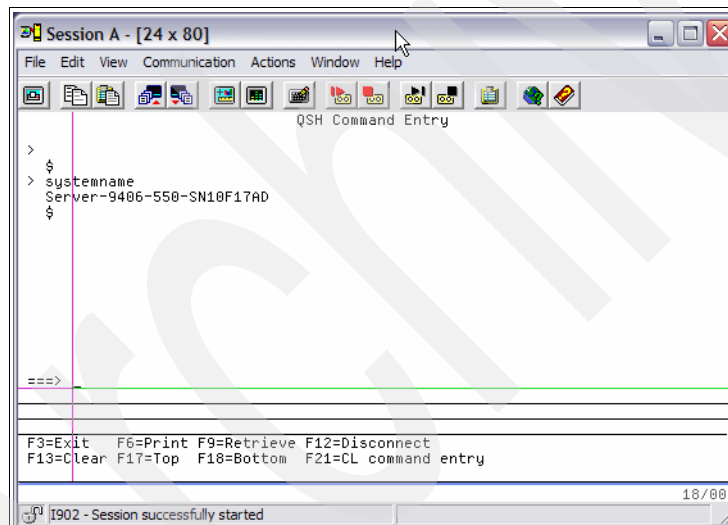
Refer to the script shown in Figure 11-17.



```
Session A - [24 x 80]
File Edit View Communication Actions Window Help
QSH Command Entry
>
$
> ls
cmm          mem-move     mem-status   systemname
mem-add      mem-remove   proc-status  testfile
$
> cat systemname
PATH=$PATH:/QOpenSys/usr/bin:/usr/ccs/bin:/QOpenSys/usr/bin/X11:/usr/sbin:./:/usr/bin
ssh -T 9.5.92.92 lssyscfg -r sys -F name$
===>
F3=Exit  F6=Print F9=Retrieve F12=Disconnect
F13=Clear F17=Top  F18=Bottom F21=CL command entry
18/007
I902 - Session successfully started
```

Figure 11-17 *systemname* - script

- ▶ Run the script (*systemname*) from the QSHELL command prompt (Figure 11-18).



```
Session A - [24 x 80]
File Edit View Communication Actions Window Help
QSH Command Entry
>
$
> systemname
Server-9406-550-SN10F17AD
$
===>
F3=Exit  F6=Print F9=Retrieve F12=Disconnect
F13=Clear F17=Top  F18=Bottom F21=CL command entry
18/007
I902 - Session successfully started
```

Figure 11-18 *systemname*

You can see the output of the script **Server-9406-550-SN10F17AD** as shown here in Figure 11-18.

Similarly, you can write the script to logon to the HMC and perform the specific tasks.

Note: To see the command syntax, logon to the HMC and type `lssyscfg --help` and press Enter.

To move the memory from one partition to other, follow the procedure below:

- ▶ To see the partition memory configuration, run the script `mem-status` as shown below in Example 11-2, Figure 11-19, and Figure 11-20.

Example 11-2 mem-status

```
PATH=$PATH:/QOpenSys/usr/bin:/usr/ccs/bin:/QOpenSys/usr/bin/X11:/usr/sbin:./usr/bin
ssh -T 9.5.92.92 lshwres -m Server-9406-550-SN10F17AD -r mem --level lpar
```




Figure 11-19 mem-status-script

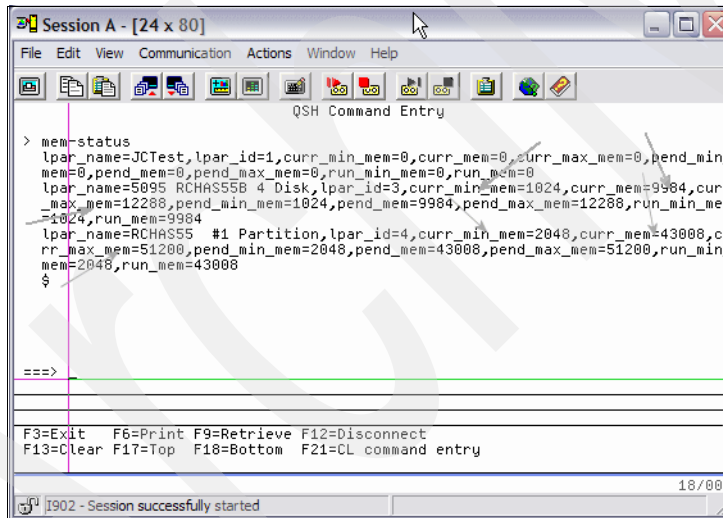


Figure 11-20 mem-status

Figure 11-20 above shows the details of partition and the same is explained in Table 11-1.

Table 11-1 Memory allocation

Partition Name	Memory - Min	Memory - Curr	Memory - Max
JCTest	0	0	0
5095 RCHAS55B 4 Disk	1024	9984	12288
RCHAS55 #1 Partition	2048	43008	51200

- ▶ Move the memory from the partition (5095 RCHAS55B 4 Disk) to partition (RCHAS55 #1 Partition) by executing the script `mem-move` as shown in Figure 11-22.

This script (Example 11-3) moves the memory of size 1024 Mb from partition id 3 (5095 RCHAS55B 4 Disk) to partition id 4 (RCHAS55 #1 Partition).

Example 11-3 mem-move

```
PATH=$PATH:/QOpenSys/usr/bin:/usr/ccs/bin:/QOpenSys/usr/bin/X11:/usr/sbin:./usr/bin
```

```
ssh -T 9.5.92.92 chhwres -r mem -m Server-9406-550-SN10F17AD -o m --id 3 --tid 4 -q 1024
```

Figure 11-21 shows the script.

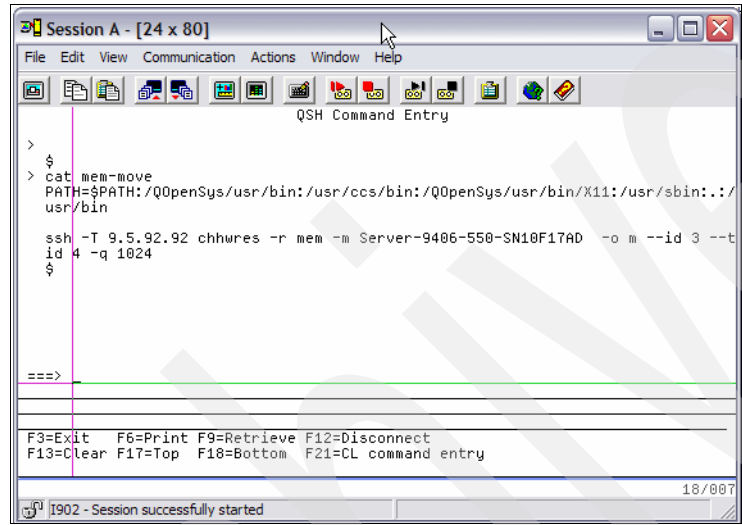


Figure 11-21 mem-move -script

Figure 11-22 shows the script executed.

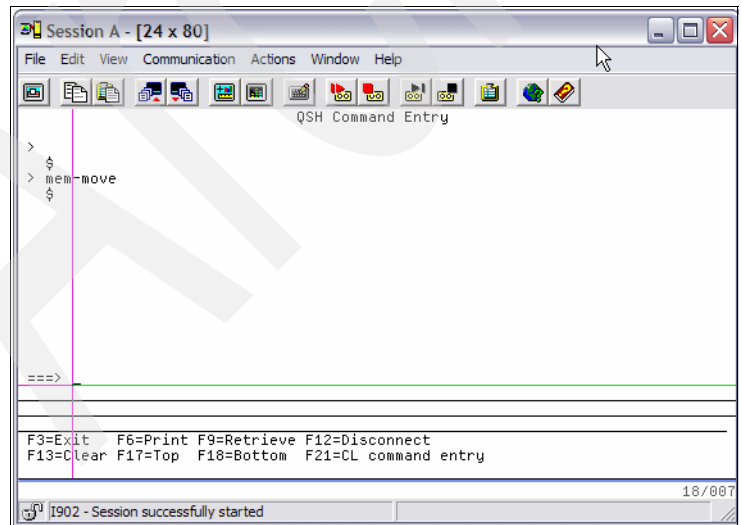


Figure 11-22 mem-move-script executed

Figure 11-21 shows the results.

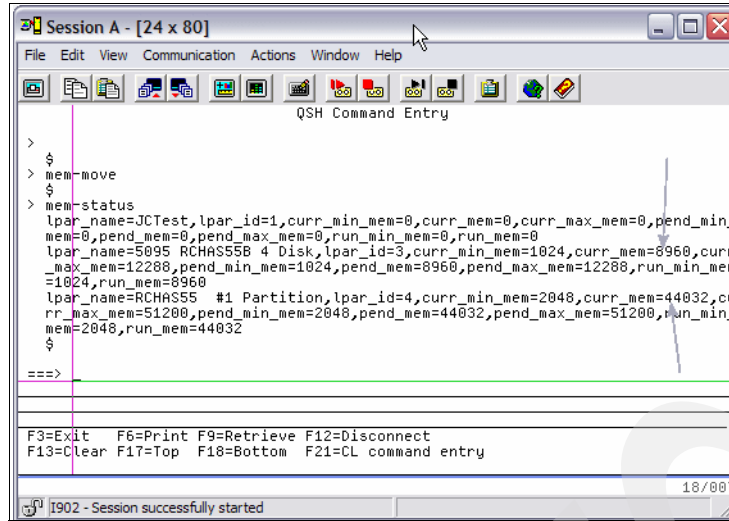


Figure 11-23 mem-status-after script execution

Table 11-2 Memory allocation after the memory movement

Partition Name	Memory - Min	Memory - Curr	Memory - Max
JCTest	0	0	0
5095 RCHAS55B 4 Disk	1024	8960	12288
RCHAS55 #1 Partition	2048	51200	51200

11.5 Scheduling the DLPAR function from i5/OS

Scripts can be scheduled to run from the i5/OS using the job scheduler. Scripts can be written from the desktop, and using the operation navigator, drag and drop them to the desired IFS directory (for example, /home/fred) in i5/OS.

Scheduling the memory movement

You can do the scheduling of memory movement using the following procedure:

- ▶ Logon to i5/OS using the user name *fred* and from the main menu, run the command `wrkjobscode` and press Enter, then press F6 to add the entry (Figure 11-24).

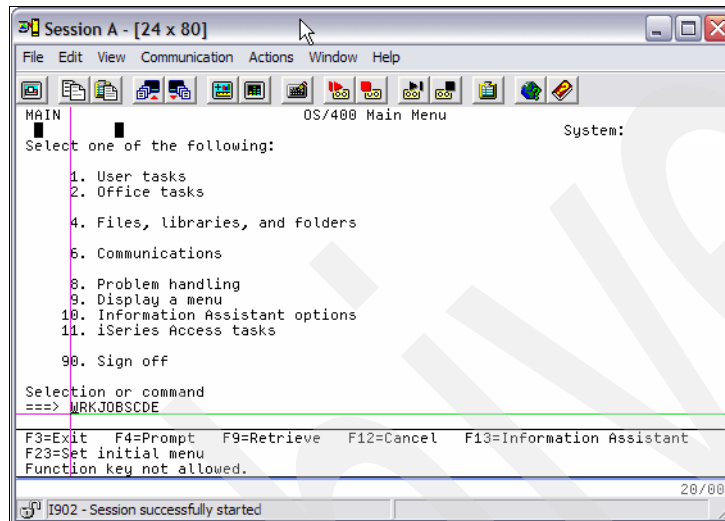


Figure 11-24 `wrkjobscode`

- ▶ Enter the job name in the *Job name* field (for example, LPARTEST).

Enter the `qsh` command in the *Command to run* field as shown in Figure 11-25 and press Enter.

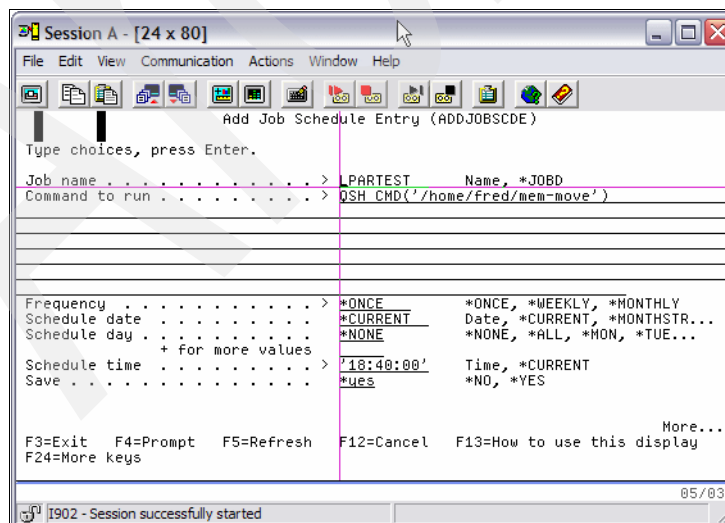


Figure 11-25 `addjobscode`

Note: The script shown in Figure 11-21 is scheduled here.

Figure 11-26 shows the added job schedule entry.

```
Session A - [24 x 80]
File Edit View Communication Actions Window Help
Work with Job Schedule Entries 05/20/05 18:36:04

Type options, press Enter.
2=Change 3=Hold 4=Remove 5=Display details 6=Release
8=Work with last submission 10=Submit immediately

-----Schedule-----
Opt Job Status Date Time Frequency Recovery Action Next Submit
Date
--- LPARTEST SCD 05/20/05 18:40:00 *ONCE *SBMRLS 05/20/05
--- PERFNAYDR HLD *ALL 00:30:00 *WEEKLY *SBMRLS 05/21/05
--- PERFNAYDS HLD *SUN 03:30:00 *WEEKLY *SBMRLS 05/22/05
--- PERFNAYPG HLD *MONTHSTR 02:30:00 *MONTHLY *SBMRLS 06/01/05
--- PERFNAYPM HLD *ALL 19:30:00 *WEEKLY *SBMRLS 05/20/05
--- QSHJHEARTBT SCD *ALL 00:11:00 *WEEKLY *NOSBM 05/21/05
--- QS9SACOL SCD *ALL 00:41:00 *WEEKLY *NOSBM 05/21/05
More...

Parameters or command
===>
F3=Exit F4=Prompt F5=Refresh F6=Add F9=Retrieve
F11=Display job queue data F12=Cancel F17=Top F18=Bottom

12/002
I902 - Session successfully started
```

Figure 11-26 *addjobscde-added*

The history log shows the memory size changes after the scheduled operation is completed (Figure 11-27).

```
Session A - [24 x 80]
File Edit View Communication Actions Window Help
Display History Log Contents

Job 006170/FRED/LPARTEST started on 05/20/05 at 18:40:00 in subsystem QBATCH
Job 006170/FRED/LPARTEST submitted for job schedule entry LPARTEST number 000
Job 006170/FRED/LPARTEST ended on 05/20/05 at 18:40:36; 1 seconds used; end c
Main storage size change to 44032.00M in progress.
Main storage size changed to 44032.00M.
Job 006191/FRED/QZSHS started on 05/20/05 at 18:40:45 in subsystem QINTER in
Job 006191/FRED/QZSHS ended on 05/20/05 at 18:48:42; 1 seconds used; end cod

Press Enter to continue. Bottom

F3=Exit F10=Display all F12=Cancel

01/001
I902 - Session successfully started
```

Figure 11-27 *History log (System i5 #1 Partition)*

11.5.1 Scheduling the i/o movement

In this section we discuss how to schedule i/o movement.

Adding the i/o to the partition

The following procedure shows the how to schedule the i/o removal, adding, and movement. From Figure 11-28 below, we can see two systems. In this section, we are using the system *FelixComplex*.

Refer to Chapter 11, “OpenSSH” on page 341 for information on creating user ids in the HMC and i5/OS as well as the `ssh` authentication procedure.

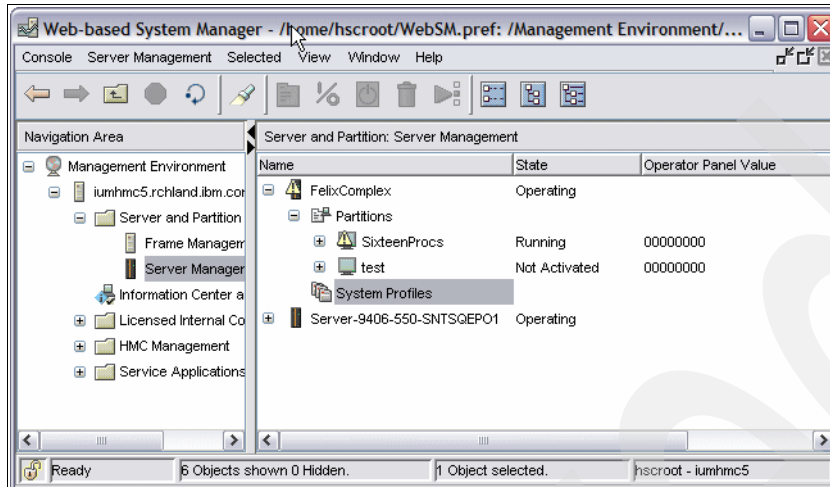


Figure 11-28 FelixComplex System

- ▶ Logon to i5/OS using the user name *fred* and enter the command `QSH` from the main menu command line to enter into the QSHELL environment.
- ▶ Create a script as shown in Example 11-4 from the windows desktop, and then through the operation navigator, drag and drop the file into the `/home/fred` directory.

Example 11-4 lpar-details

```
PATH=$PATH:/QopenSys/usr/bin:/usr/ccs/bin:/QopenSys/usr/bin/X11:/usr/sbin:./
#####
ssh -T 9.5.17.228 lssyscfg -r lpar -m FelixComplex$
```

- ▶ To see the LPAR details, run the script `lpar-details` as shown in Figure 11-29.




Figure 11-29 lpar-details executed

- ▶ To check the i/o resources of the system (FelixComplex) shown in Figure 11-28 above, run the script `iodetails` as shown below in Figure 11-30.

You can write a script as shown in Example 11-5 to see the iodetails of the partition.

Example 11-5 iodetails-script

```
PATH=$PATH:/QOpenSys/usr/bin:/usr/ccs/bin:/QOpenSys/usr/bin/X11:/usr/sbin:./usr/bin
#####
ssh -T 9.5.17.228 lshwres -r io --rsubtype slot -m FelixComplex -F
phys_loc,description,lpar_name,drc_index,bus_id --header
```

Figure 11-30 shows the script.

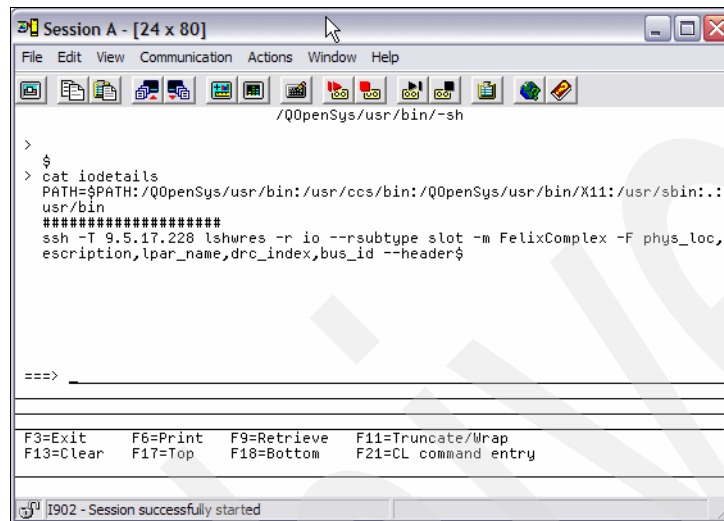


Figure 11-30 iodetails-script

The results are shown in Figure 11-31 on page 358 through Figure 11-36 on page 361.



Figure 11-31 iodetails-executed

QSH Command Entry

```
C07,Empty slot,null,2103000F,15
C08,Empty slot,null,2104000F,15
C09,Empty slot,null,2105000F,15
C11,I/O Processor,null,21010013,19
C12,PCI Ultra4 SCSI Disk Controller,null,21020013,19
C13,Empty slot,null,21030013,19
C14,Empty slot,null,21040013,19
C15,Empty slot,null,21050013,19
C01,I/O Processor,null,21010014,20
C02,PCI Ultra4 SCSI Disk Controller,null,21020014,20
C03,Empty slot,null,21030014,20
C04,Empty slot,null,21040014,20
C05,I/O Processor,null,21010015,21
C06,PCI Ultra4 SCSI Disk Controller,null,21020015,21
```

==>

```
F3=Exit F6=Print F9=Retrieve F12=Disconnect
F13=Clear F17=Top F18=Bottom F21=CL command entry
```

Figure 11-32 iodetails - continued 1

QSH Command Entry

```
C07,Empty slot,null,21030015,21
C08,Empty slot,null,21040015,21
C09,Empty slot,null,21050015,21
C11,I/O Processor,null,21010010,16
C12,PCI Ultra4 SCSI Disk Controller,null,21020010,16
C13,Empty slot,null,21030010,16
C14,Empty slot,null,21040010,16
C15,Empty slot,null,21050010,16
C01,Empty slot,null,21010011,17
C02,Empty slot,null,21020011,17
C03,Empty slot,null,21030011,17
C04,Empty slot,null,21040011,17
C05,I/O Processor,null,21010012,18
C06,PCI Ultra4 SCSI Disk Controller,null,21020012,18
```

==>

```
F3=Exit F6=Print F9=Retrieve F12=Disconnect
F13=Clear F17=Top F18=Bottom F21=CL command entry
```

Figure 11-33 iodetails - continued 2

QSH Command Entry

```
C07,Empty slot,null,21030012,18
C08,Empty slot,null,21040012,18
C09,Empty slot,null,21050012,18
C11,I/O Processor,null,21010016,22
C12,PCI Ultra4 SCSI Disk Controller,null,21020016,22
C13,Empty slot,null,21030016,22
C14,Empty slot,null,21040016,22
C15,Empty slot,null,21050016,22
C01,Empty slot,null,21010017,23
C02,Empty slot,null,21020017,23
C03,Empty slot,null,21030017,23
C04,Empty slot,null,21040017,23
C05,I/O Processor,null,21010018,24
C06,PCI Ultra4 SCSI Disk Controller,null,21020018,24
```

==>

```
F3=Exit F6=Print F9=Retrieve F12=Disconnect
F13=Clear F17=Top F18=Bottom F21=CL command entry
```

Figure 11-34 iodetails - continued 3

QSH Command Entry

```
C07,PCI Ultra4 SCSI Disk Controller,null,21030018,24
C08,Empty slot,null,21040018,24
C09,Empty slot,null,21050018,24
C11,PCI I/O Processor,SixteenProcs,2101000A,10
C12,PCI Ultra4 SCSI Disk Controller,SixteenProcs,2102000A,10
C13,Empty slot,SixteenProcs,2103000A,10
C14,Empty slot,SixteenProcs,2104000A,10
C15,Empty slot,SixteenProcs,2105000A,10
C01,PCI I/O Processor,SixteenProcs,2101000B,11
C02,Empty slot,SixteenProcs,2102000B,11
C03,PCI Ultra4 SCSI Disk Controller,SixteenProcs,2103000B,11
C04,PCI 10/100/1000Mbps Ethernet UTP 2-port,SixteenProcs,2104000B,11
C05,PCI I/O Processor,SixteenProcs,2101000C,12
C06,PCI Ultra4 SCSI Disk Controller,SixteenProcs,2102000C,12
```

==>

```
F3=Exit F6=Print F9=Retrieve F12=Disconnect
F13=Clear F17=Top F18=Bottom F21=CL command entry
```

Figure 11-35 iodetails - continued 4

```

QSH Command Entry

C07,SCSI bus controller,SixteenProcs,2103000C,12
C08,Empty slot,SixteenProcs,2104000C,12
C09,Empty slot,SixteenProcs,2105000C,12
$

===>
F3=Exit  F6=Print  F9=Retrieve  F12=Disconnect
F13=Clear  F17=Top  F18=Bottom  F21=CL command entry

```

Figure 11-36 *iodetails - continued 5*

Figure 11-31 through Figure 11-36 have shown the output of the `iodetails` script. In this output, IOP and IOA at location C11 and C12 at BUS ID13 are not allocated to any partition (that is, `lpar_name` is null).

- To add this IOP and IOA to the partition, run the script `io-add` as shown in Example 11-6. This will add the IOP and IOA (C11 and C12) to the *SixteenProcs* partition.

You can run this script `io-add` from the QSHELL environment interactively or you can schedule to run the script from the i5/OS. The following procedure shows how to schedule this script from i5/OS.

Example 11-6 io-add-script

```

PATH=$PATH:/QOpenSys/usr/bin:/usr/ccs/bin:/QOpenSys/usr/bin/X11:/usr/sbin:./usr/bin
#####
ssh -T 9.5.17.228 chhwres -r io -m FelixComplex -o a --id 1 -l 2101000D&
#####
ssh -T 9.5.17.228 chhwres -r io -m FelixComplex -o a --id 1 -l 2102000D

```

- ▶ From the i5/OS main menu, run the command `wrkjobsjde` and press Enter, then press F6 to add the entry (Figure 11-37).

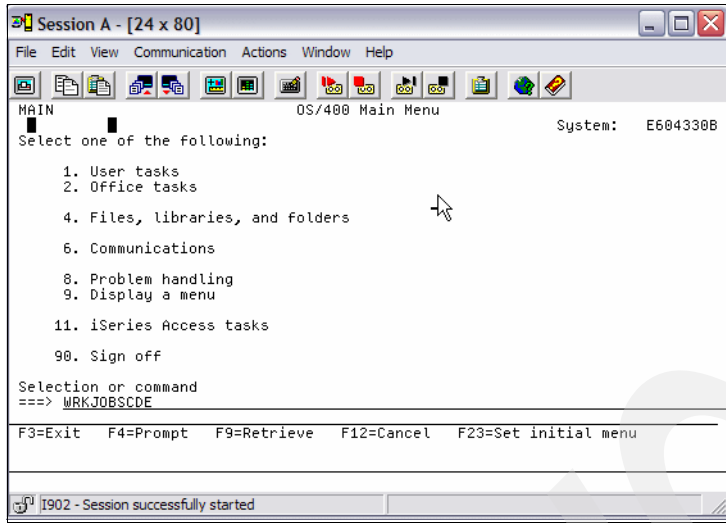


Figure 11-37 `wrkjobsjde -io-add`

- ▶ Enter the job name in the **Job name** field (for example, IOADD).
- ▶ Enter the `qsh` command in the **Command to run** field as shown in Figure 11-38 and press Enter.

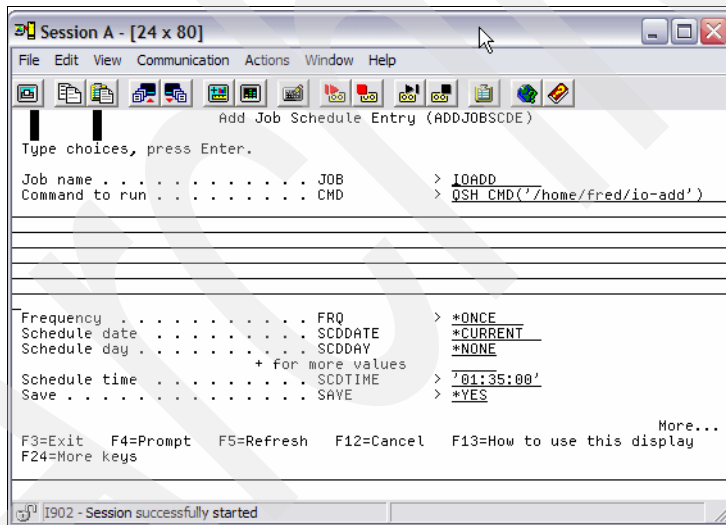


Figure 11-38 `io-add-scheduled`

Figure 11-39 shows the added job schedule entry.

```

Session A - [24 x 80]
File Edit View Communication Actions Window Help
Work with Job Schedule Entries E604330B 05/24/05 01:29:06

Type options, press Enter.
2=Change 3=Hold 4=Remove 5=Display details 6=Release
8=Work with last submission 10=Submit immediately

-----Schedule-----
Opt Job      Status Date      Time      Frequency  Recovery Action  Next
-----
  -- IOADD     SCD    05/24/05  01:35:00  *ONCE     *SBMRLS  05/24/05
  -- QS9AUTOPTF SCD    *FRI     22:30:00  *WEEKLY   *SBMRLS  05/27/05
  -- QS9AUTOTST SCD    *FRI     03:00:00  *WEEKLY   *SBMRLS  05/27/05

Parameters or command
==>
F3=Exit  F4=Prompt  F5=Refresh  F6=Add  F9=Retrieve
F11=Display job queue data  F12=Cancel  F17=Top  F18=Bottom

I902 - Session successfully started
  
```

Figure 11-39 io-add-scheduled-entry

Note: To see the command syntax, logon to the HMC, and from the command line, type `chhwres --help` and press Enter.

Once the scheduled activity is completed, you can check the history log (which indicates the completion of the job) as shown in Figure 11-40.

```

Session A - [24 x 80]
File Edit View Communication Actions Window Help
Display History Log Contents

Job 005879/FRED/IOADD started on 05/24/05 at 01:35:00 in subsystem QBATCH in
Job 005879/FRED/IOADD submitted for job schedule entry IOADD number 000003.
Job 005879/FRED/IOADD ended on 05/24/05 at 01:36:01; 1 seconds used; end code

Press Enter to continue.
F3=Exit  F10=Display all  F12=Cancel

I902 - Session successfully started
  
```

Figure 11-40 io-add histlog

- ▶ Run the script `iodetails` from the QSHHELL to see the resource status as shown in Figure 11-41.

```

QSH Command Entry

>
$
> iodetails
phys_loc,description,lpar_name,drc_index,bus_id
C11,PCI Mag Media Controller,SixteenProcs,2102000D,13
C12,PCI Ultra4 SCSI Disk Controller,SixteenProcs,2102000D,13
C13,Empty slot,null,2103000D,13
C14,Empty slot,null,2104000D,13
C15,Empty slot,null,2105000D,13
C01,Empty slot,null,2101000E,14
C02,Empty slot,null,2102000E,14
C03,Empty slot,null,2103000E,14
C04,Empty slot,null,2104000E,14
C05,I/O Processor,null,2101000F,15

===>
F3=Exit   F6=Print  F9=Retrieve F12=Disconnect
F13=Clear F17=Top   F18=Bottom  F21=CL command entry

```

Figure 11-41 `iodetails` -after adding the i/o

In Figure 11-41 above, the Slots C11 and C12 are assigned to the partition **SixteenProcs**.

Note: This is the end of the procedure for i/o add.

Removing the i/o from the partition

As shown in Figure 11-41, the IOP and IOA at location C11 and C12 at BUS ID13 is now allocated to partition **SixteenProcs**.

To remove this IOP and IOA from the partition, either you can run the script `io-remove` as in Example 11-7 from the QSHHELL, or you can schedule the job as shown in Figure 11-43, Figure 11-44, and Figure 11-45.

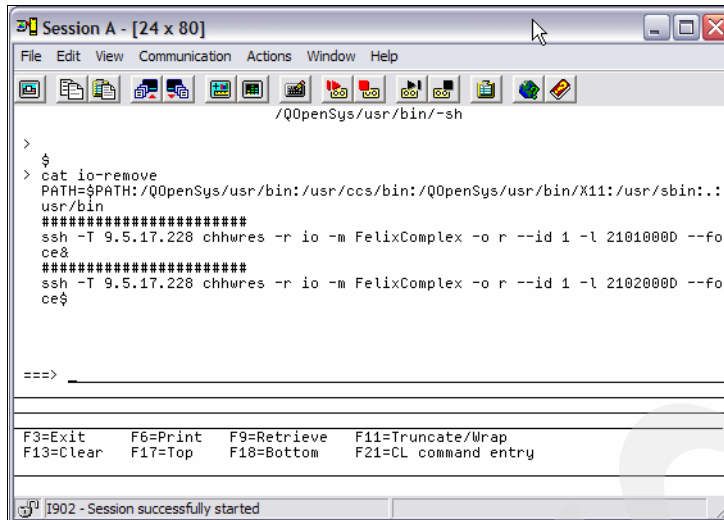
Example 11-7 `io-remove` script

```

PATH=$PATH:/QOpenSys/usr/bin:/usr/ccs/bin:/QOpenSys/usr/bin/X11:/usr/sbin:./usr/bin
#####
ssh -T 9.5.17.228 chwres -r io -m FelixComplex -o r --id 1 -l 2101000D --force&
#####
ssh -T 9.5.17.228 chwres -r io -m FelixComplex -o r --id 1 -l 2102000D --force

```

The removal script is shown in Figure 11-42.



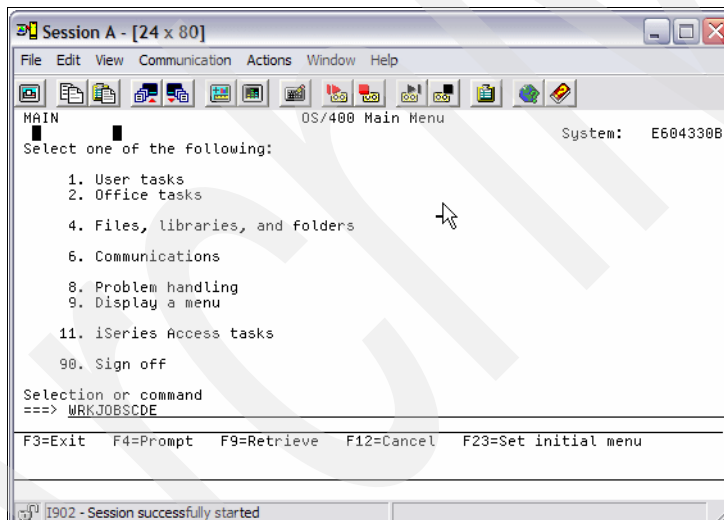
```
Session A - [24 x 80]
File Edit View Communication Actions Window Help
/QOpenSys/usr/bin/~sh
>
>
> cat io-remove
PATH=$PATH:/QOpenSys/usr/bin:/usr/ccs/bin:/QOpenSys/usr/bin/X11:/usr/sbin:./usr/bin
#####
ssh -T 9.5.17.228 chhwres -r io -m FelixComplex -o r --id 1 -l 21010000 --for ce8
#####
ssh -T 9.5.17.228 chhwres -r io -m FelixComplex -o r --id 1 -l 21020000 --for ce8
====>

F3=Exit   F6=Print   F9=Retrieve F11=Truncate/Wrap
F13=Clear F17=Top    F18=Bottom F21=CL command entry

I902 - Session successfully started
```

Figure 11-42 `io-remove` script

- ▶ From the i5/OS main menu, run the command `wrkjobsdcde` and press Enter, then press F6 to add the entry as shown in Figure 11-43.



```
Session A - [24 x 80]
File Edit View Communication Actions Window Help
MAIN OS/400 Main Menu System: E604330B
Select one of the following:
  1. User tasks
  2. Office tasks
  4. Files, libraries, and folders
  6. Communications
  8. Problem handling
  9. Display a menu
 11. iSeries Access tasks
 90. Sign off
Selection or command
====> WRKJOBSCDE

F3=Exit   F4=Prompt   F9=Retrieve   F12=Cancel   F23=Set initial menu

I902 - Session successfully started
```

Figure 11-43 `wrkjobsdcde -io-remove`

The removal script is shown in Figure 11-44.

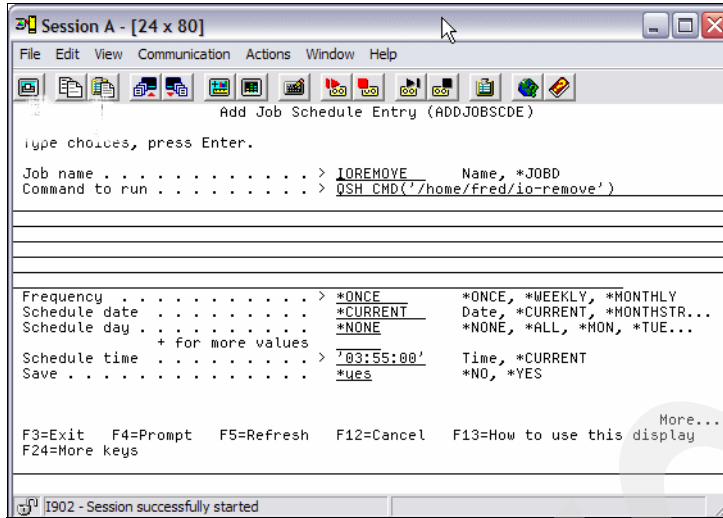


Figure 11-44 io-remove - scheduled

Figure 11-45 shows the added job schedule entry.

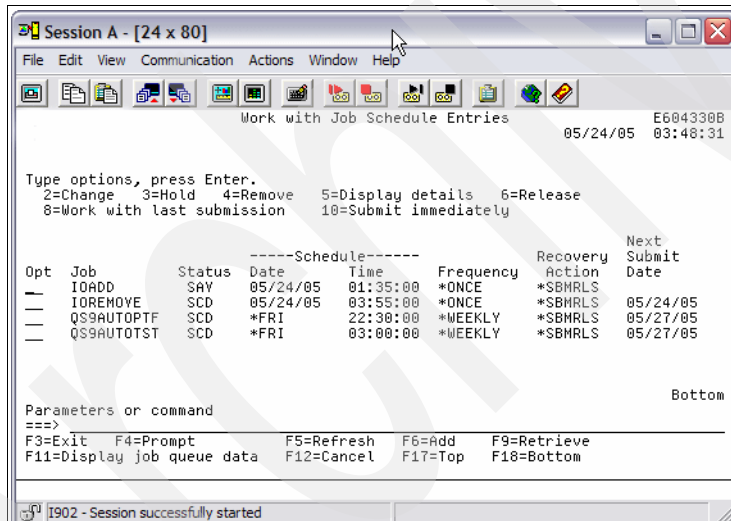


Figure 11-45 io-remove - scheduled entry

Once the scheduled activity is completed, you can check the history log (which shows the completion of the job) as shown in Figure 11-46.

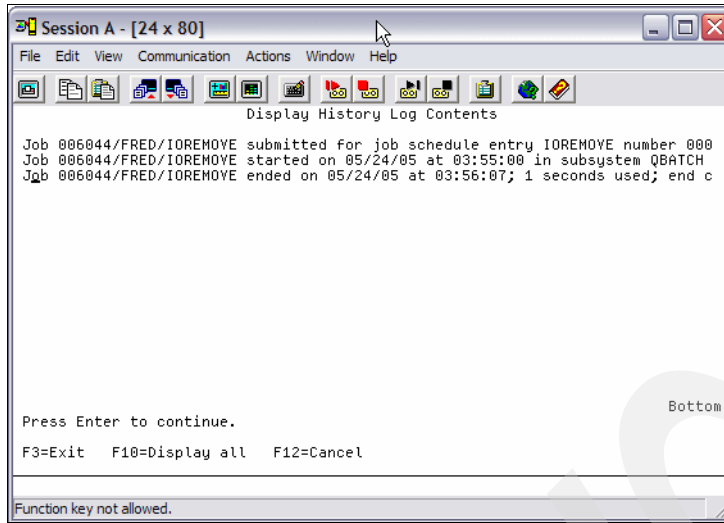


Figure 11-46 *io-remove histlog*

Figure 11-47 shows the `iodetails` after removing the `i/o`.

```
QSH Command Entry
>
$
> iodetails
phys_loc,description,lpar_name,drc_index,bus_id
C11,PCI I/O Processor,null,2101000D,13
C12,PCI Mag Media Controller,null,2102000D,13
C13,Empty slot,null,2103000D,13
C14,Empty slot,null,2104000D,13
C15,Empty slot,null,2105000D,13
C01,Empty slot,null,2101000E,14
C02,Empty slot,null,2102000E,14
C03,Empty slot,null,2103000E,14
C04,Empty slot,null,2104000E,14
C05,I/O Processor,null,2101000F,15

===>
F3=Exit F6=Print F9=Retrieve F12=Disconnect
F13=Clear F17=Top F18=Bottom F21=CL command entry
```

Figure 11-47 *iodetails - after removing the i/o*

Note: This is the end of the procedure for `i/o` removal.

Moving the i/o from one partition to another

To move the i/o from one partition to another, we need to know the partition ids and the drc index of the particular IOP or IOA.

You can see the partition details in Figure 11-29 on page 357.

To move the IOP (C11) and IOA(C12) from the *SixteenProc* partition to the *test* partition, run the script as shown in Example 11-8.

Example 11-3 shows the details taken from Figure 11-29 on page 357 and Figure 11-41 on page 364.

Table 11-3 i/o details

Partition Name	LPAR Id	IOP drc_index	IOA drc_index
SixteenProcs	1	2101000D	2102000D
Test	2		

Example 11-8 shows the script to do this function.

Example 11-8 io-move -script

```
PATH=$PATH:/QOpenSys/usr/bin:/usr/ccs/bin:/QOpenSys/usr/bin/X11:/usr/sbin:./usr/bin
#####
ssh -T 9.5.17.228 chhwres -r io -m FelixComplex -o m --id 1 -l 2101000D --tid 2
#####
ssh -T 9.5.17.228 chhwres -r io -m FelixComplex -o m --id 1 -l 2102000D --tid 2
```

Note: To see the command syntax, logon to the HMC, and from the command line, type: **chhwres --help** and press Enter.

11.6 Scheduling the DLPAR function from Windows

HMC can be accessed from many other operating systems with the `ssh` client utility program installed. One of the popular utilities is PuTTY, which is a free `ssh` client utility program that can be downloaded from the Internet.

More information about PuTTY can be found in:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/>

You can download PuTTY and related utilities from the URL:

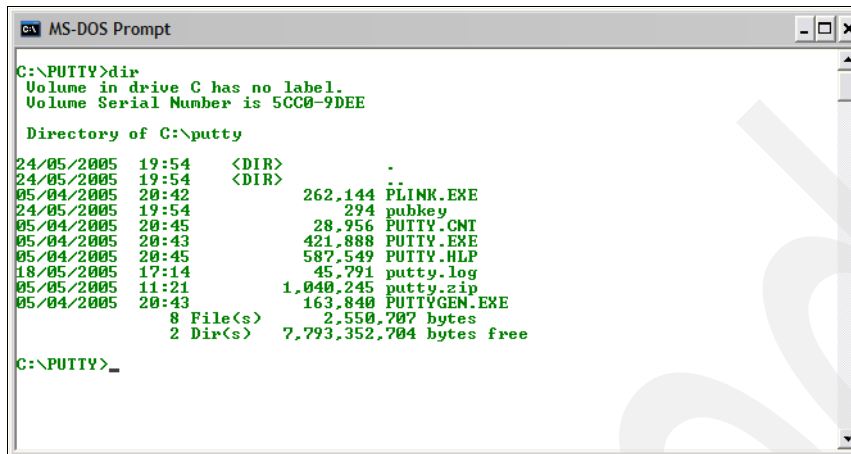
<http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>

The following procedure shows how use the PuTTY utility program to communicate to HMC from the Windows client.

From the Windows client, create the folder:

```
c:\putty
```

Download the PuTTY utility programs (`putty.exe` and `plink.exe`) from the Internet to the folder `c:\putty` as shown in Figure 11-48.



```
MS-DOS Prompt
C:\PUTTY>dir
Volume in drive C has no label.
Volume Serial Number is 5CC0-9DEE

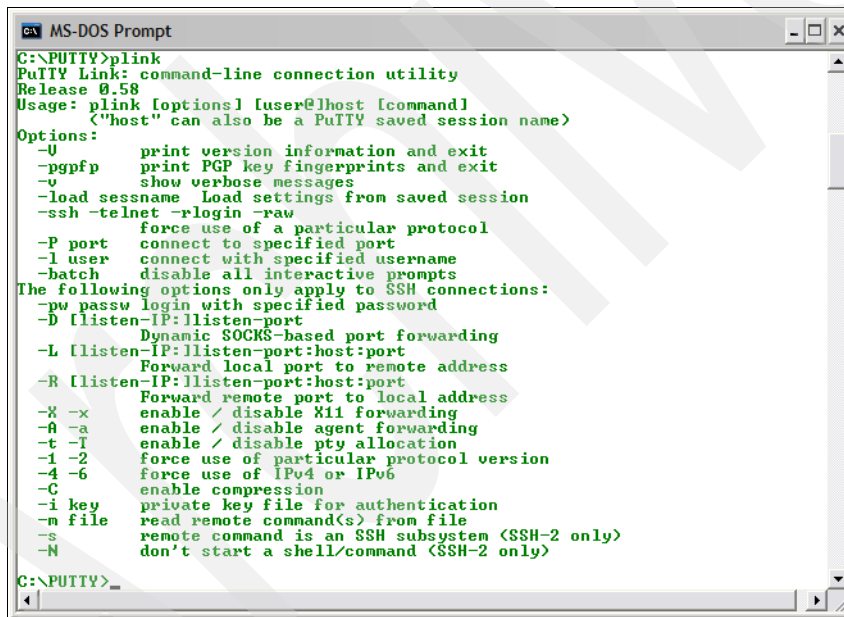
Directory of C:\putty

24/05/2005  19:54  <DIR>          .
24/05/2005  19:54  <DIR>          ..
05/04/2005  20:42             262,144  PLINK.EXE
24/05/2005  19:54             294     pubkey
05/04/2005  20:45             28,956  PUTTY.CMT
05/04/2005  20:43             421,888  PUTTY.EXE
05/04/2005  20:45             587,549  PUTTY.HLP
18/05/2005  17:14             45,791  putty.log
05/05/2005  11:21            1,040,245  putty.zip
05/04/2005  20:43            163,840  PUTTYGEN.EXE
           8 File(s)          2,550,707 bytes
           2 Dir(s)          7,793,352,704 bytes free

C:\PUTTY>
```

Figure 11-48 `putty` - Folder

Type `plink` and press Enter from `c:\putty` to see the command syntax as shown in Figure 11-49.



```
MS-DOS Prompt
C:\PUTTY>plink
PuTTY Link: command-line connection utility
Release 0.58
Usage: plink [options] [user@]host [command]
        <"host" can also be a PuTTY saved session name>

Options:
-U          print version information and exit
-pgpfp     print PGP key fingerprints and exit
-v         show verbose messages
-load sessname Load settings from saved session
-ssh -telnet -rlogin -raw
-P port    force use of a particular protocol
-p port    connect to specified port
-l user    connect with specified username
-batch     disable all interactive prompts
The following options only apply to SSH connections:
-pw passwd login with specified password
-D [listen-IP:]listen-port
           Dynamic SOCKS-based port forwarding
-L [listen-IP:]listen-port:host:port
           Forward local port to host:port
-R [listen-IP:]listen-port:host:port
           Forward remote port to local address
-K -x     enable / disable X11 forwarding
-A -a     enable / disable agent forwarding
-t -T     enable / disable pty allocation
-i -2     force use of particular protocol version
-4 -6     force use of IPv4 or IPv6
-C        enable compression
-i key    private key file for authentication
-m file   read remote command(s) from file
-s        remote command is an SSH subsystem <SSH-2 only>
-N        don't start a shell/command <SSH-2 only>

C:\PUTTY>
```

Figure 11-49 `Plink` command syntax

Archived



Using Work Management to influence System i5 resources

This chapter describes the new options added to OS/400 V5R3 Work Management to influence System i5 performance for specific workloads. These new options are processor and memory affinity on some multi-processors models of System i5 systems.

12.1 Main storage and processor affinity concept

In some environments and system configurations, processes and threads can achieve improved affinity for memory and processor resources. This improved level of affinity can result in improved performance. You can tune the main storage affinity level setting on your server at the process level and at the thread level.

Important: Tuning main storage affinity levels may improve performance in some environments or system configurations, or degrade it in others.

12.1.1 Processor affinity concept

Other things being equal, it is desirable to dispatch a thread on the processor it last used. This dispatching criterion is called processor affinity. The level of emphasis on processor affinity can vary.

The highest possible degree of processor affinity is to bind a thread to a specific processor. Binding means that the thread will be dispatched to that processor only, regardless of the availability of other processors.

This technique can be useful for CPU-intensive programs that experience few interrupts. It can sometimes be counterproductive for ordinary programs, because it may delay the redispach of a thread after an I/O until the processor to which the thread is bound becomes available. If the thread has been blocked for the duration of an I/O operation, it is unlikely that much of its processing context remains in the caches of the processor to which it is bound. It would probably be better served if it were dispatched to the next available processor.

12.1.2 Memory affinity concept

In this section, we present a basic view of memory affinity.

Each Multi-Chip Module (MCM) contains 8 processors, making a processor node that is connected to memory cards. As long as the processors access these local memory cards, the performance is optimum.

Now, with multiple MCMs, a processor on one node may need to access memory cards connected to another processor node (another MCM). So, instead of using a local memory access, the bus interconnecting the MCMs is used to perform a remote memory access, and the performance is not optimum. See Figure 12-1.

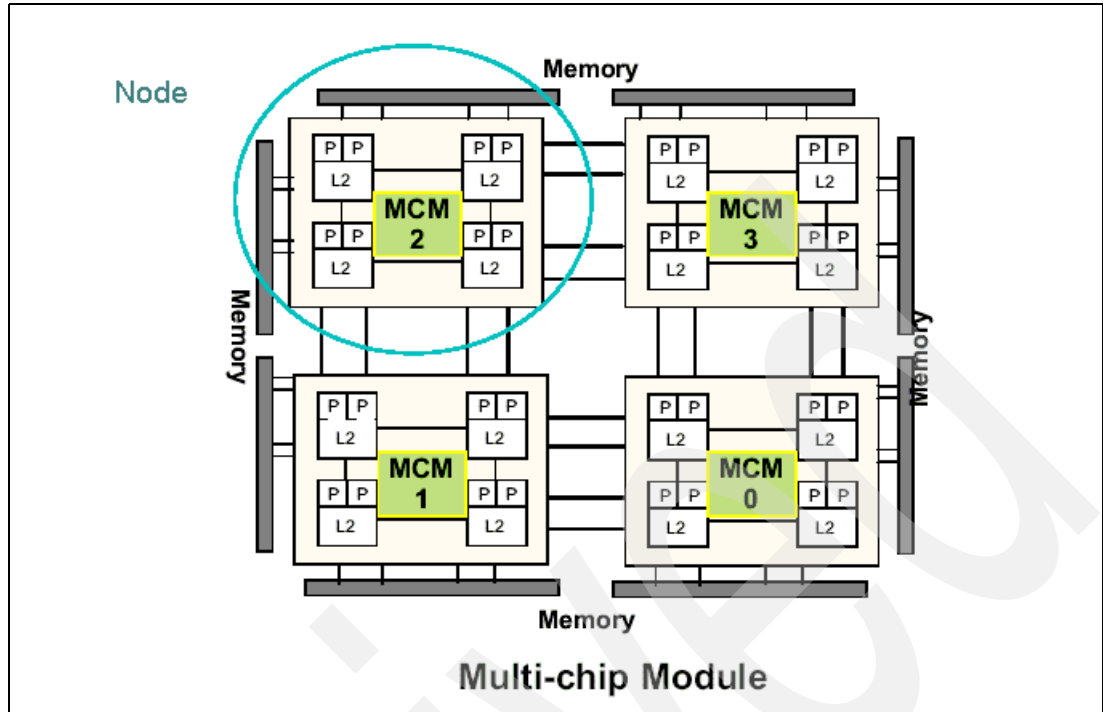


Figure 12-1 Processors and memory layout for n-way PowerPC MCMs

The memory affinity support recognizes the relationship between processors, memory, and multichip modules (MCMs) in SMP machines such as the IBM @server System i5. The support provides improved performance to some high performance computing applications.

Memory affinity is a special purpose option for improving performance on IBM System i5 machines that contain multiple multichip modules. But memory affinity is neither recommended nor beneficial for general use. Most systems run better with memory affinity disabled. IBM testing has discovered that memory affinity can have an unpredictable effect on system and application performance. Some applications benefit from memory affinity, others can degrade.

12.2 Work Management enhancements for memory and processor affinity

The following performance system values are new for i5/OS V5R3:

- ▶ QTHDRSCAFN — Thread affinity
- ▶ QTHDRSCADJ — Automatically adjust thread resources

The following parameters are new for the Add Routing Entry (ADDRTGE) command:

- ▶ THDRSCAFN — Thread resources affinity
- ▶ RSCAFNGRP — Resources affinity group

12.2.1 QTHDRSCAFN (thread affinity)

This specifies whether secondary threads will have affinity to the same group of processors and memory as the initial thread or not. It also specifies the degree to which the system tries to maintain the affinity between threads and the subset of system resources they are assigned.

A change made to this system value takes effect immediately for all jobs that become active after the change, but only if they retrieve their affinity values from the system value. Jobs and threads that are currently active will continue to run with the affinity values in which they were started. New jobs and threads that get their affinity values from jobs that are currently active, (for example, batch immediate jobs or secondary threads that inherit attributes from the initial thread) will continue to use the values stored in the initiating job or thread, rather than the current system value.

In addition, you can use the *automatically adjust thread resources* (QTHDRSCADJ) system value to distribute the use of system resources.

Thread affinity

This specifies whether or not secondary threads are grouped with the initial thread. Select one of the following choices:

- ▶ **Secondary threads use same processors and memory as initiating thread (*GROUP)**
This indicates that secondary threads have a preference for the same group of processors and memory as the initial thread. Multi-threaded workloads that share data between the threads within a job may perform better when using this option.
- ▶ **Secondary threads do not use same processors and memory as initiating thread (*NOGROUP)**
This indicates that secondary threads do not have a preference for the same group of processors and memory as the initial thread. Workloads that use the full processing power of a system may perform better when using this option.
- ▶ **Level of affinity**
This specifies the degree to which the threads try to use the subset of system resources in which they have a preference. Select one of the following choices:
 - Threads will only use resources they have affinity to (*HIGH).
 - Threads will only use the subset of system resources they have affinity to, and will wait until they become available.
- ▶ **Threads use any resource (*NORMAL)**
Threads will use any processor or memory in the system if the resources they have affinity to are not readily available.

12.2.2 QTHDRSCADJ (automatically adjust thread resources)

This specifies whether the system should dynamically make adjustments to the affinity of threads currently running on the system (1) or not (0). If some resources are being used more than others, the system may reassign some of the threads running on the more heavily used resources to have affinity to the less used resources.

If you have specified to group secondary threads using the thread affinity (QTHDRSCAFN) system value, the threads within one process will all be moved as a group. If jobs have been logically grouped with the routing entry or prestart job entry, the entire group will be moved together. If this option is selected, the system will attempt to distribute work evenly across all the system's resources.

12.2.3 ADDRTGE command — new parameters

In the following sections we provide a description of the new parameters.

Thread resources affinity (THDRSCAFN)

This specifies the affinity of threads to system resources.

► Element 1: Group, single values

- ***SYSVAL**

When a job is started using this routing entry, the thread resources affinity value from the QTHDRSCAFN system value will be used.

- ***NOGROUP**

Jobs using this routing entry will have affinity to a group of processors and memory. Secondary threads running under the job will not necessarily have affinity to the same group of processors and memory.

- ***GROUP**

Jobs using this routing entry will have affinity to a group of processors and memory. Secondary threads running under the job will all have affinity to the same group of processors and memory as the initial thread.

► Element 2: Level

- ***NORMAL**

A thread will use any processor or memory if the resources it has affinity to are not readily available.

- ***HIGH**

A thread will only use the resources it has affinity to, and will wait until they become available if necessary.

Resources affinity group (RSCAFNGRP)

This specifies whether or not jobs using this routing entry will be grouped together having affinity to the same system resources (processors and memory). A value of *YES for this parameter will take precedence over the QTHDRSCAFN system value when set to *NOGROUP.

► ***NO**

Jobs that use this routing entry will not be grouped together.

► ***YES**

Jobs that use this routing entry will be grouped together such that they will have affinity to the same system resources. Jobs that share data in memory may perform better if they have affinity to the same resources.

Archived



Virtual Partition Manager

In this chapter we discuss the following topics:

- ▶ Introduction to Virtual Partition Manager for eServer System i5 systems
- ▶ Planning for Virtual Partition Manager
- ▶ Getting started with Virtual Partition Manager
- ▶ Preparing your system for Virtual Partition Manager
- ▶ Creating Linux partitions using Virtual Partition Manager
- ▶ Establishing network connectivity for Linux partitions
- ▶ Setting up i5/OS virtual I/O resources for Linux partitions

13.1 Introduction to Virtual Partition Manager for eServer System i5 systems

With the recently announced System i5 processor based eServer i5 systems, IBM is delivering the 3rd generation of logical partitioning for the iSeries family of servers. The new partitioning capabilities enable customers to further simplify their infrastructures.

The IBM Virtualization Engine™, which provides support for logical partitioning and resource virtualization through i5/OS, is included with eServer System i5 systems.

The Virtual Partition Manager (VPM) introduces the capability to create and manage Linux partitions without the use of the Hardware Management Console (HMC). The Virtual Partition Manager supports the needs of small and medium customers that want to add simple Linux workloads to their eServer System i5 system.

The Virtual Partition Manager supports environments with one i5/OS partition and up to four Linux partitions. In addition, the single i5/OS partition must own and manage all of the I/O resources.

This document is intended to assist IBM Business Partners and clients that are planning to deploy Linux partitions on eServer System i5 systems using the Virtual Partition Manager.

The HMC remains the premier management tool for logical partitioning on the eServer System i5 system. It is required to support more robust configurations: multiple i5/OS partitions, AIX 5L partitions, flexible hardware resource assignments, and dynamic resource movement.

The key differences between the Virtual Partition Manager and the HMC are highlighted in Table 13-1.

Table 13-1 Differences between VPM and HMC

	Virtual Partition Manager	Hardware Management Console
Operating systems supported	i5/OS and Linux	i5/OS, Linux and AIX 5L
Maximum number of partitions	5(1 i5/OS and 4 Linux)	254
Uncapped partition support	Yes	Yes
Dynamic resource movement	No	Yes
I/O support for Linux	Virtual	Virtual and Direct
Maximum # of Virtual Ethernet connections	4	4096
Maximum virtual disk per partition	64 TB	64 TB

This document is intended to assist users that are planning to deploy Linux partitions on eServer System i5 systems using the Virtual Partition Manager.

13.2 Planning for Virtual Partition Manager

Virtual Partition Manager is enabled by enhancing the partition management tasks in the Dedicated Service Tools (DST) and System Service Tools (SST) for i5/OS V5R3. This capability is enabled only for eServer i5 systems, allowing you to create up to a maximum of four Linux partitions in addition to the one i5/OS partition that owns all of the I/O resources for the Linux partitions.

Keep in mind that the intended use of Virtual Partition Manager allows deployment of a limited number of Linux partitions on eServer i5 systems without requiring an HMC, where the entire I/O is hosted by a single i5/OS partition.

13.2.1 Planning considerations

We strongly recommend that you fully understand the planning considerations required to enable Virtual Partition Manager on your eServer i5 system. Following are the planning considerations:

- ▶ You need a minimum of an eServer System i5 system with i5/OS V5R3, and PTF MF34753 (or its supersede) applied to enable Virtual Partition Manager.

Important: Prerequisite PTFs MF34946 and MF34962 must be permanently applied before doing a SAVSYS of the i5/OS partition. A D-mode IPL of the Virtual Partition Manager activated i5/OS partition will fail with an SRC of A6005001 with media that does not have these PTFs permanently applied.

- ▶ An IPL is required to apply MF34753 (or its supersedes).
- ▶ Ensure that you have the latest firmware level (SF230_113 [PTF MH00265] or above) downloaded for the eServer System i5 system.
- ▶ A maximum of four Linux partitions are supported.
- ▶ I/O for all Linux partitions must be managed by a single i5/OS instance.
- ▶ Linux uses virtual I/O to access disk, tape, DVD, and Ethernet resources owned by the i5/OS partition.
- ▶ Up to a maximum of four Virtual Ethernet connections may be configured for each Linux partition or for the i5/OS partition.
- ▶ Direct attach I/O cannot be installed in any Linux partition, including any additional Ethernet Adapters for a Firewall gateway.
- ▶ Only a single i5/OS instance can exist on eServer i5. If you need more than one i5/OS instance on your eServer i5, use a Hardware Management Console (HMC).
- ▶ Linux partition creation and management can only be performed through DST or SST tasks. iSeries Navigator does not support Virtual Partition Manager.
- ▶ Automatic processor balancing between Linux and i5/OS partitions is supported through uncapped processor pool. Dynamic movement of resources such as processor, memory, and I/O is not supported. Restart the Linux partition for the changes to take effect.
- ▶ HMC functionality, such as Capacity Upgrade on Demand (On/Off CoD, and Reserve CoD) is not enabled through DST or SST tasks. Only permanently activated processors may be allocated to any of the partitions.
- ▶ Partition configuration data cannot be saved through DST or SST tasks. Ensure that hardcopy prints are kept with configuration screens should you need to recreate the partitions.

- ▶ Migration of partition configuration data to HMC is not available. If HMC is deployed at a future stage, you need to recreate the Linux partitions. The data stored through virtual I/O on i5/OS remains unaffected.
- ▶ Initially, the Virtual Partition Manager configuration screens are only available in English and are not translated.
- ▶ Virtual Partition Manager cannot be managed through services offerings such as LPAR Toolkit or similar LPAR management tools provided by various IBM business partners. Likewise, there are no Application Programming Interfaces (API) or SSH capabilities.
- ▶ You cannot use Virtual Partition Manager on an eServer System i5 system that is configured using an HMC.

Note: An IPL of eServer i5 system may be required for partition creation.

Once you have carefully understood the above planning considerations, you can start planning for your eServer System i5 system with Linux partitions using Virtual Partition Manager.

13.2.2 Planning for Linux logical partitions

The IBM eServer Hardware Information Center contains important information that you need to understand as you start your preparations for deploying Linux partitions. Begin with the “Partitioning for Linux” section:

http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/index.htm?info/iphae/linuxpartitions.htm

This section provides information about hardware requirements, logical partition planning tasks, and details on how to create a Linux logical partition using i5/OS virtual I/O resources. Some of the important planning steps are documented in the next section.

13.2.3 Perform capacity planning

Sizing guidelines for partitions with Linux workloads can be derived through using the IBM Workload Estimator (WLE).

<http://www-912.ibm.com/supporthome.nsf/document/16533356>

WLE is a tool that runs in a Browser or on a personal computer (PC). It estimates the computer resources required for traditional workloads as core applications and new workloads, as Linux, WebSphere® and Lotus® workloads. For Linux, partition sizing information is available for file serving, Web serving, and network application workloads.

13.2.4 Design and validate your partition configuration

Use the Logical Partition Validation Tool (LVT) to help you design a partitioned system. You can download a copy from the following Web address:

http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/index.htm?info/iphat/iphatlvt.htm

The LVT provides you with a validation report that reflects your system requirements while not exceeding logical partition recommendations. Keep in mind that you can only define hosted Linux partitions using the Virtual Partition Manager. It is also important to understand and plan for the Hypervisor Memory requirements.

13.2.5 Understand use of virtual resources for Linux partitions

To learn how your Linux logical partitions can communicate with other devices through virtual I/O, see the Virtual I/O devices topic in the Information Center at the following Web address:

http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/index.htm?info/iphbi/iphbivirtualnative.htm

13.2.6 Ordering your eServer i5 system

To simplify the implementation of your eServer System i5 system, we recommend that you order the Novell or Red Hat Linux distribution with your eServer System i5 system, order feature code #0142 Linux specify, and order feature code #0496 Force i5/OS Preload.

Beginning May 2005, the IBM configuration tool (eConfig) is enhanced to support the preload of i5/OS on your eServer System i5 system, even when you elect to configure Linux partitions. Prior to this, when you ordered your eServer i5 and specified a Linux partition, the IBM manufacturing plant did not preload i5/OS and only installed i5/OS microcode on the load source disk unit. The remaining disk configuration and i5/OS installation tasks were then performed at the client location after the partitions were defined.

The #0496 FORCE i5/OS PRELOAD feature code preloads i5/OS on a new server across all disk units, even though logical partitioning has been specified. This feature code cannot be used if the Linux partition has dedicated disk controllers and drives included in the ordered configuration.

We recommend that you include this new feature code to preload i5/OS when you order the Linux distributions from Red Hat or Novell with the eServer System i5 system and virtual storage.

In the case of Virtual Partition Manager, since all of the Linux partitions will use virtual I/O, it is important that you take advantage of this IBM configuration support enhancement so that your new eServer i5 system comes fully loaded with i5/OS, ready for you to start deploying your Linux partitions.

13.3 Getting started with Virtual Partition Manager

In this section, we review the pre-requisites that you need to plan for and implement before defining your Linux partitions using Virtual Partition Manager. We recommend that you study the concepts for Partitioning for Linux for eServer i5 systems, documented in the eServer Information at the following Web address:

http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/index.htm?info/iphbi/iphbikickoff.htm

13.3.1 Minimum configuration requirements

The following requirements apply for Linux partitions created on an eServer i5 system. Each partition requires the following components:

- ▶ Processor unit:
 - 0.10 processing units allocated out of a shared processing pool
- ▶ Memory:
 - A minimum of 128MB of memory or region size (whichever is largest) is needed.
 - Hypervisor memory set aside from your total memory capacity—available for managing logical partitions. Your Logical Partition Validation Tool (LVT) gives you an indication of this memory requirement. For more information, see the memory section in the Information Center at the following Web address.
- ▶ Disk storage:
 - 1GB served through Virtual I/O by i5/OS
- ▶ Network interface:
 - One virtual Ethernet adapter for communicating to i5/OS or other Linux partitions
- ▶ Partition management console:
 - iSeries Operations Console, LAN Console or 5250 (Twinaxial) Console for partition creation and management
- ▶ Software requirements:
 - i5/OS V5R3 or above
 - Minimum of 2.3.0 version firmware applied. For information about minimum software requirements, and download instructions, visit the IBM eServer Prerequisite Web-site at:

http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/index.htm?info/iphathat1parmemory

http://www-912.ibm.com/E_DIR/ESERVERPREREQ.NSF

Select the **Software** tab with the options shown as Figure 13-1.

- Ensure that you have a minimum of PTF MF34753 (or its supersede) applied to enable Virtual Partition Manager. If this PTF is not already installed, please download and apply them before starting your partition setup.
- Prerequisite PTFs MF34946 and MF34962 must be permanently applied before doing a SAVSYS of the i5/OS partition. A D-mode IPL of the Virtual Partition Manager activated i5/OS partition will fail with an SRC of A6005001 with media that does not have these PTFs permanently applied.

Note: These minimum configurations may vary significantly depending on the workload enabled for each of the Linux partitions, and based on the Linux distribution and version that you install. Please refer to a specific Linux distribution's guideline for the actual values.

Figure 13-1 Prerequisite tool selections

13.3.2 Complete initial setup of your eServer i5

Before you define Linux partitions and load Linux distribution, you need to complete the Initial Server Setup tasks using either the predefined or customized setup checklists. You can find the checklists at the following Web address:

http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/index.htm?info/initialserversetup

Use the checklists initially for standalone servers, without specifying logical partitioning. Even though you will create Linux partitions at a later stage, for the purposes of setting up your server, you want to complete the installation so that it does not require a Hardware Management Console (HMC).

With your new eServer System i5 system, i5/OS currently owns all of the processor, memory, and I/O resources. The next few chapters provide step-by-step instructions for the basic configuration tasks required to remove resources from i5/OS and set up Linux partitions using Virtual Partition Manager. We will also discuss how to define the Network Server Storage Space (NWSSTG), and Network Server Storage Description (NWSD) for virtual storage and enable virtual Ethernet for network communications between multiple partitions.

13.4 Preparing your system for Virtual Partition Manager

This section provides step-by-step instructions on how you can remove logical resources from i5/OS using the Virtual Partition Manager, in preparation for defining new Linux partitions. With your new eServer i5 system, by default, the i5/OS partition currently owns all of the processor, memory, and I/O resources.

You can invoke the Virtual Partition Manager either through Dedicated Service Tools (DST) or System Service Tools (SST) tasks. The advantage of using SST is that you can bring your eServer i5 system to full operational mode rather than having to do your partition definitions in a restricted state where the rest of the operating system has not started, and users are not able to use the system.

We recommend that you manage your Linux partition, creation, and management through SST tasks, which enables you to use your i5/OS environment concurrently. After you define the partitions, a system IPL is required to complete the removal of resources from i5/OS.

13.4.1 Removing resources from i5/OS

1. From the Main menu or i5/OS command line, start the System Service tools by typing STRSST, as seen in Figure 13-2.

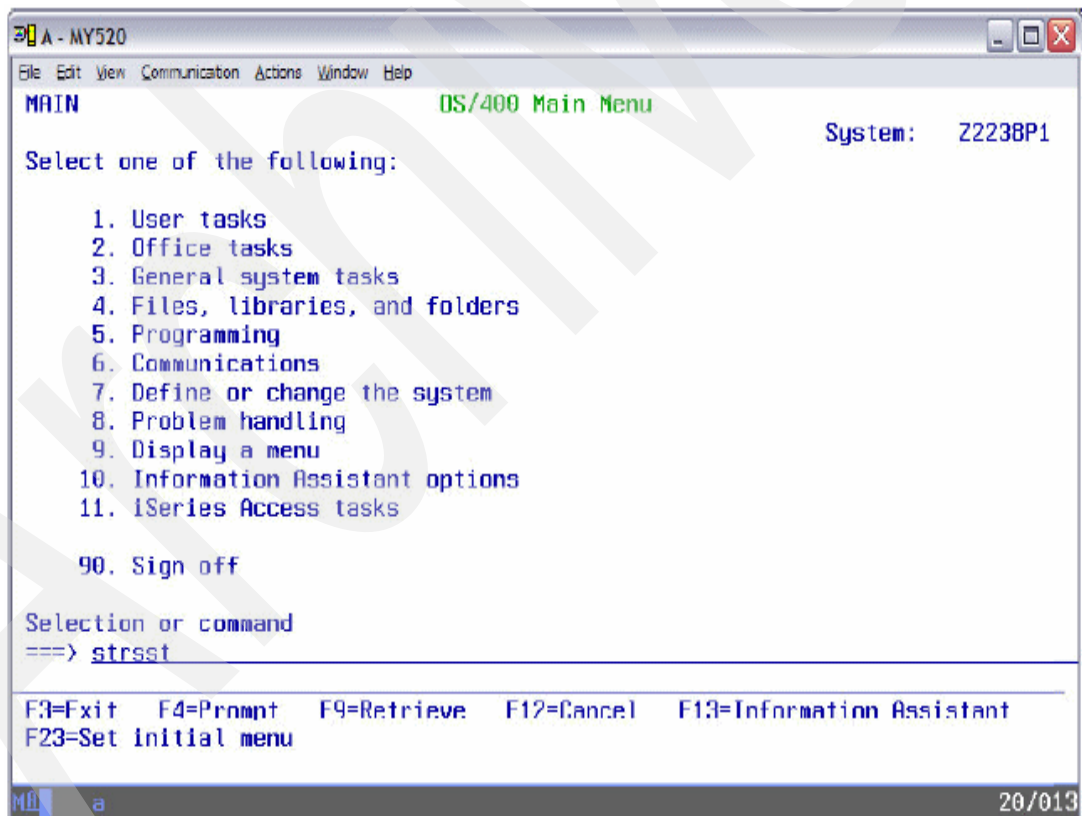


Figure 13-2 Main menu

2. Enter your user ID and a password as seen in Figure 13-3. This assumes that your Security Officer already created a DST/SST user profile for you to use, and provided adequate privileges to allow you to perform partition creation and management tasks.

For information about how to create DST user profiles, refer to the InfoCenter article on Managing service tools user ID and passwords at the following Web address:

<http://publib.boulder.ibm.com/infocenter/iseres/v5r3/ic2924/index.htm?info/rzamh/rzamhmanagst>

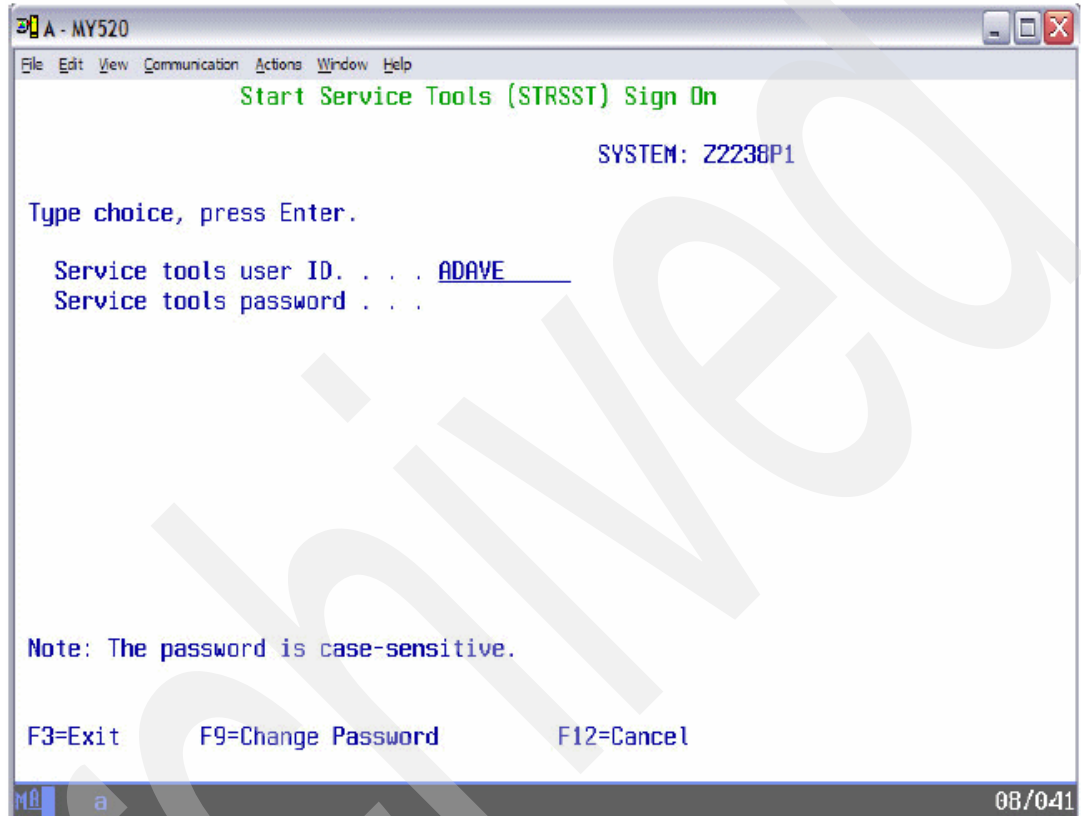


Figure 13-3 Start Service Tools sign-on screen

3. From the Start Service Tools (SST) menu, select option **5. Work with system partitions** as seen in Figure 13-4.

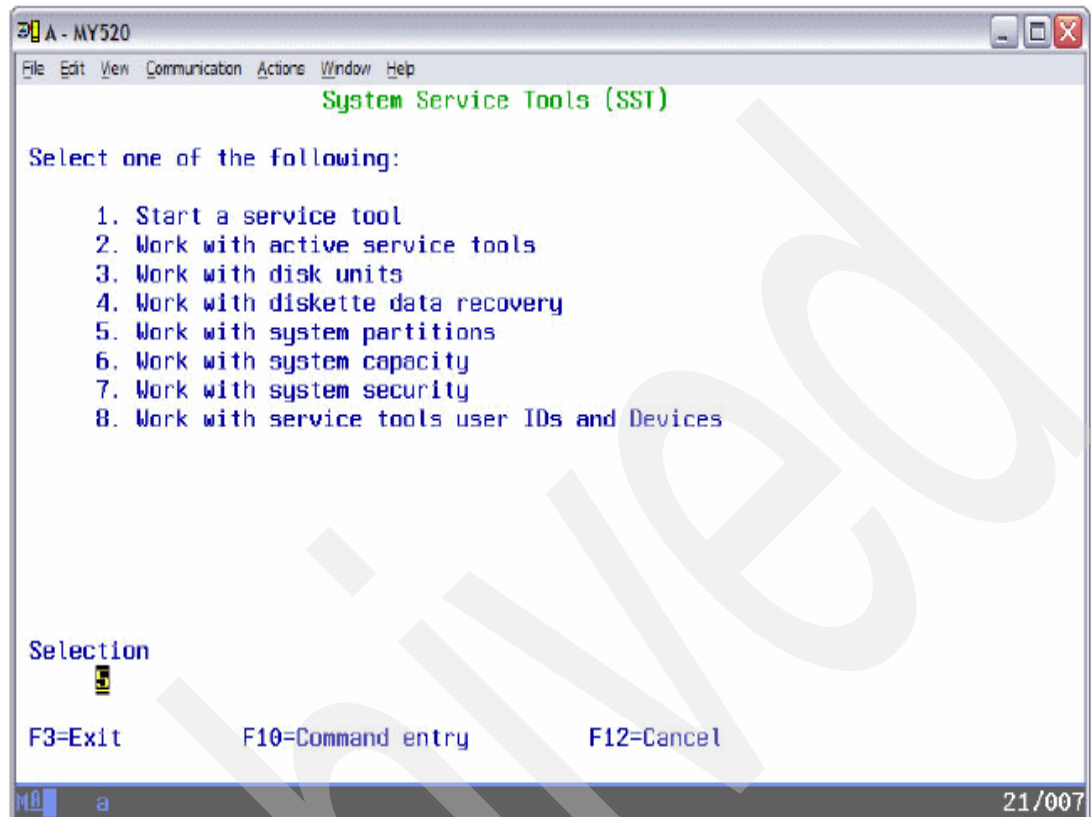


Figure 13-4 System Service Tools (SST)

An informational message appears as shown in Figure 13-5. This message appears when you enter the option to Work with System Partitions for the first time, or when you clear all partition configuration data.

Note: If another session is currently using the Virtual Partition Manager, an error dialog appears indicating that the tool is already in use.

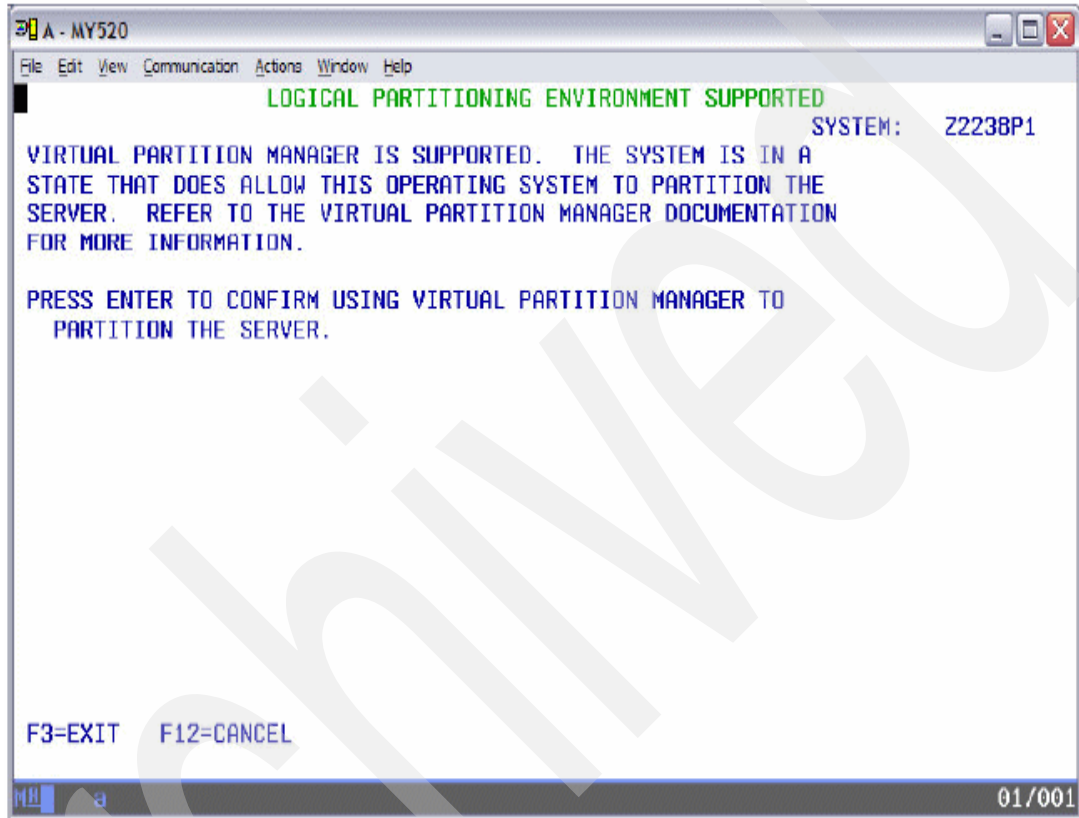


Figure 13-5 Initial informational display

4. Press Enter at the message. The Logical Partition management tasks appear, as seen in in Figure 13-6.

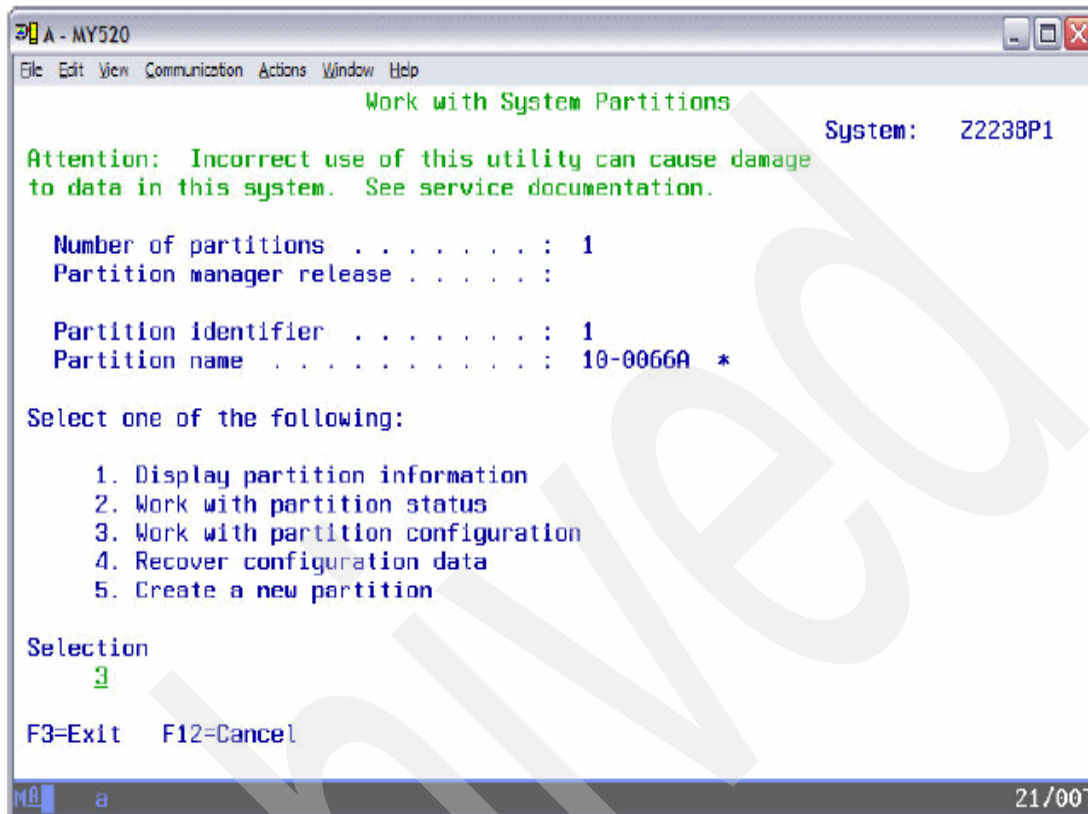


Figure 13-6 Work with System Partitions

5. Select option **3. Work with partition configuration**, as shown in Figure 13-6.

The objective of the next few steps is to ensure that we remove processing and memory resources from the i5/OS partition so that we can create Linux partitions. It is important that you complete the sizing exercise for your overall CPU and memory utilization for your i5/OS workloads to account for removal of such resources. Also take into account the overheads associated with micro-partition processor allocation. For more information, see the Whitepaper on i5/OS LPAR performance on POWER4 and System i5 systems.

Note: Option 1: Display Partition Information is not enabled for Virtual Partition Manager. Instead, use Option 2: Work with Partition Status, or use Option 3: Work with Partition Configuration to manage logical partitions using Virtual Partition Manager.

- From the Work with Partition Configuration menu, select option **2. Change partition configuration** for your i5/OS instance as shown in Figure 13-7.

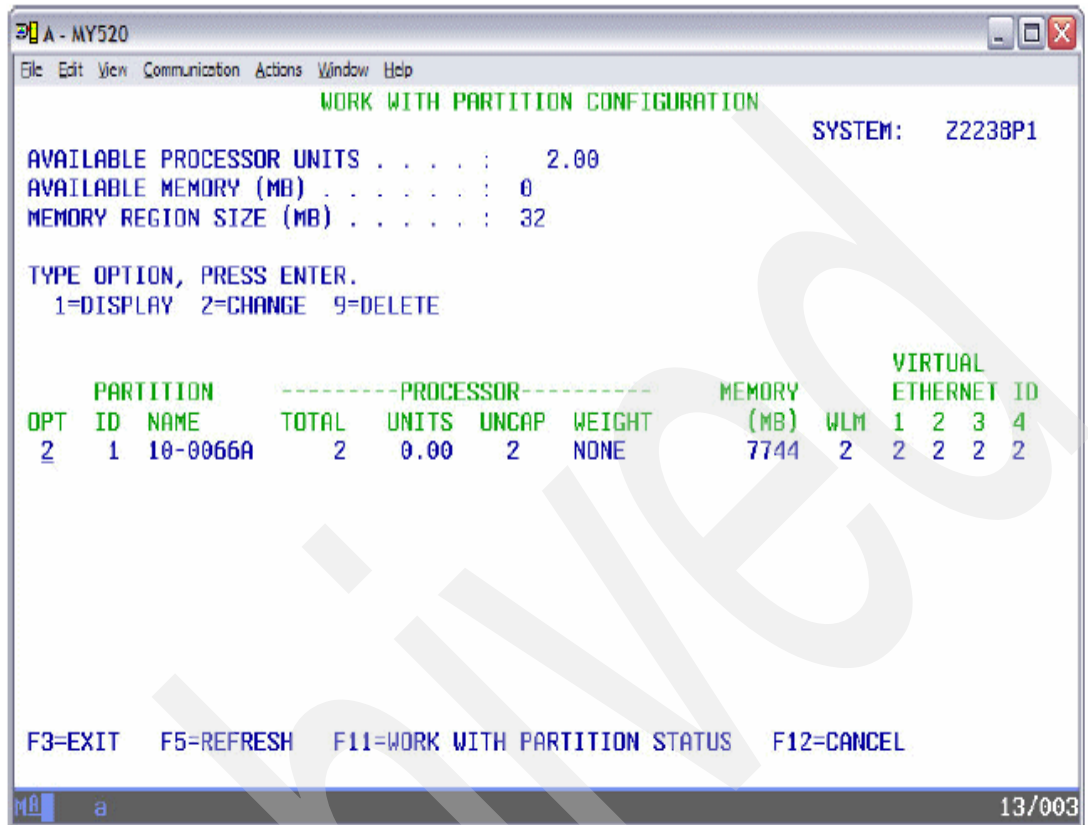


Figure 13-7 Work with Partition Configuration

With a new system, one i5/OS partition will be defined and the *Change Partition Configuration display* will show the defaults as shown in Figure 13-8. This is where we will remove resources from your i5/OS instance so that additional Linux partitions can be created.

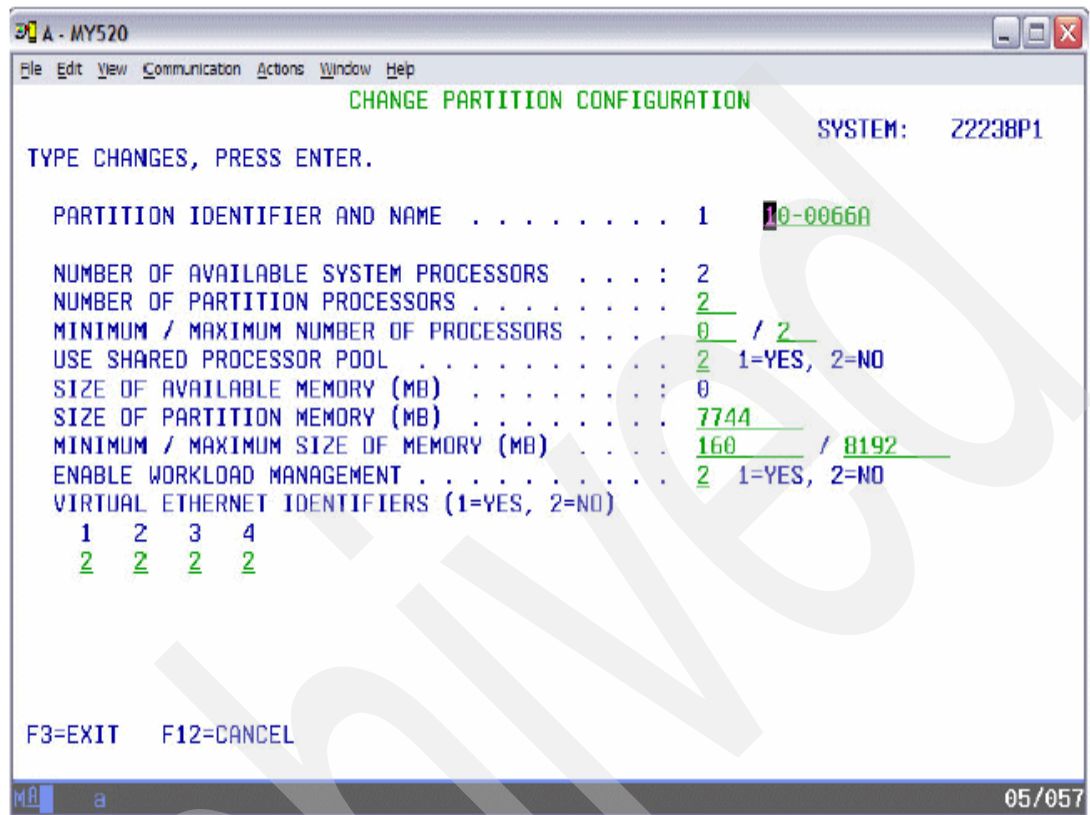


Figure 13-8 Change Partition Configuration

7. Make several changes here, based on the resources you want to set for the i5/OS partition. You need to assign the CPU and memory allocation according to the planning you completed with the Logical Partitioning Validation Tool (LVT). We examine all of the changes step-by-step as highlighted in Figure 13-9.




Figure 13-9 Change Partition Configuration Options

- ▶ **Minimum / Maximum Number of Processors:** Here, you enter the value for the minimum number of processors that you will need for your i5/OS environment, and the maximum number of processors that will be available for i5/OS. In our sample environment, we had 2 processors and we left the maximum value at 2 – enabling i5/OS to utilize all of the processors if the Linux partitions are not being used.
- ▶ **Use Shared Processor Pool:** The shared processor pool is a group of physical processors that provide processing capacity that can be shared among multiple logical partitions. The shared processor pool allows you to assign partial processors to a logical partition. In our example, since we are going to create four additional Linux partitions, we would not have sufficient resources to allocate a full processor for every partition. Another advantage for changing the partition configuration to use the shared processor pool is that the partition will have the capability to use processor resources unused by other partitions.

Note: When the current value of 2 is changed to 1, indicating that the shared processor pool is desired, the system prompts for new shared processor pool values once the Enter key is pressed.

- ▶ **Shared Processor Pool Units:** Specifies the total number of processing units that will be available to the partition after the resources are removed. In this example, the i5/OS partition will be left with 100 processing units, or a full processor after we have removed the CPU resources.
- ▶ **Minimum / Maximum Shared Processor Pool Units:** A minimum of 0.10 processing units is required for every full processor that may be utilized for the given partition. In our example, since we have a maximum of 2 processors that the i5/OS partition can utilize, the minimum value is therefore required to be set to 0.20 processing units with a maximum of 200 processing units, or 2 processors.
- ▶ **Uncapped processing:** By changing the value from 2=NO, to 1=YES, the partition may utilize excess shared processor cycles. The quantity assigned is dependent on the setting of the uncapped processing weight field that follows. If the value is 2=NO, then this partition will not be assigned excess shared processor cycles and will be a capped partition.

A capped partition indicates that the logical partition will never exceed its assigned processing capacity. The capped mode could be used if you know that a software application would never require more than a certain amount of processing power. Any unused processing resources will be used only by the uncapped partitions in the shared processor pool.

A partition using the uncapped mode indicates that the partition's assigned current processing capacity can be exceeded, up to the partition's current virtual processor setting when the shared processor pool has any unused processing power.

Note: Considering that Virtual Partition Manager does not support dynamic movement of resources, you might want to define your partitions as uncapped partitions to automatically assign unused processing units to a partition that needs the resources.

- ▶ **Uncapped Processing Weight:** With Virtual Partition Manager, since you can only create 4 additional Linux partitions, this range is restricted to a set of predefined values: 0=None, 64=Low, 128=Medium and 255=High. By setting the uncapped weight (255 being the highest weight), any available unused capacity is distributed to contending logical partitions in proportion to the established value of the uncapped weight. The default uncapped weight value is 128. For example, if partition 2 had an uncapped weight of 64 and partition 3 had an uncapped weight of 255, partition 3 would get up to four times the unused processing resources that partition 2 received. A value of zero will be required to be set for partitions that are capped.

- **Size of Partition Memory:** Linux partitions require a minimum of 128 megabytes. In this example, the value indicates the amount of main storage that remains with i5/OS partition. Make the new value multiples of the value set as the LMB size during your initial set up using ASMI. For example, you cannot set a value of 6700 because it gives an error message, like the one show in Figure 13-10. See the memory section, at the following Web address, in eServer Information Center for additional details:

http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/index.htm?info/iphat/iphat1parmemory

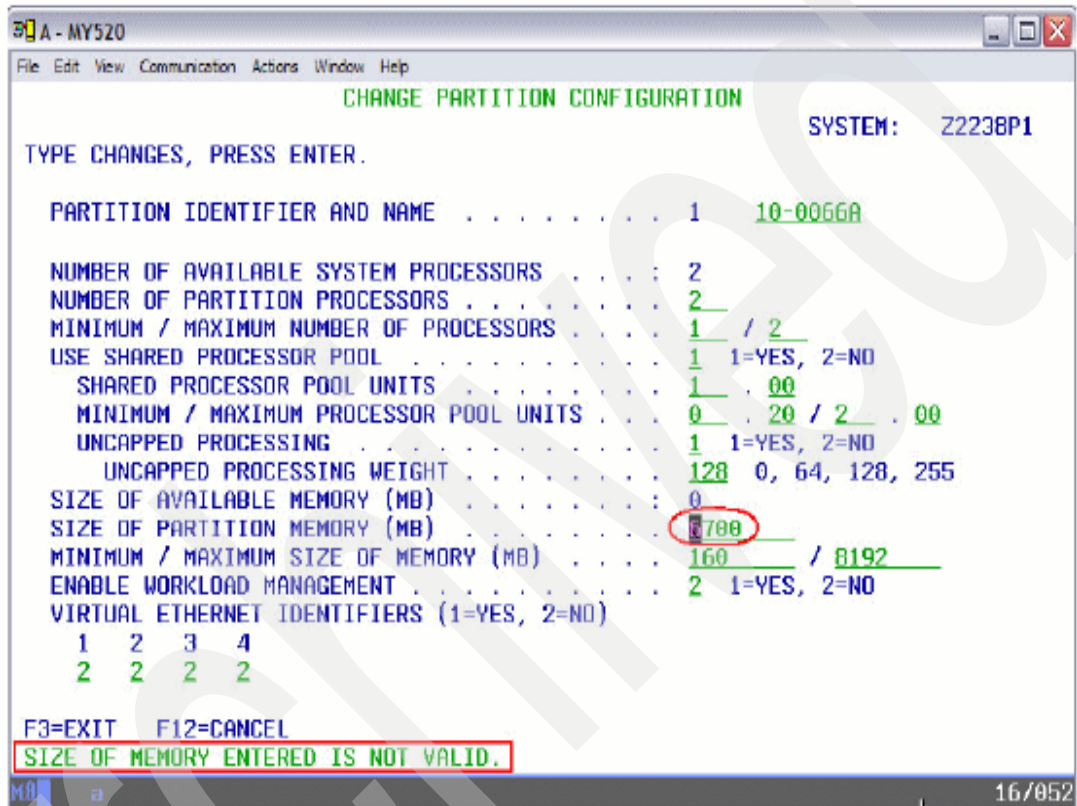


Figure 13-10 Change Partition Configuration Memory Error

- ▶ **Enable Workload Manager:** The default value for Virtual Partition Manager is set to 2 = No, meaning that the partition is not allowed to use future workload management tool within the partition to automatically adjust resource assignments, such as the IBM Enterprise Workload Manager.
 - ▶ **Virtual Ethernet Identifiers:** A value of 1 indicates you are enabling one of the virtual Ethernet communications ports for inter-partition communications between Linux or i5/OS partitions. In this example, we changed the value for one of the ports to 1, which will create a virtual Ethernet adapter under i5/OS with a resource type of 268C and an associated resource name that will be used during the set up of Ethernet line description.
8. Remove CPU and memory resources, enable the i5/OS partition to take advantage of shared processor pool and uncapped processors, and enable a virtual Ethernet port for inter-partition communications. Press Enter. The message screen shows in Figure 13-11.




Figure 13-11 Confirm Changed Partition

9. Press Enter again on the Confirm Changed Partition screen, which will complete the changes required on the i5/OS partition (Figure 13-12).




Figure 13-12 Partition Change Successful

Notice that the changes made to i5/OS resources requires an IPL of the system, as indicated by the "<" in Figure 13-12. There is no need to perform this task at present, you can perform the IPL once you define all of the Linux partitions.

Also notice that changes to memory allocation in i5/OS partition are not immediate.

10. Use option **1** against the new partition you created, and then function F9 to toggle between CURRENT and PENDING values. You will see that the memory is gradually removed from the owning partition. When you remove large amounts of main storage from the owning partition, the removal may take a considerable time.

Note: Notice that memory is released quicker if you define your partitions when the system is in a restricted state through DST.

11. Figure 13-13 shows the available CPU and memory resources for creating new partitions. In the next section, we use these resources to define new Linux partitions.

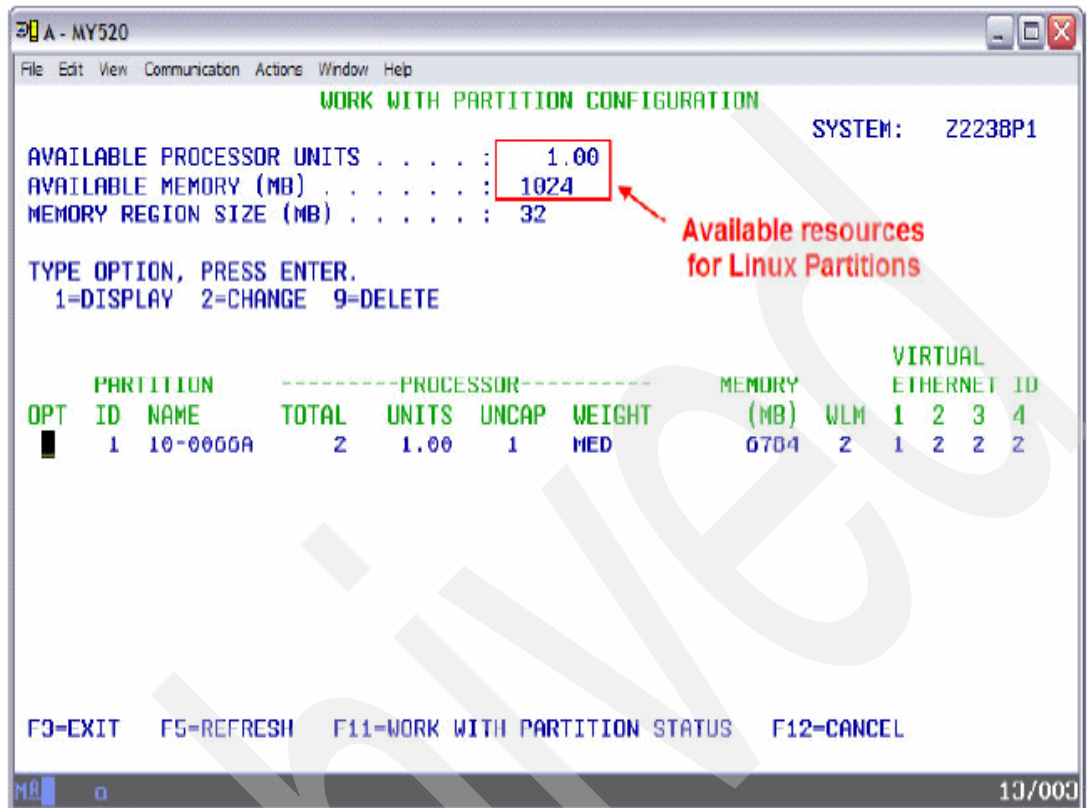


Figure 13-13 Available Resources for Creating New Partitions

12. You are now ready to create Linux partitions using Virtual Partition Manager, and the resources that have been removed from i5/OS partition. Even though the Work with System Partitions display shows that you may require a System IPL, you should wait until all of the new Linux partitions are defined. The next chapter looks at step by step instructions on how to define a Linux partition using Virtual Partition Manager.

13.4.2 Understanding virtual I/O slots and requirements for IPLs

When defining Linux partitions using Virtual Partition Manager, you must think in terms of setting up a brand new server — except that it does not have its own hardware. However the requirement still remains. Resources such as memory, CPU, serial SCSI, Ethernet LAN, and serial client adapter are just some of the resources you need on your computer. In the case of eServer i5, the partition hypervisor virtualizes all of these resources through enabling virtual I/O slots.

Depending on the number of Linux partitions you define and on the number of virtual I/O slots you enable (such as virtual SCSI, virtual Ethernet LANs), you may or may not notice that an IPL is required.

By default, the i5/OS partition is enabled with 8 virtual I/O slots. Internally, two of these virtual slots are taken up by virtual serial, and one for virtual management device. This leaves five free slots. When you create a partition, it takes up to 2 virtual I/O in the i5/OS partition: one for the virtual SCSI server adapter, and one for the virtual serial client adapter. When you add the i5/OS partition to a virtual Ethernet LAN, it uses one additional virtual slot.

When you exceed the allocation of five virtual I/O slots, the partition hypervisor sends a command to allocate more virtual I/O slots. At this point, it just gets everything it will ever need. Therefore, once you go beyond the eight available virtual I/O slots, the system sets itself up to use all of the available virtual I/O slots and enables them during the next IPL. Hence, notice that when you define more than one Linux partition — and enable a couple of virtual Ethernet LANs — the Partition Configuration Screen shows you that you need to IPL your system to implement the changes.

Note that once you get all of the allocations you need, even if you remove the LPARs or delete them, the system does not reset the virtual I/O slot count.

In summary, notice that during Linux partition creation, you need an IPL of the system to activate the changes in virtual I/O slots. This is only required once, and its best to perform it during the initial set up of your LPARs.

13.5 Creating Linux partitions using Virtual Partition Manager

This chapter provides step by step instructions for creating a Linux partition on eServer i5 system using the Virtual Partition Manager, after you have removed resources from the i5/OS partition.

13.5.1 Defining new Linux partitions

1. From the *Work with System Partitions* display, take option **5. Create a new partition** as shown in Figure 13-14.

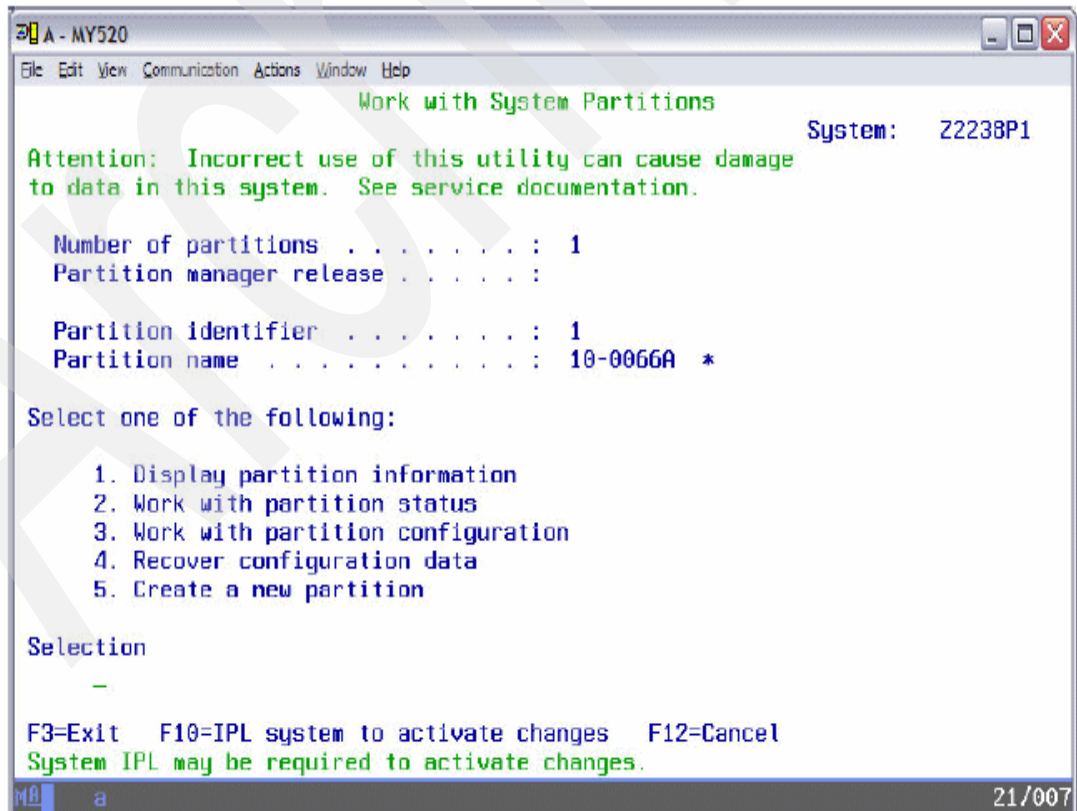


Figure 13-14 Work with System Partitions

2. Assign the values for creating the new Linux partition as per your Logical Partitioning Validation Tool (LVT) output, as shown in Figure 13-15. A brief explanation of each of the value is also provided here:
 - **Partition Identifier and Name:** Enter the partition name for your Linux partition. You can also change the partition identifier should you choose to. In this example, the default given (next partition identifier number) by the system is selected.
 - **Number of Partition Processors:** If you are defining a Linux partition with dedicated processors, then enter the value of the processors that will be dedicated to that partition. For partitions utilizing shared processor pool, you can enter the value of the number of virtual processors that you plan to utilize in your applications. This value will also affect how you set up your Min/Max number of processors, and also how you set up your Min/Max shared processor pool units later. For example, if you set a value of 4, then the maximum number of processors in the next field (as shown in **3** below) should also be set to 4. Also, you will need at least 0.40 processing units specified in step **4**.
 - **Minimum / Maximum number of Processors:** Enter the value for minimum number of processors that you will need for your Linux environment, and the maximum number of processors that will be available for Linux.
 - **Shared Processor Pool Units:** Assuming that you are going to use shared processor pool, you now need to specify the total number of processing units that will be available to the partition during start up. In our example, since we are going to create four additional Linux partitions, we would not have sufficient resources to allocate a full processor for every partition to be set up as having a dedicated CPU.

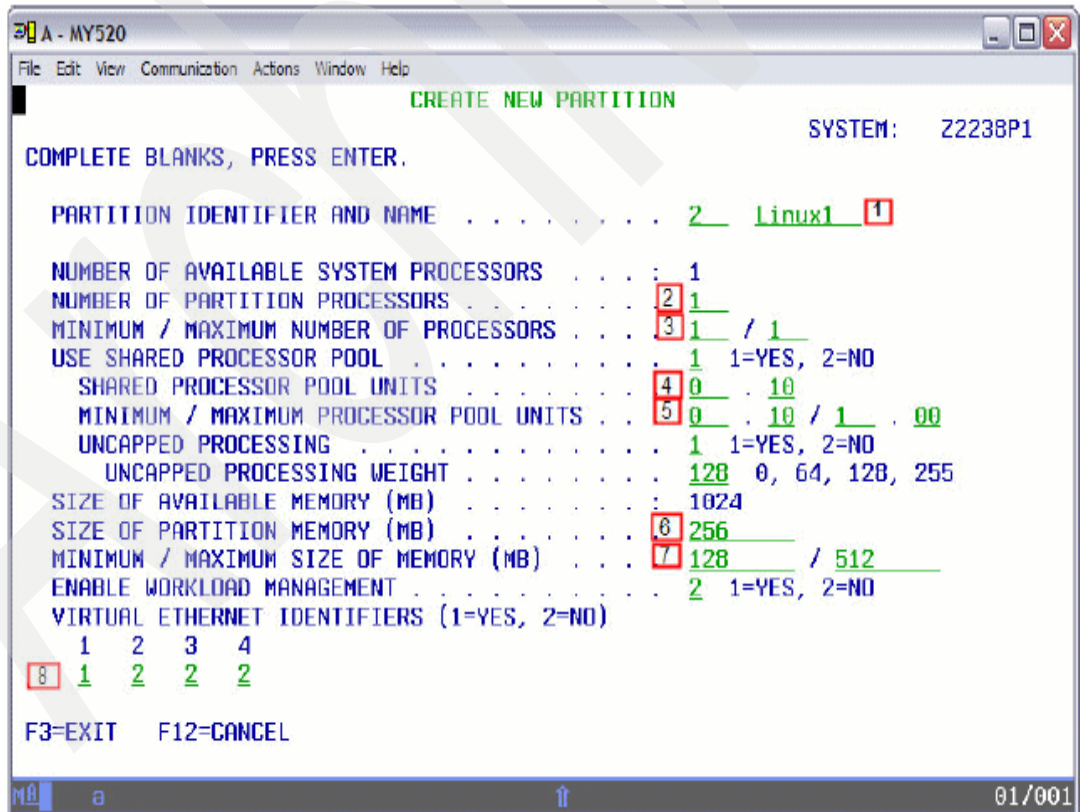


Figure 13-15 Create New Partition

- **Minimum / Maximum Shared Processor Pool Units:** A minimum of 0.10 processing units is required for every full processor that may be utilized for the given partition. Assign the values appropriately based on the range you want your partitions to utilize unused processing cycles.
- **Uncapped processing:** You have the option to have your Linux partition shared capped, or shared uncapped. See the shared processor section in the IBM Information Center for more information about capped and uncapped processors.

<http://www-1.ibm.com/servers/eserver/series/perfmgmt/pdf/lparperf.pdf>

By default, an uncapped processing weight of 128 (medium) is assigned. A value of 255 indicates that the partition has a higher priority to utilize idle CPU cycles compared to a partition with 128 (medium) or 64 (low). For example, if partition 3 had an uncapped weight of 64 and partition 4 had an uncapped weight of 255, partition 4 gets up to four times the unused processing resources that partition 3 received. A value of zero is required to be set for partitions that are capped.

Note: Considering that Virtual Partition Manager does not support dynamic movement of resources, you might want to define your partitions as uncapped partitions to automatically assign unused processing units to a partition that needs the resources.

- **Size of Partition Memory:** The minimum memory requirement for Linux partition is 128 megabytes. At the following Web address, see the memory section in IBM eServer Information Center for additional details.

http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/index.htm?info/iphathat1parmemory

Set the appropriate values based on your planning done with Logical Partitioning Validation Tool (LVT).

- **Virtual Ethernet Identifiers:** A value of 1 indicates you are enabling one of the virtual Ethernet communications ports for inter-partition communications between Linux or i5/OS partitions. You can select up to 4 virtual Ethernet communications links.

- Once these values are set, you will get a confirmation display as shown below in Figure 13-16. Press Enter, and you will be returned back to the *Work with System Partitions* display. You can repeat the above steps to define another Linux partition, if necessary.




Figure 13-16 Confirm New Partition

- Once you have defined all of your partitions, you can view them using Option 3 from *Work with System Partitions* as shown below in Figure 13-17. In this example, Linux4 was defined as a capped processor partition.

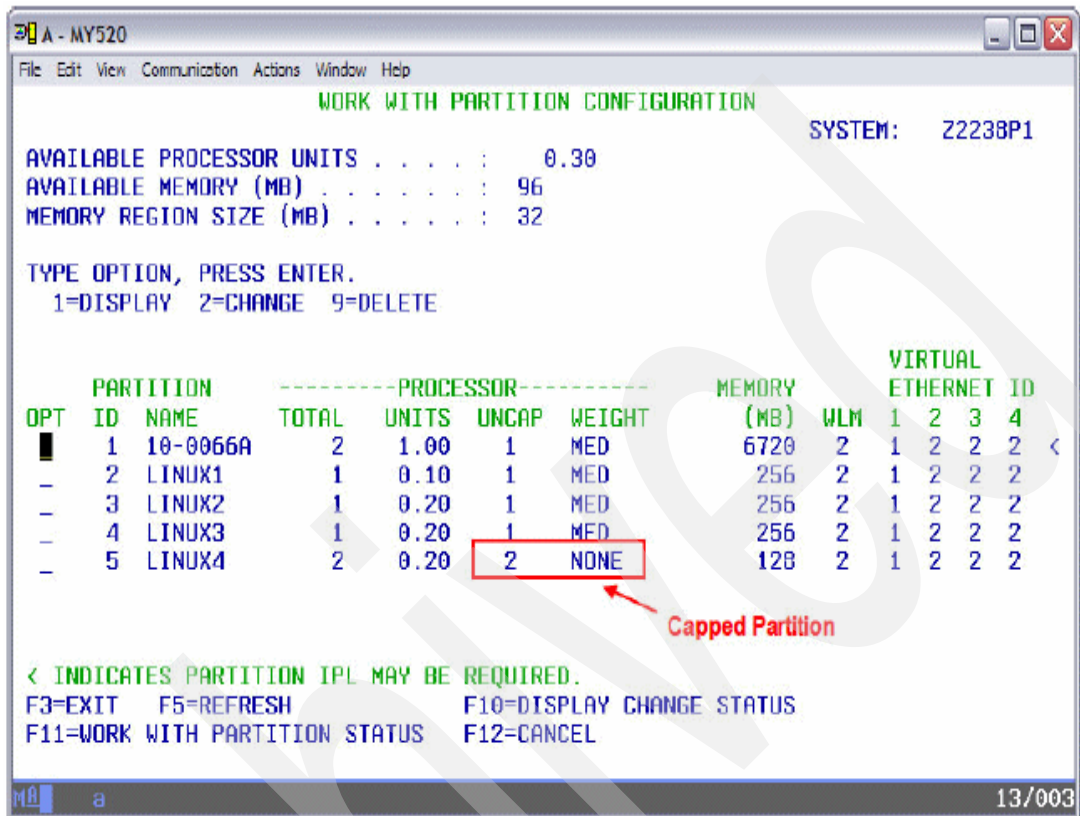


Figure 13-17 View of All new Partitions Defined

5. You can either update the partitions to change any resource configurations, or delete them and recreate them. Keep in mind that you can only change or delete one partition at a time. If you wanted to start all over again and clean up all of the configuration information, you can use the option to Clear Partition Configuration Data as discussed in the “Recover configuration data” on page 403 section.
6. Exit out of the Work with Partition Configuration display to return back to the Work with System Partitions. You can now start the system IPL to finish the updates required to the hypervisor for enabling the new Linux partitions.
7. Select the F10 function key as shown in Figure 13-18.

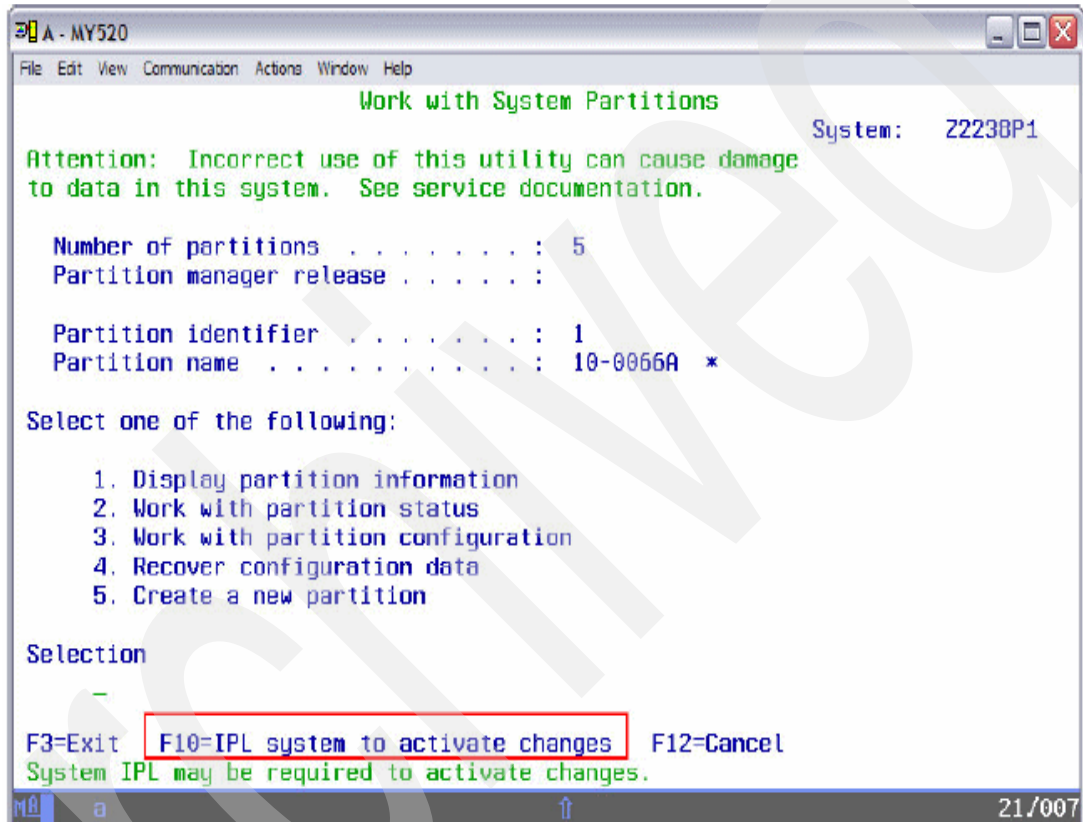


Figure 13-18 Ready to IPL from Work with System Partitions

Recover configuration data

If for some reason you want to restart from the very beginning by deleting all of the partitions, you may want to do so by taking option 5 from the Work with System Partitions display, and then selecting option 7. **Clear configuration data**, as shown here in Figure 13-19.

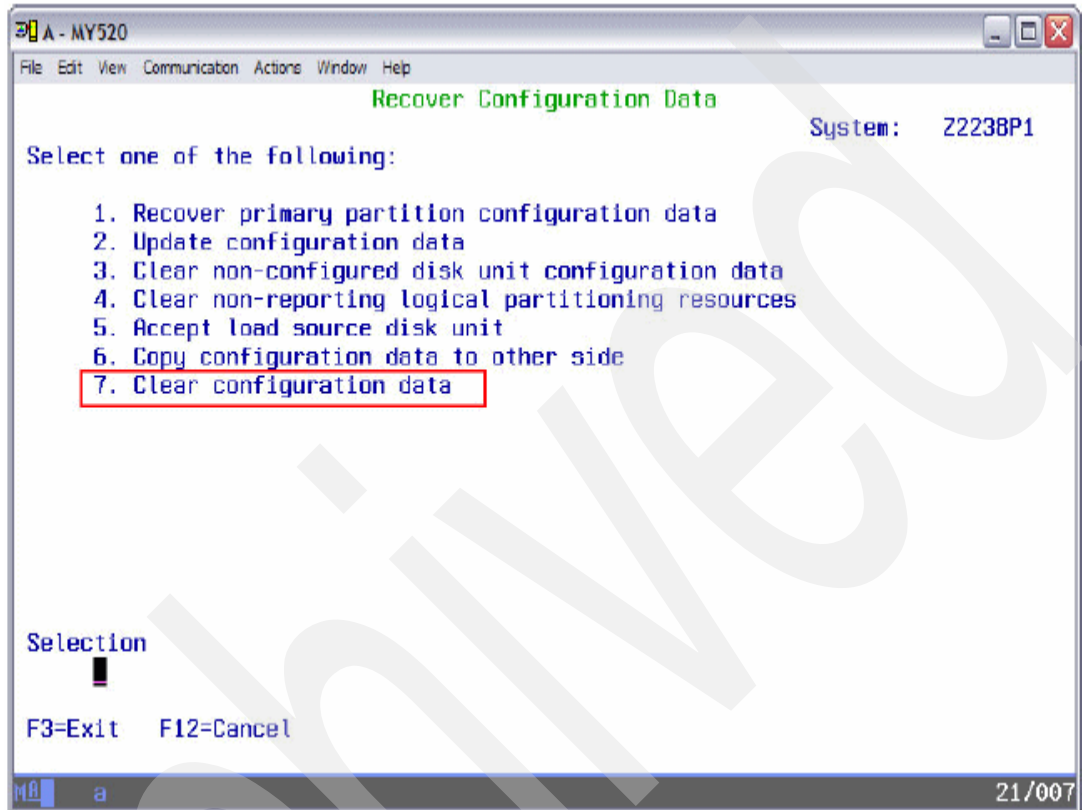


Figure 13-19 Clear configuration data

Take care when taking this option, as it completely removes all Linux partition configurations. You also need to take this action when migrating from Virtual Partition Manager to HMC managed Linux partitions.

Migration considerations when moving to HMC

The following steps outline some of the planning considerations for migrating off of the Virtual Partition Manager to HMC managed Linux partitions. Please note that you cannot save and restore partition configuration data; instead, you must create them in their entirety. However, you do not need to recreate your data saved in Linux partitions through the Network Server Storage Space. Point to the same virtual SCSI and virtual I/O resources when partitions are defined through HMC:

1. Power off all of the Linux partitions.
2. Start SST or DST, and select the option to **Work with system partitions**.

If you do not already have your LVT prints, you may want to take note of the resource allocations you have for each of your Linux partitions such as CPU, memory, Min/Max values, number of virtual Ethernet LANs, and the eWLM settings. Keep in mind that you cannot print the configuration data through any report from the Virtual Partition Manager displays.
3. Clear the Linux partition configurations using option 4 from the Work with System Partitions display. The Recovery Partition Configuration screen appears.
4. Select option **7. Clear partition configuration data**.
5. IPL your server after clearing the partition configuration data.
6. Connect HMC to your server.
7. Create a new profile for the Linux partition that contains the physical I/O, memory, and processors desired.
8. Perform the following actions for each Linux partition:
 - Create a Virtual SCSI Server adapter in the i5/OS partition.
 - Create a Virtual Serial client adapter in the i5/OS partition that points to the Virtual serial Server in the Linux partition in slot 0.
 - Create a new partition for Linux. You want to place the Virtual SCSI Client adapter in Virtual slot 2 because data in the NWSD points to this location.
 - Create any Virtual Ethernet devices you had in the Linux partition.
 - Make the Linux partition's power controlling partition to be the i5/OS partition.
 - Power on and Power off the Linux partitions. This drives the configuration down to the hypervisor.
9. Vary on the NWSD in the i5/OS partition.

13.6 Establishing network connectivity for Linux partitions

This chapter provides the steps for configuring support in the i5/OS partition for exposing the virtual LAN connections of the Linux partitions to an external network via a physical network connection in the i5/OS partition.

Note: There are multiple ways to configure virtual network access to a physical network including Proxy ARP, Network Address Translation, and Static Routing via a Linux partition. This document only presents setup instructions for the Proxy ARP method.

One method you can use for bridging virtual network connections to an external network is through the use of Proxy ARP. With Proxy ARP, the physical network adapter of the i5/OS partition responds to all network requests within a defined range and then routes those addresses to the virtual LAN segment.

There are a number of steps that you must complete for a Proxy ARP configuration, which include the following actions:

- ▶ Define a virtual LAN segment that places all of the Linux partitions and the i5/OS partition on the same virtual LAN. Complete this through the LPAR definition (discussed earlier).
- ▶ Create an Ethernet Line Descriptor for the virtual LAN adapter defined for the i5/OS partition.
- ▶ Create a TCP/IP interface for the virtual LAN adapter defined for the i5/OS partition. This TCP/IP interface is “associated” with the TCP/IP interface of the physical adapter allocated to the i5/OS partition.
- ▶ Define network parameters in the Linux partition(s) including IP Address, subnet mask, and gateway (which is the address of the i5/OS partition on the virtual LAN).

13.6.1 Defining i5/OS components that support Proxy ARP

Following are the detailed steps for defining the i5/OS components that support Proxy ARP:

1. Display the communication resources in the i5/OS partition (WRKHDWRSC *CMN) to determine the resource that was created for the virtual network adapter in the i5/OS partition.

The communication resource used for the virtual LAN connection is of a type 268C, enabled and found when you use the `WRKHDWRSC *CMN` command, as shown in Figure 13-20. The corresponding resource name (CMNxx) is used in the Ethernet Line Descriptor.

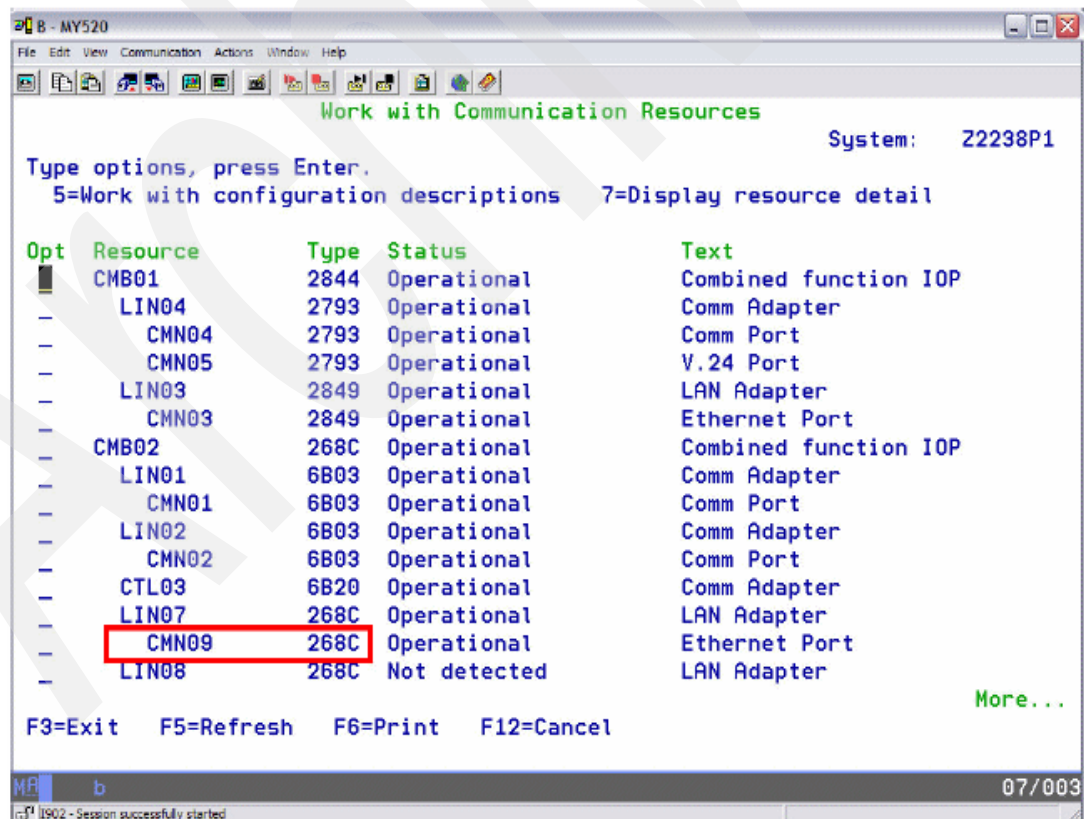


Figure 13-20 Work with Hardware Resources

2. Use the **CRTLINETH** command to create the Ethernet line descriptor. See Figure 13-21.




Figure 13-21 Create Ethernet Line Description

- **Line Description:** Enter the name for the Ethernet line description. For virtual Ethernets it is a common practice to start the name with 'VRT' and include the virtual Ethernet number in the name. As an example, if you are creating the line descriptor for the 1st virtual Ethernet, a common name to use would be 'VRTETH01'.
- **Resource Name:** The resource name is the name of the communication resource for the virtual network connection on the virtual LAN. This is the resource that was identified earlier with the **WRKHDWRSC *CMN** command.

- Once the Ethernet line descriptor is created, it needs to be varied on. You can accomplish this with the Work with Line Descriptors (**WRKLNDD**) command as seen in Figure 13-22.

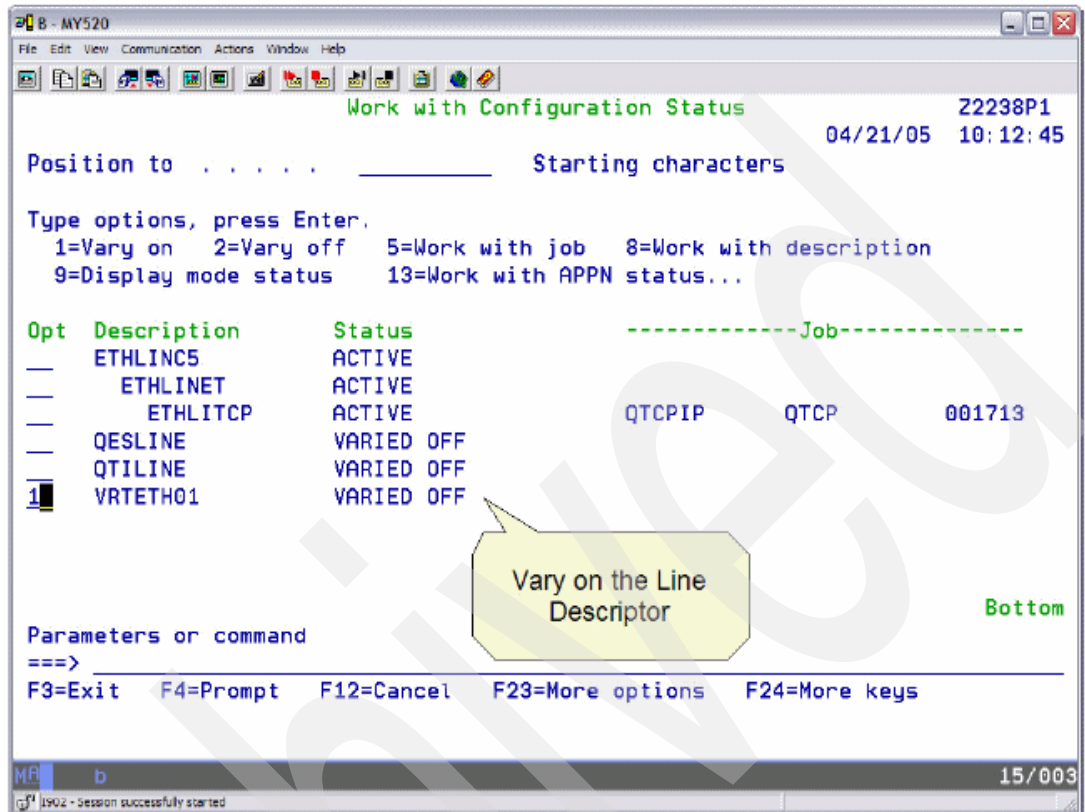


Figure 13-22 Work with Configuration Status

- Create the TCP/IP interface for the i5/OS network adapter for the virtual LAN. To create the TCP/IP interface, type the command **ADDTCPIFC** (Add TCP/IP Interface) as seen in Figure 13-23.




Figure 13-23 Add TCP/IP interface

- **Internet Address:** Type the address of the network interface for the i5/OS partition on the virtual LAN.

Note: This is the address that Linux partitions use as their gateway (or route) to the external network.

- **Line description:** The line description is the name of the Ethernet line descriptor (defined earlier) for the network adapter on the virtual LAN.
- **Subnet Mask:** The subnet mask defines the size of the network to which the interface is being. For virtual LAN, the subnet mask, along with the address of the interface determine the range of addresses for which are proxies.
- **Associated local interface:** The associated local interface is the address of the TCP/IP interface for the real (physical) network adapter in the i5/OS partition. The associated local interface routes traffic between the external (physical) network and the virtual network.

Note: Setting up of a Proxy ARP environment requires that you define a set of addresses that meet a number of specific rules including starting and ending boundaries as well as network size. Use a subnet calculator to help define the network range.

5. Once the TCP/IP interface is created, start it. You can start the interface through the Configure TCP/IP (**CFGTCP**) command. After you specify the command, select option 1 to Work with TCP/IP Interfaces, and then select option 9 to start the newly defined interface as seen in Figure 13-24.


```

Work with TCP/IP Interfaces
System: 22238P1

Type options, press Enter.
  1=Add  2=Change  4=Remove  5=Display  9=Start  10=End

Opt  Internet      Subnet      Line      Line
     Address      Mask       Description Type
-----
  9  9.5.110.46    255.255.255.128  ETHLINC5  *ELAN
     9.5.110.49    255.255.255.240  VRTETH01  *ELAN
     127.0.0.1     255.0.0.0      *LOOPBACK *NONE

```

Figure 13-24 Start TCP/IP interface

6. Proxy ARP requires that TCP/IP packets flow between two network interfaces (the virtual and real/physical interface). This requires that “IP datagram forwarding” be enabled. Enter the command **CHGTCPA** (Change TCP/IP Attributes), and change the value of IP datagram forwarding to ***YES** as seen in Figure 13-25.

```

Change TCP/IP Attributes [CHGTCPA]

Type choices, press Enter.

TCP keep alive . . . . . 120 1-40320, *SAME, *DFT
TCP urgent pointer . . . . . *BSD *SAME, *BSD, *RFC
TCP receive buffer size . . . . . 8192 512-8388608, *SAME, *DFT
TCP send buffer size . . . . . 8192 512-8388608, *SAME, *DFT
TCP R1 retransmission count . . . . . 9 1-15, *SAME, *DFT
TCP R2 retransmission count . . . . . 16 2-16, *SAME, *DFT
TCP minimum retransmit time . . . . . 250 100-1000, *SAME, *DFT
TCP closed timewait timeout . . . . . 120 0-14400, *SAME, *DFT
TCP close connection message . . . . . *THRESHOLD *SAME, *THRESHOLD, *ALL...
UDP checksum . . . . . *YES *SAME, *YES, *NO
Path MTU discovery:
  Enablement . . . . . *YES *SAME, *DFT, *NO, *YES
  Interval . . . . . 10 5-40320, *ONCE
IP datagram forwarding . . . . . *YES *SAME, *YES, *NO
IP source routing . . . . . *YES *SAME, *YES, *NO
IP reassembly time-out . . . . . 10 5-120, *SAME, *DFT

F3=Exit F4=Prompt F5=Refresh F10=Additional parameters F12=Cancel
F13=How to use this display F24=More keys

More...

```

Figure 13-25 Change TCP/IP attributes

Note: After you enable datagram forwarding, ping the address of the i5/OS TCP/IP interface on the virtual LAN. If the interface responds, then the set up of Proxy ARP is complete and is ready for use by the Linux partitions.

13.7 Setting up i5/OS virtual I/O resources for Linux partitions

There are three components you need create in i5/OS to support Linux partitions with hosted resources. This chapter provides instructions for defining the Network Server Descriptor and Network Storage Space.

13.7.1 Network Server Descriptor

The Network Server Descriptor defines a number of parameters for the Linux environment including startup location as well as startup parameters.

Use the following steps to define the Network Server Descriptor:

1. Start the Create Network Server Descriptor using the **CRTNWS** command as seen in Figure 13-26.




Figure 13-26 Create network server descriptor

- **Network server description:** This is the user-defined name for the Network Server.
- **Resource name:** The Resource name indicates the Virtual SCSI server adapter that provides virtual I/O resources (virtual disk [NWSSTG], virtual CD/DVD, virtual tape) to the Linux partition that has the corresponding Virtual SCSI client adapter. *AUTO indicates that the system determines the resource name of the first (and in this case the only) Virtual SCSI server adapter for the partition.
- **Network server type:** For Linux partitions, the Network server type is always set to *GUEST.
- **Partition:** Partition is the name of the partition created with the Virtual Partition Manager via SST.
- **Code page:** For Linux partitions the code page is always set to 437.

Note: The Code page defined in the Network Server Descriptor has no relationship to the Language support provided by Linux. The Code Page defined here is used for internal communications between the Hypervisor and the Logical Partition.

2. After you complete the first page of parameters, press the <PgDn> key twice to display the last page of parameters as seen in Figure 13-27.

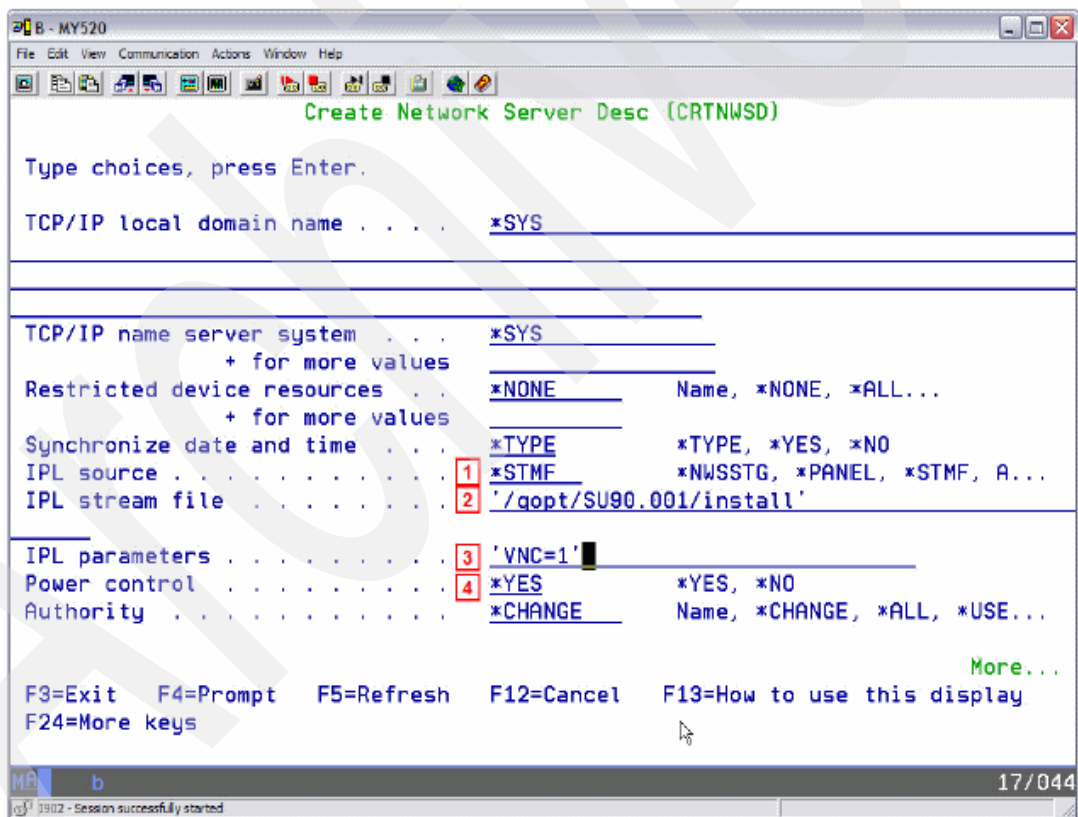


Figure 13-27 Network Server Descriptor

- **IPL source:** The IPL source indicates where to look for the initial boot file. A *NWSSTG setting indicates that the initial boot file is in the bootable disk partition of the first disk linked to the Network Server Descriptor. A *STMF setting indicates that the initial boot file is a stream file located in the IFS. When the setting is *STMF the path indicated by **2** is used.

Note: The installation of Linux is typically performed with IPL Source set to *STMF. Post-installation changes the setting to *NWSSTG to boot off of the Linux kernel written to the virtual disk by the installation program.

- **IPL stream file:** The IPL stream file is used when **1** is set to *STMF. The path indicates the path to the initial boot file.

Note: The initial boot file for an installation of Linux is indicated in the documentation provided with the distribution

- **IPL parameters:** The IPL parameters field defines parameters that are sent directly to the Linux kernel. The IPL parameters are defined in the Linux-distributor documentation.

Note: Both Novell / SuSE SLES 9 and Red Hat RHEL 4 support graphic-based installations. The IPL parameters field can be used to indicate that a graphical-based installation is requested.

Note: The parameters defined in IPL parameters are sent directly to the Linux operating system; therefore, you must enter them exactly as specified in the Linux documentation including upper/lower case.

3. After you complete the fields, press Enter to create the Network Server Descriptor.

13.7.2 Network Server Storage Space

The Network Server Storage Space (NWSSTG), also referred to as virtual disk, provides the disk resource to the Linux partition. The Network Storage Space is an object within the Integrated File System (IFS) that is abstracted to the Linux partition such that it appears, and is used, as an actual disk drive.

Use the following steps to create the Network Server Storage Space:

1. Type the Create Network Server Storage Space command, **CRTNWSSTG**, which creates the Network Server Storage Space. Entering the command displays Figure 13-28.




Figure 13-28 Create Server Storage Space

- **Network server storage space:** The Network server storage space is a user-defined name given to the network server storage space.
- **Size:** The size field indicates the size (in megabytes) for the virtual disk. The size can be anywhere from 1 megabyte to 1 terabyte. For a Linux installation, the installer uses the size of the virtual disk to determine the default package selection.

Note: For both Novell / SuSE SLES 9 and RedHat RHEL 4, a size of 4 GB (4096) supports a default package selection (during installation) that installs the most common open-source applications, including SAMBA, for file serving and Apache for Web Serving.

- **Format:** For Linux partitions, set the format to *OPEN. During the Linux installation, the “disk” represented by the virtual disk is formatted to have the number of disk partitions that installation requires.
2. After you complete the fields, press Enter to cause the Network Server Storage Space creation. At this point the space requested is reserved in the IFS and the Network Server Storage Space is available for use.

- Associate the Network Server Storage Space with the Network Server, by linking the storage space to the network server. Type the Add Server Storage Link command, **ADDNWSSTGL** as seen in Figure 13-29.

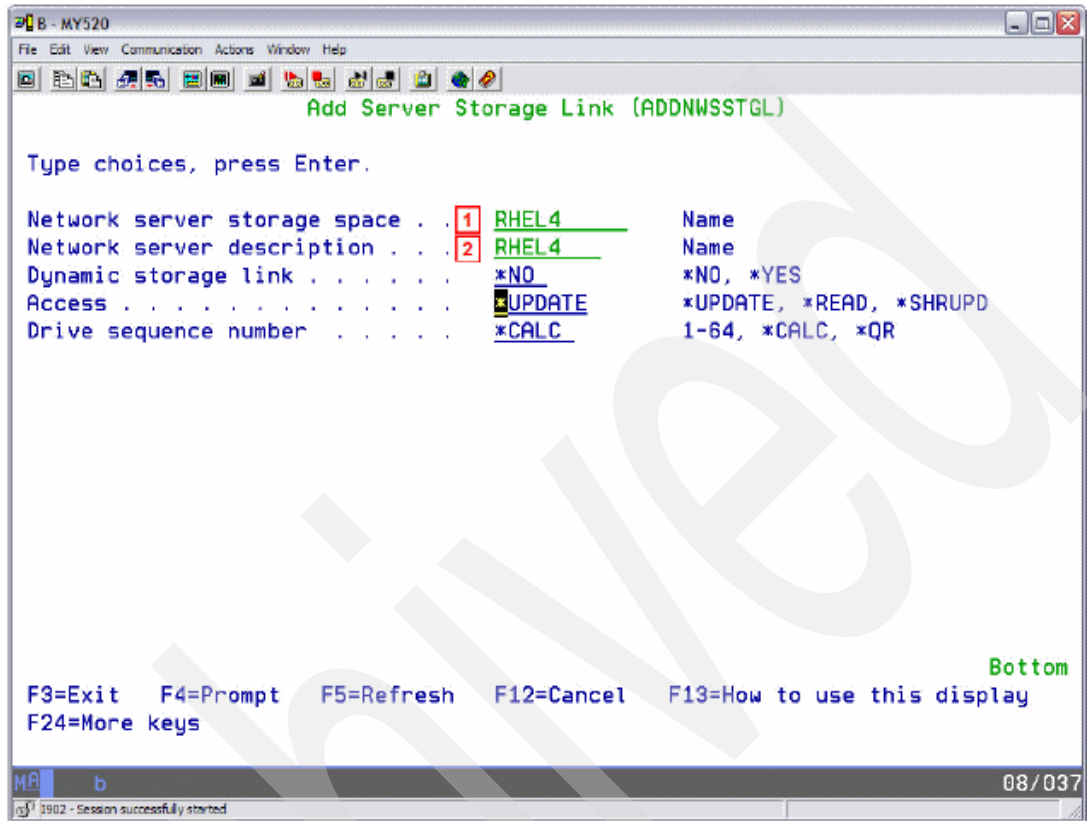


Figure 13-29 Add Server Storage Link

- **Network server storage space:** The name of the Network Server Storage Space to be linked to the Network Server defined in 2.
- **Network server description:** The name of the Network Server to which the storage space defined in 1 is linked. The Network Server is the component that provides access to the storage space(s), via the virtual SCSI adapters that are linked to the Network Server.

Virtual console access

Access to the Linux console is provided through the hosting i5/OS partition via a TCP/IP-based application. Access to the console is limited to DST user ids that were granted “remote panel key authority”. This section provides the instructions for defining the DST user and accessing the virtual console.

Use the following steps to Create the Dedicated System Server Tools user with the correct authorities:

1. DST users are created through System Server Tools. Type the command **STRSST** to start System Server Tools. When the System Service Tools menu is displayed, select option 8 (Work with service tools user IDs and Devices).
2. When the Work with Service Tools User IDs and Devices menu is displayed, select option 1 (Service Tools user IDs).
3. When the Work with Service Tools User IDs window is displayed, enter option 1 to create a new user (Figure 13-30).

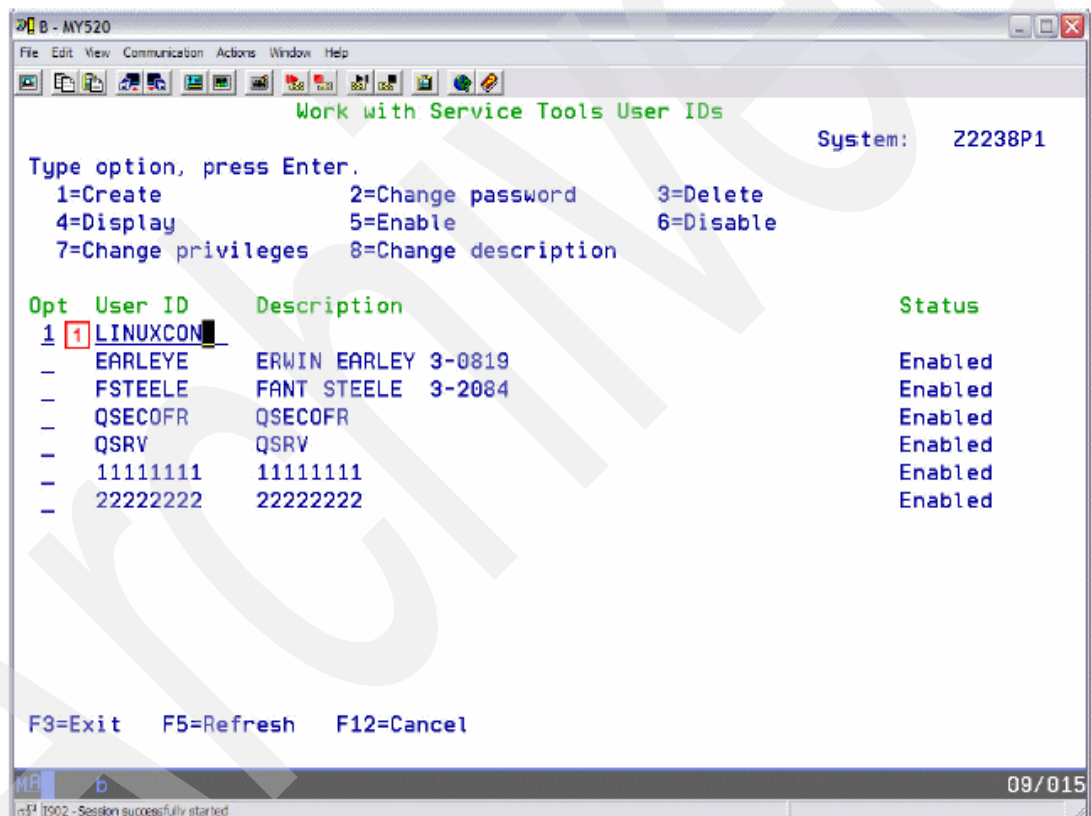


Figure 13-30 Work with Service Tools IDs

- **User ID:** The user-id being created for virtual console access.

4. Press Enter. The Create Service Tools User ID screen appears, as shown in Figure 13-31.




Figure 13-31 Create Service Tools ID

- **Password:** Type the password for the user-id being created. This is the password used to access the virtual console.
 - **Set password to expire:** Type 2 to indicate that the password should not be set to expired.
5. Press Enter to complete the DST user definition.
 6. After you create the DST user, modify the authorities for the user to include the remote panel key authority.

- On the Work with Server Tools User IDs, select option 7 (Change Privileges) for the user just created as seen in Figure 13-32.

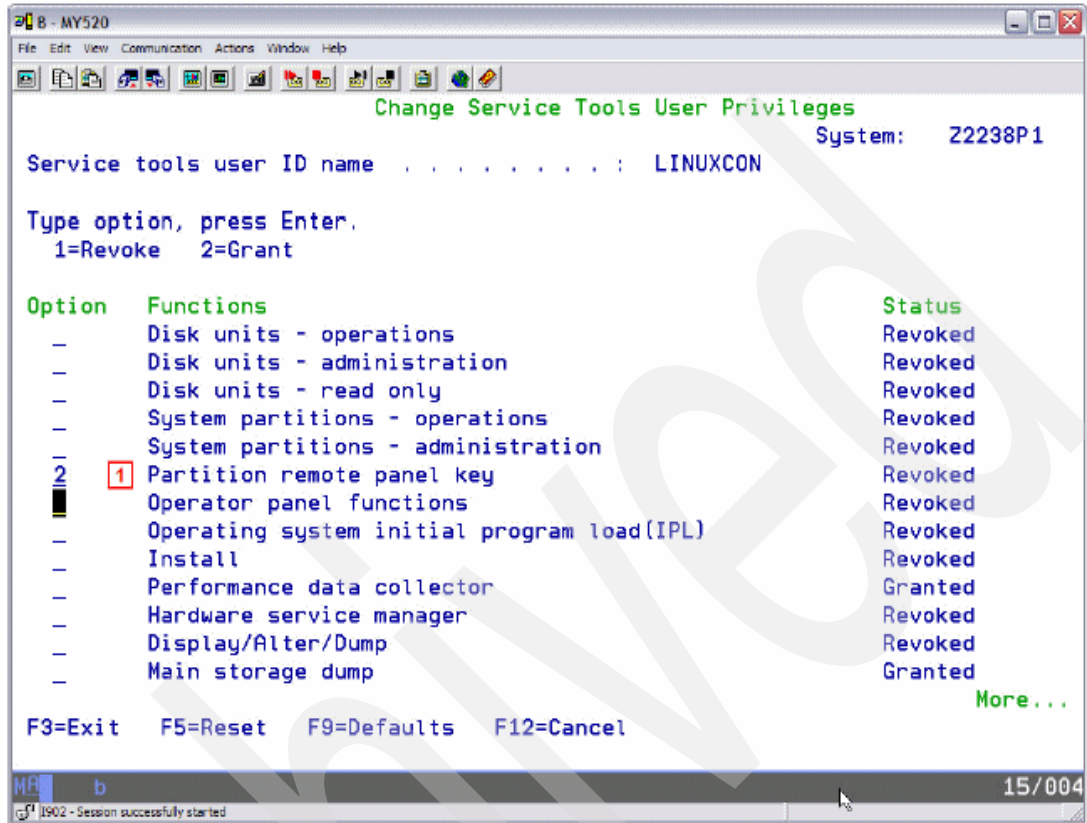


Figure 13-32 Change Service Tools User Privileges

- Partition remote panel key:** This is the authority that needs to be granted for virtual console access.

Note: In addition to the Partition remote panel key authority, the user id also requires “System partitions—operations” and “System partitions—administration” authority.

- After the required authority has been granted, press Enter.
- Press the F3 key.
- The DST is created and configured. Press F3 three times.
- Press Enter to exit out of System Service Tools.

13.7.3 Starting the Linux installation

At this point all of the components are defined for the Linux environment. You can start the Linux installation. This section provides the steps for accessing the virtual console and starting the Network Server.

- Prior to starting the Network Server, access to the virtual console needs to be achieved so that messages generated by the installer can be viewed. To accomplish this, use a telnet client, such as the windows telnet client, and access port 2301 of the i5/OS hosting partition. As an example, if the IP address of the hosting i5/OS partition is 192.168.10.10 then access to the virtual console is accomplished by entering the command ‘telnet 192.168.10.10 2301’.

2. A list of Linux partitions is provided. Type the number that corresponds to the partition for which you want to access the console.
3. When prompted for the OS/400 service tools user id, type the DST user that was created for virtual console access.
4. When prompted for the OS/400 service tools password, type the password defined for the DST user.
5. After the Virtual Console is accessed, the Network Server can be varied on. A vary on of the server is accomplished through the Work With Configuration Status command. Type the command `WRKCFGSTS *NWS` to work with the configuration of the Network Servers as seen in Figure 13-33.

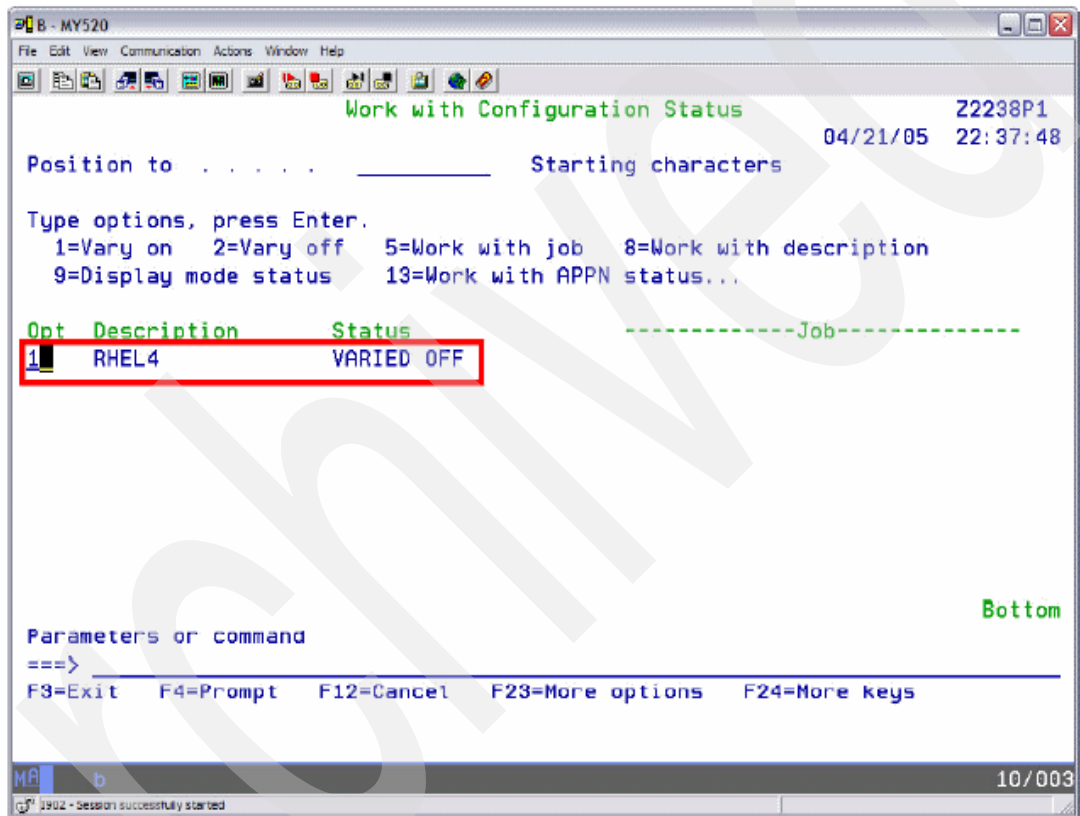


Figure 13-33 work with Configuration Status

6. An option '1' on the Network Server causes the network server to be varied on. As part of vary on process, a request is sent to power-on the partition defined in the Network Server Descriptor. After the partition starts, the Linux startup messages are displayed in the virtual console window.

Note: Detailed Linux installation instructions are provided in the distributor documentation.

13.8 Virtual medial management

This section covers virtual medial management.

13.8.1 Linux native backup with virtual tape

The OS/400 tape drive can be used by Linux for Linux-based save/restore of files and directories in a hosted partition.

Linux oriented backup has the same attributes as i5/OS oriented backup on file and directory level. The only difference is that the backup files are not saved to files in the NFS directory, but directly on tape. Tape is directly connected to i5/OS environment but can be used in Linux as virtual tape.

The following steps must be done to successfully save Linux files and directories on tape connected to an i5/OS partition:

1. Initialize tape for Linux. On i5/OS, the standard code for coding letters and numbers used is EBCDIC, whereas on Linux it is ASCII. This means that, in the scenario of Linux oriented backup, you need to strictly sort out tape cartridges for i5/OS and for Linux.

The following command is an example of initializing the tape cartridge for Linux:

```
INZTAP DEV(TAP02) NEWVOL(LNXBCK) CHECK(*NO) CODE(*ASCII) CLEAR(*YES)
```

2. Find the name convention for i5 tape in Linux. Use the `dmesg | grep st` command in `ssh PuTTY` to find the name of the tape you want to use for Linux backup. In our scenario we have the two tape devices, `st0` and `st1`, and the messages during boot regarding tapes are shown in Figure 13-34.

```
st: Version 20040318, fixed bufsize 32768, s/g segs 256
Attached scsi tape st0 at scsi0, channel 2, id 0, lun 0
st0: try direct i/o: yes (alignment 512 B), max page reachable by HBA 294912
Attached scsi tape st1 at scsi0, channel 2, id 1, lun 0
st1: try direct i/o: yes (alignment 512 B), max page reachable by HBA 294912
```

Figure 13-34 `dmesg` - tape messages during boot

3. Assume that the order of tape devices is the same in i5/OS and Linux and we want use TAP02 (not TAP01) in i5/OS terminology, meaning we need to find `st1` tape in Linux. Find where the `st1` is located in Linux. Use `ssh PuTTY` and the command:

```
find / -name st1
```

We receive the following output as shown in Figure 13-35.

```
rchas10d:~ # find / -name st1
/sys/class/scsi_tape/st1
/dev/st1
```

Figure 13-35 Find the tape location

Because `st1` is a device, the `/dev/st1` is the right path to tape.

4. Vary off tape on i5/OS. Use 5250 screen and command:

```
VRYCFG CFGOBJ(TAP02) CFGTYPE(*DEV) STATUS(*OFF) RANGE(*OBJ)
```

5. Use `ssh PuTTY` screen to check status of tape. Use the command:

```
mt -f /dev/st1 status
```

We receive the output as shown in Figure 13-36.

```
rchas10d:~ # mt -f /dev/st1 status
drive type = Generic SCSI-2 tape
drive status = 805306880
sense key error = 0
residue count = 0
file number = 0
block number = 0
Tape block size 512 bytes. Density code 0x30 (unknown).
Soft error count since last status=0
General status bits on (41010000):
BOT ONLINE IM_REP_EN
```

Figure 13-36 Tape status

Tip: When you receive error messages such as “mt: /dev/st1: Input/output error” try the following possibilities:

- ▶ Check at first to see if the tape cartridge isn't ejected.
- ▶ Use the command `sg_start /dev/sgx 1`

In the above expression, you can find the value of *x* by completing information from the `sg_map` command and the `dmesg | grep sg` and `dmesg | grep st` commands.

5. Type the following command, where *tape_device* is the name of the tape device you want to back up to:

```
backupios -tape tape_device
```

This command creates a bootable tape that can be used to restore the Virtual I/O Server.

Firmware maintenance

This chapter describes the various options available for maintaining both HMC and managed system firmware levels. We show you, through examples, some of the main firmware update options.

We discuss the different methods of updating the HMC to a new software level as well as installing individual fix packs. We also cover the backup of the HMC to help with the recovery process in the event of a disaster. Finally, we show the two ways of updating an i5 managed system's firmware, either through the HMC (out-of-band) or via an i5/OS service partition (in-band) with special MHxxxxx PTFs.

The following main topics are covered in this chapter:

- ▶ How to determine the HMC installed software level
- ▶ HMC backup of critical console data
- ▶ Installing an individual HMC fix
- ▶ Firmware overview

14.1 HMC firmware maintenance

The HMC software level has to be maintained just like the i5/OS operating system and managed system firmware (SP). HMC firmware is packaged as a full Recovery CD set or as a Corrective Service pack/fix image.

The HMC recovery CDs are bootable images and can be used to perform a complete recovery of the HMC (scratch install) or an update to an existing HMC version.

The HMC update packages are available on CDs or as downloadable zip files. The downloadable zip 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).

Important: An upgrade installation from the recovery CDs might be preferable to updating the HMC with the corrective service files for the following reasons:

- ▶ It is the only method that can be used to update the file system of HMC Version 4.2 and earlier from ext2 to ext3.
- ▶ It will reduce the size of the critical console backup file, allowing quicker backup operations. The critical console data backup saves the HMC user settings, profiles, partitions, etc. as well as any system updates since the last install or upgrade from the HMC recovery CDs. For more details, see 14.1.4, “HMC backup of critical console data” on page 427.

You should consult the IBM HMC support Web site on a regular basis to see if there any new updates available for the HMC. Use the link below to access the HMC support Web site.

<http://techsupport.services.ibm.com/server/hmc/System i5>

14.1.1 How to determine the HMC installed software level

There are various ways to display the current HMC software level depending on whether you are using the true HMC console, the Websm client, a restricted shell terminal, or an ssh client.

HMC console

From the true HMC console, select **Help** → **About Hardware Management Console** from the HMC desktop toolbar.

A panel similar to the one shown in Figure 14-1 is displayed. This screen shows the installed HMC version, release, and build. Notice that any individual fixes that have been applied to the HMC are also shown here.

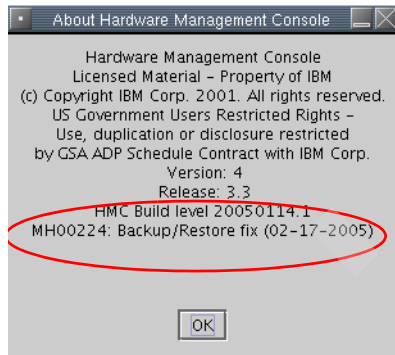


Figure 14-1 HMC console software level

HMC Websm Client

From the HMC Websm client, select **Licensed Internal Code Maintenance** → **HMC Code Update** in the HMC Navigation area. A screen similar to the one shown in Figure 14-2 is displayed.

Important: You cannot see if any individual HMC fixes have been applied to the HMC through the Websm client. You can only see the version and release level of the HMC.

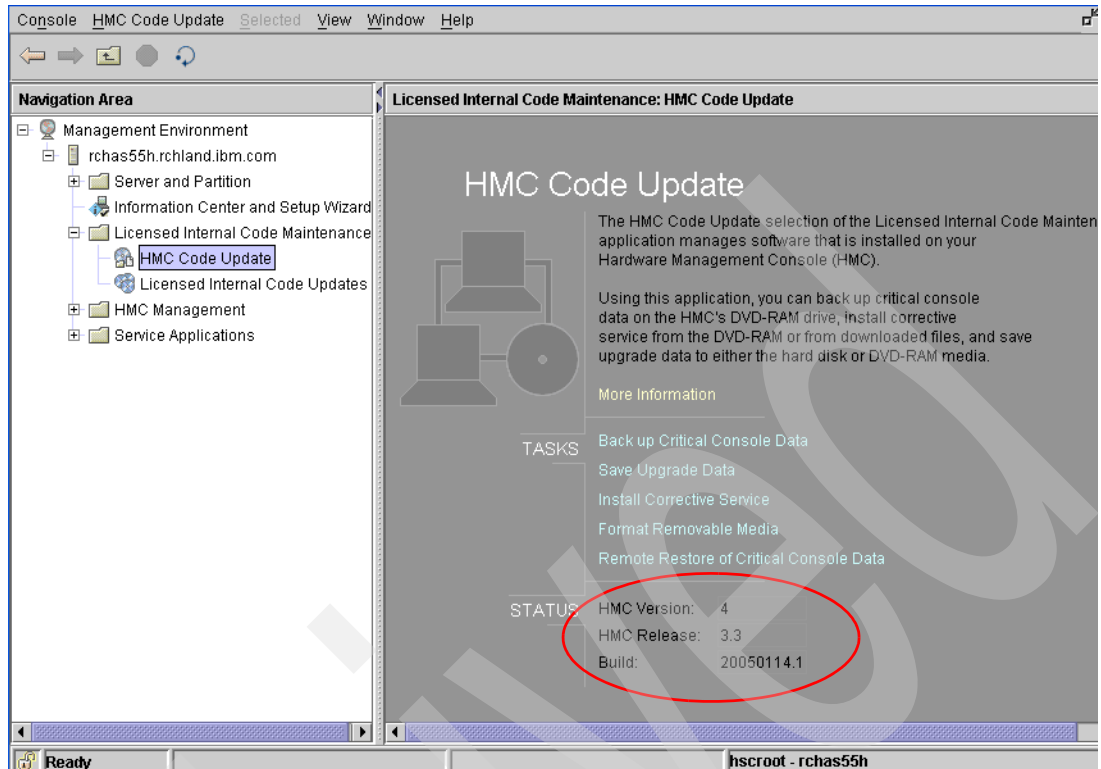


Figure 14-2 Websm HMC software level

HMC ssh client/restricted shell terminal

You can start an `ssh` client to the HMC (see “Scheduling the DLPAR function from Windows” on page 368) or a restricted shell terminal (see “Initial tour of the desktop” on page 56). By using the `lshmc -V` command, we can see the installed HMC software level. See Figure 14-3.

```

login as: hscroot
Sent username "hscroot"
hscroot@rchas55h's password:
Last login: Mon Feb 21 11:50:36 2005 from console
hscroot@rchas55h:~> lshmc -V
"version= Version: 4
  Release: 3.3
HMC Build level 20050114.1
MH00224: Backup/Restore fix (02-17-2005)"
hscroot@rchas55h:~>

```

Figure 14-3 ssh client - lshmc -V

14.1.2 Obtaining HMC update software

This section looks at the different possibilities of obtaining the HMC firmware updates.

You should always consult the HMC support Web site when planning to update the HMC. There may be a preferred method to update the HMC depending on the functions contained in the newer release. The following link can be used to access the HMC support Web site:

<http://techsupport.services.ibm.com/server/hmc/System i5>

Figure 14-4 shows the HMC support Web site for HMC version 4 release 4.

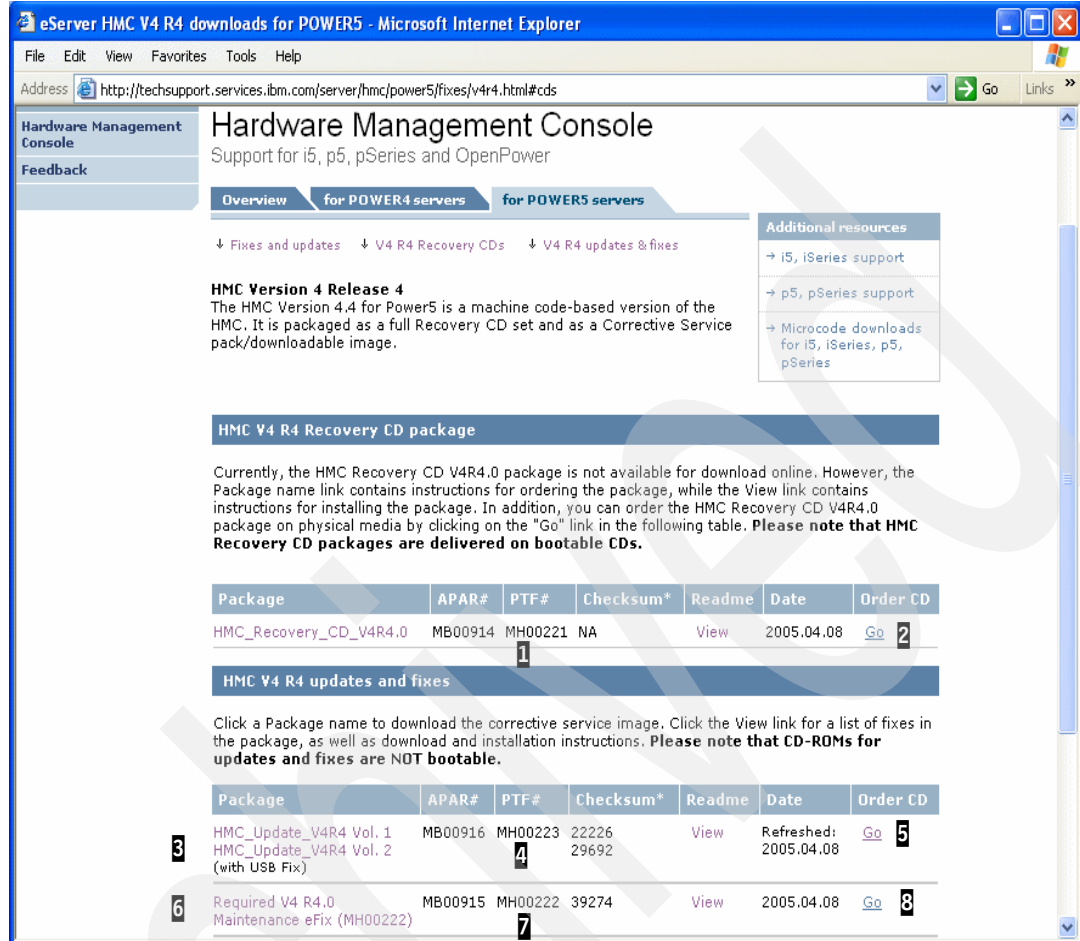


Figure 14-4 HMC support Web site

Obtaining the HMC recovery CDs

From Figure 14-4 you can see that the HMC recovery CDs are not available to download and can only be ordered from IBM and sent via post. You can either order the PTF number **1** MH00221 (this PTF number is unique to V4R4 and will change with each new release of the HMC recovery CDs) by using the normal i5/OS command SNDPTFORD, Fix Central or by clicking the **Go** link **2** and ordering through the Internet. You must have a registered IBM ID to order the HMC recovery CDs through the Internet.



Obtaining the HMC Update packages

If your HMC has a VPN connection to the Internet, you may choose to perform the HMC update directly from the IBM support FTP server. It is important to understand that some of the HMC update packages are large (over 1GB) and may take time to download. You may wish to download these update packages to one of your company's own internal FTP servers first and then perform the HMC update from this internal server. As shown in Figure 14-4, by clicking the HMC update package **3**, you can download these packages to your PC and place them on your company's own internal FTP server. You can then decide to either install these packages via a FTP server or by burning the zip files to CD-R or DVD-R/DVD-RAM media.


Important: You must NOT unzip the HMC update files when burning to external media, The HMC will itself, unpack these files during the install process.

You can also use the external IBM FTP server to download the HMC_Update_VxRyM0_n.zip files to one of your company's servers (such as i5/OS IFS). The external FTP server site is:

<ftp://techsupport.services.ibm.com/eserver/pseries/hmc/fixes/>



If your company has limited access or bandwidth to the Internet for FTP, you may decide to order the HMC update packages on CD format and have them sent to you via the post. You can order the PTF number MH00223  (this PTF number is unique to V4R4 and will change with each new release of the HMC update packages) by using the normal i5/OS command SNDPTFORD, Fix Central, or by clicking the **Go** link  and ordering through the Internet. You must have a registered IBM ID to order the HMC update packages through the Internet.

Obtaining HMC individual fixes

If your HMC has a VPN connection to the Internet, you may choose to install the HMC fix directly from the IBM support FTP server, or you can also download the fix packages to one of your company's own internal FTP servers first and then perform the HMC update from this internal server. As shown in Figure 14-4, by clicking the HMC fix , you can download the fix package to your PC and place them on your company's own internal FTP server. You can then decide to either install these packages via the FTP server or by burning the zip file to CD-R or DVD-R/DVD-RAM media.

You can also use the external IBM FTP server to download the MHxxxx.zip files to one of your company's servers (such as i5/OS IFS). The external FTP server site is:

<ftp://techsupport.services.ibm.com/eserver/pseries/hmc/fixes/>

If you company has limited access or bandwidth to the Internet for FTP, you may decide to order the HMC fix package on CD format and have them sent to you via the post. You can order the PTF number MH00222  (this PTF number is unique to V4R4 and will change with each new release of the HMC update packages) by using the normal i5/OS command SNDPTFORD or by clicking the **Go** link  and ordering through the Internet. You must have a registered IBM ID to order the HMC fixes through the Internet.

14.1.3 Changing the i5/OS default FTP attributes

If you decide to use an i5/OS partition as your FTP server when installing HMC fixes, you may need to change the i5/OS default FTP settings. The HMC requires the i5/OS FTP server to the unix format.

Important: Changing the i5/OS FTP attributes is a global change and affects all users and any applications using the traditional i5/OS FTP naming convention of QSYS.LIB.

Enter the following command from the i5/OS command line:

```
CHGFTP NAMEFMT(*PATH) CURDIR(*HOMEDIR) LISTFMT(*UNIX)
```

Also, any data sent from the HMC will be stored in the i5/OS user profile used in the HMC FTP process. If you want to change the home directory for that profile, enter the following command on the i5/OS command line using the relevant i5/OS profile name and home directory:

```
CHGUSRPRF USRPRF('user profile name') HOMEDIR('/home/mydir')
```


14.1.4 HMC backup of critical console data

We recommend taking a backup of the HMCs Critical Console Data (CCD) before proceeding with the HMC upgrade process. This backup will save all the HMC user data, user preferences, any partition profile data backup on the HMC disk, and various log files.

Important: The backup CCD is not a complete save of the HMC, as it contains user data as well as any system updates/fixes since the last install or upgrade from the HMC recovery CDs. If you need to reinstall the HMC from scratch, you must use the original IBM HMC recovery CDs and then restore the CCD backup.

For example, Figure 14-5 shows two HMCs (A) and (B) that follow the two possible HMC upgrade paths.

HMC (A) was installed with the V4R4 recovery cd set and then upgraded to V4R5 using the V4R5 corrective service update packages. The next CCD backup will save all your user data as well as all the updates installed since V4R4. In this case the CCD backup would be over 1GB in size.

HMC (B) was installed with the V4R4 recovery CD set and then upgraded to V4R5 using the V4R5 recovery cd set. The next CCD backup will save all your user data as well as all the updates installed since the last install/upgrade from the recovery cd set (V4R5). In this case the CCD backup could be less than 100 MB in size.




Figure 14-5 Critical Console Data Backup

There are three methods of backing up the HMC CCD:

- ▶ Back up to DVD on the local system.
- ▶ Send the backup critical data to the remote site (FTP).
- ▶ Back up to the mounted remote system (NFS).

In this section we show how to back up the CCD to the local HMC DVD drive and to a remote i5/OS FTP server.

Important: If you need to recover your HMC CCD from a FTP server or NFS during a HMC scratch install, you will need to reconfigure the HMC network settings before you are able to connect to the relevant remote system. Saving to DVD-RAM eliminates this step.

Back up Critical Console Data to DVD on local system

There is a DVD drive supplied with the HMC (eServer BACKUP DVD) which can be used to back up the HMC CCD. The only compatible DVD format for writing to, is DVD-RAM.

The DVD must be formatted in the DVD-RAM format before data can be saved to the DVD.

Important: The CCD backup to DVD media may be a long running process, depending on the amount of user data and system updates/fixes to be saved. In our experience this can be anything from a few minutes to many hours.

To format the DVD-RAM media

The following steps show how to format the DVD-RAM disk:

1. Place a DVD-RAM disk in to the HMC DVD drive.
2. In the HMC Navigation area, under your managed system, click **Licensed Internal Code Maintenance**.
3. Then click **HMC Code Update**.
4. In the right-hand window, click **Format Removable Media**.
5. Then select the **Format DVD** radio button.
6. Select **Backup/restore**.
7. Then click the **Format** button.

The DVD-RAM disk should be formatted in a few seconds, after which you will receive a "Format DVD has been successfully completed - ACT0001F" message.

Back up to formatted DVD media

Use the following steps to back up the CCD to the formatted DVD media:

1. In the HMC Navigation area, click **Licensed Internal Code Maintenance**.
2. Then click the **HMC Code Update**.
3. In the right-hand window, click **Back up Critical Console Data**.
4. Select the **Back up to DVD on local system** radio button and click the **Next** button.
5. Enter some valid text in the description window and click **OK**.

Send back up critical data to remote site

Use the following steps to back up the CCD to a remote FTP server. In our example, we show how to back up the CCD to an i5/OS Integrated File System(IFS), although any remote FTP server (UNIX or Windows) could also be used.

The HMC creates a file on the FTP server (IFS) with the following syntax:

```
HMCBackup_ date.time.tgz
```

Important: The ability to send the CCD to an i5/OS FTP server requires the HMC to be running at V4R5 or later.

1. In the HMC Navigation area, click **Licensed Internal Code Maintenance**.
2. Then click **HMC Code Update**.
3. In the right-hand window, click **Back up Critical Console Data**.
4. Select the **Send back up critical data to remote site** radio button and click the **Next** button.
5. Enter the host name/IP address of the i5/OS system, along with a valid i5/OS user id and password. You may also enter a useful text description for the backup in the window provided. The backup file will be stored in the home directory of the i5/OS user id entered here. Figure 14-6 shows an example of this.

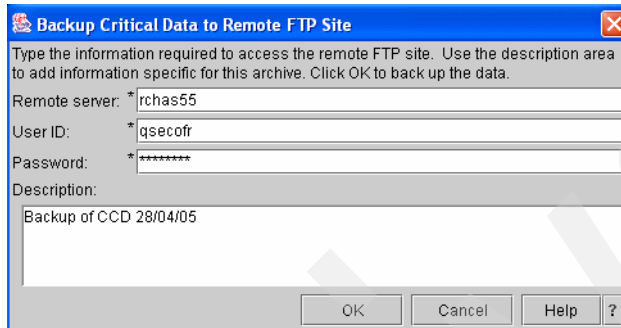


Figure 14-6 Backup of CCD to i5/OS

Click **OK** when you have completed the required fields.

6. The CCD backup first saves to the local HMC disk, before sending it to the FTP server. You receive a completion message when the save has finished.
7. You can view the CCD backup file in the i5/OS IFS by using the WRKLNK command. Figure 14-7 shows the WRKLNK for the QSECOFR home directory after the HMC CCD has completed.

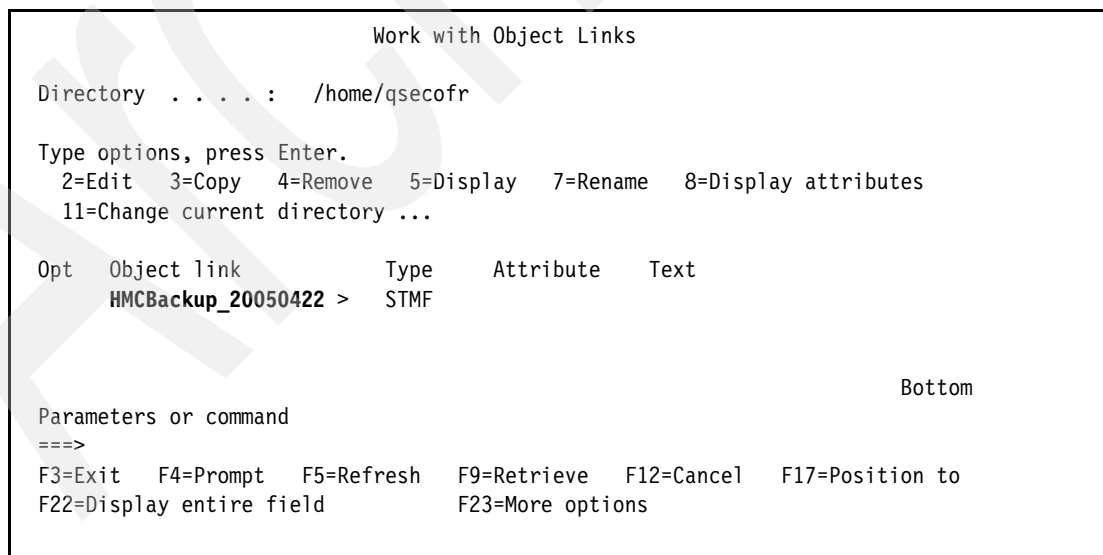


Figure 14-7 WRKLNK - HMC CCD Backup

14.1.5 Updating the HMC software

The media you are using to update the HMC will depend on the method used to upgrade your HMC. We show you two ways of upgrading the HMC. The first method will be using the HMC update packages from an FTP server. The second update method will be using the HMC recovery CDs.

Important: The examples in this section are based on the HMC upgrade to V4R4M0 and may change with future upgrades.

Updating the HMC software level from a FTP server

In our example, as our HMC is not connected to the Internet, we have already downloaded the relevant HMC_Update_VxRyMz_n.zip files to our i5/OS partition and stored them in the /home/qsecofr directory in the IFS (Figure 14-8).

```
Work with Object Links

Directory . . . . . : /home/qsecofr

Type options, press Enter.
  2=Edit  3=Copy  4=Remove  5=Display  7=Rename  8=Display attributes
  11=Change current directory ...

Opt  Object link      Type  Attribute  Text
-----
HMC_Update_V4R4.0_ > STMF
HMC_Update_V4R4.0_ > STMF

Bottom

Parameters or command
===>
F3=Exit  F4=Prompt  F5=Refresh  F9=Retrieve  F12=Cancel  F17=Position to
F22=Display entire field  F23=More options
```

Figure 14-8 WRKLNK for HMC update files

Important: We recommend that you have a backup of the critical console data before proceeding with the HMC updates.

The next steps show how to update the HMC from an i5/OS partition. These steps can be performed from the physical HMC or through the Websm client:

1. In the HMC Navigation area, click **Licensed Internal Code Maintenance**.
2. Then click **HMC Code Update**.
3. In the right-hand window click **Install Corrective Service**.
4. The Install Corrective Service panel appears (Figure 14-9). You need to select the second radio button (Download the corrective service file from a remote system) and fill in the supplied fields.
 - **Remote site** - This is the FTP server *hostname* or *IP* address.
 - **Patch file** - Enter the name of the relevant HMC_Update_VxRyMz_n.zip file.
 - **User ID** - The user profile used to connect to the remote system.
 - **Password** - The user profiles password used to connect to the remote system.

In our example we enter the information (also shown in Figure 14-9):

Remote Site - rchas55 (this is our i5/OS partition name).

Patch file - /HMC_Update_V4R4M0_1.zip (this is the first HMC update file in the IFS).

User ID - qsecofr (the i5/OS user profile used to access the i5/OS partition).

Password - xxxxxxxx (the i5/OS user profiles password).




Figure 14-9 Install Corrective Service panel

When you have completed all the fields with the correct information, click **OK** to continue.

5. The HMC working panel appears (Figure 14-10) showing the status of the install process. The HMC update data is inflated and then installed on the HMC.

The size of the HMC update zip file will depend on how long the install process takes. In our example, the install of the first HMC package took around 20 minutes.

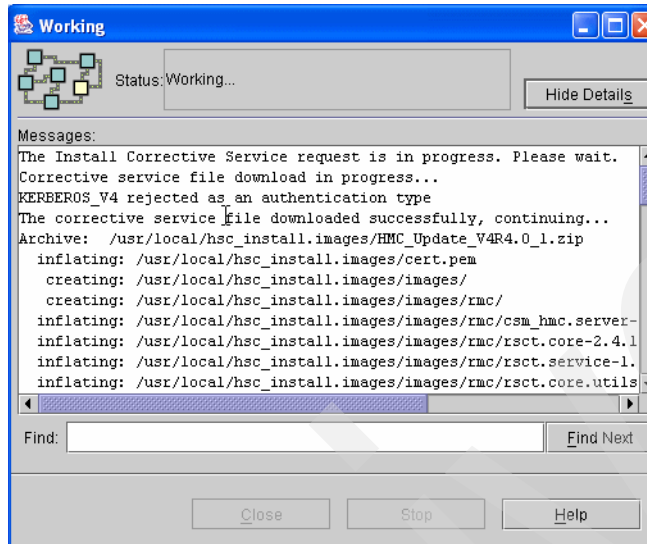


Figure 14-10 Install of HMC update via FTP server

6. When the installation of the first HMC update package has completed, you will receive a completion message (Figure 14-11)




Figure 14-11 HMC Install Corrective Service completion message

Click the selection box to automatically reboot the HMC when you select **OK**. You must reboot the HMC before installing the second HMC update package.

7. Next we repeat steps 4 through to 6 for the second HMC update package - HMC_Update_V4R4M0_2.zip.
8. This completes the update process for the HMC.

Updating the HMC software level from a HMC Recovery CD set

If you have received the HMC Recovery CD set, then you can upgrade your HMC release using these CDs.

When we update the HMC with the Recovery CD set, we are in fact replacing all the HMC data on the HMCs disk. To ensure that we keep all of our user data, such as partition/profile data, user profiles, user preferences, network configuration, etc., we must perform the Save Upgrade Data task on the HMC immediately before the HMC update. The Save Upgrade Data task is used to save the current HMC configuration to a special disk partition on the HMC. This data is then reloaded automatically during the update process.

Important: The HMC Backup Critical Console Data task is not the same as the Save Upgrade Data task.

The Save Upgrade Data task is only used during a HMC software update from the Recovery CD set.

The Backup Critical Console Data task is used to save the HMC configuration to removal media or a remote system. This backup CAN NOT be used during the HMC update. However, we recommend that you take a Critical Console Data backup before doing any software updates on the HMC, in case the update process fails.

The next steps show how to update the HMC with the Recovery CD set:

1. Save the current HMC configuration by following these steps:
 - a. In the HMC Navigation area, click **Licensed Internal Code Maintenance**.
 - b. Then select **HMC Code Update**.
 - c. In the right-hand window click **Save Upgrade Data**.
 - d. An information window opens (Figure 14-12) and prompts you to select the media type for the backup.

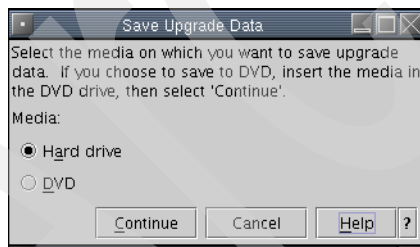


Figure 14-12 Save Upgrade Data media selection

Select the **Hard drive** radio button and click **Continue**.

Important: The Save Upgrade Data to DVD media can only be used by IBM service personnel.

- e. A confirmation window appears (Figure 14-13) before the save upgrade data is saved to the HMC disk.



Figure 14-13 Confirm Save Upgrade Data window

Click **Continue** to proceed with the save.

- f. When the save is complete, an information window opens - see Figure 14-14.




Figure 14-14 Save of Upgrade Data completed

2. Next shut down and power off the HMC.
3. Insert the first HMC recovery CD in to the HMC DVD drive and power on the HMC console.
4. The HMC boots from the DVD drive and displays the HMC Hard Disk Install/Upgrade screen as shown in Figure 14-15.

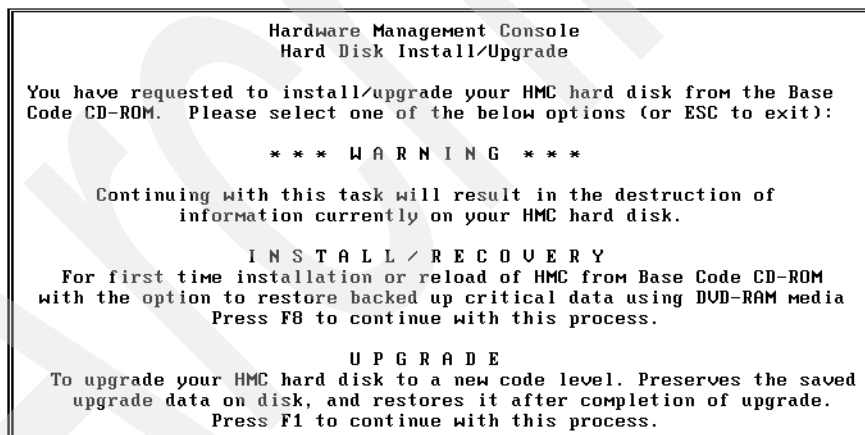


Figure 14-15 HMC Recovery cd set upgrade

Press **F1** to select the UPGRADE option.

5. A second HMC Hard Disk Upgrade screen is displayed (Figure 14-16). This screen explains the upgrade process and states that the Save Upgrade Data task must have been performed before continuing.

Important: You must NOT continue with the update process if you have not completed the HMC Save Upgrade Data task. If you upgrade the HMC without this save, all the HMC configuration and partition/profile data will be lost.




Figure 14-16 HMC Recovery cd set upgrade confirmation

Press **F1** to continue with the upgrade process.

6. An HMC progress screen starts, similar to the one in Figure 14-17.

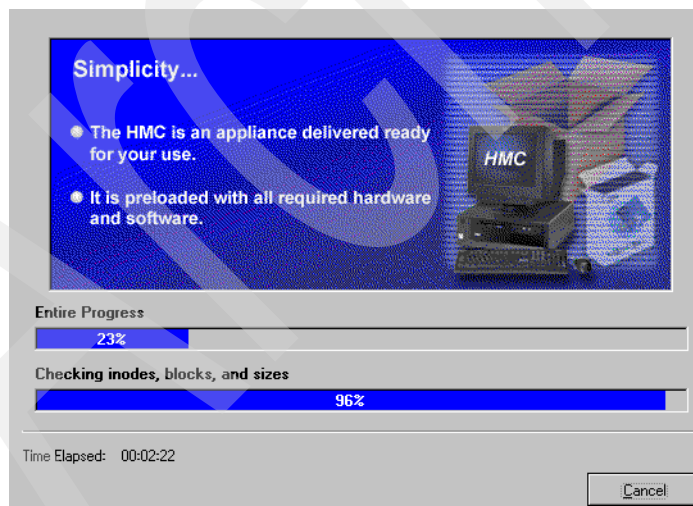


Figure 14-17 HMC update status

7. When the HMC has finished installing from the first CD, you are prompted to insert the second CD (Figure 14-18).




Figure 14-18 HMC upgrade - insert CD 2

8. Remove the first recovery CD from the HMC DVD drive and insert the second recovery CD. Press any key to continue the HMC upgrade process. The HMC will reboot and then start installing from the second CD.
9. When the HMC has finished installing from the second CD, you are prompted to install the third HMC recovery CD (Figure 14-19).




Figure 14-19 HMC upgrade - insert CD 3

Remove the second recovery CD from the HMC DVD drive and insert the third recovery CD. Enter 1 and press Enter to continue with the HMC upgrade process.

10. When the HMC has finished installing from the third CD, you are prompted to install the fourth HMC recovery CD (Figure 14-20).

```
Please select one of the following options:

1 -- Install additional software from CD media
2 -- Finish the installation at a later time

Please enter your selection : 1
```

Figure 14-20 HMC upgrade - insert CD 4

Remove the third recovery CD from the HMC DVD drive and insert the fourth recovery CD. Type **1** and press Enter to continue with the HMC upgrade process.

11. When the HMC has finished installing from the fourth CD you are prompted to either Restore the Critical Console Data from DVD or finish the HMC installation (Figure 14-21).

Important: You must NEVER insert the CCD backup DVD during the upgrade installation process.

```
Please select one of the following options:

1 -- Restore critical console data from DVD media
    Only select this option if you are performing a Hard Disk Recovery.
2 -- Finish the installation

Please enter your selection : 2
```

Figure 14-21 HMC upgrade - Finish installation

Remove the fourth HMC recovery CD media from the drive and do not insert any other CDs. Type **2** and press Enter to finish the HMC upgrade.

Note: You may have remaining CDs in the HMC Recovery CD set. These CD are normally other language translations of the Information Center and do not need to be installed at this time.

12. The HMC will reboot. At the login prompt enter a valid user ID and password and accept the License Agreement for Machine Code twice.

13. This completes the HMC code installation.

14. To verify the new HMC software level, you should use one of the methods in 14.1.1, "How to determine the HMC installed software level" on page 423

14.1.6 Installing an individual HMC fix

In our example, as our HMC is not connected to the Internet, we have already downloaded the relevant HMC fix file (MH00222.zip) to our i5/OS partition and have stored it in the /home/qsecofr directory in the IFS (Figure 14-22).

```
Work with Object Links

Directory . . . . . : /home/qsecofr

Type options, press Enter.
  2=Edit  3=Copy  4=Remove  5=Display  7=Rename  8=Display attributes
  11=Change current directory ...

Opt  Object link      Type  Attribute  Text
     MH00222.zip      STMF

Bottom
Parameters or command
===>
F3=Exit  F4=Prompt  F5=Refresh  F9=Retrieve  F12=Cancel  F17=Position to
F22=Display entire field  F23=More options
```

Figure 14-22 WRKLNK for HMC fix install

Important: We recommend that you have a backup of the critical console data before proceeding with the HMC updates.

The next steps show how to install an individual fix on the HMC from an i5/OS partition. These steps can be performed from the physical HMC or through the Websm client:

1. In the HMC Navigation area, click **Licensed Internal Code Maintenance**.
2. Then click **HMC Code Update**.
3. In the right-hand window click **Install Corrective Service**.
4. The Install Corrective Service panel appears (Figure 14-9). You need to select the second radio button (Download the corrective service file from a remote system) and fill in the supplied fields.
 - **Remote site** - This is the FTP server *hostname* or *IP* address.
 - **Patch file** - Enter the name of the relevant MHxxxx.zip file
 - **User ID** - The user profile used to connect to the remote system
 - **Password** - The user profiles password used to connect to the remote systemIn our example we enter the information (also shown in Figure 14-23):
 - **Remote Site** - rchas55 (this is our i5/OS partition name)
 - **Patch file** - MH00222.zip (this is the HMC update file in the IFS)
 - **User ID** - qsecofr (the i5/OS user profile used to access the i5/OS partition)
 - **Password** - xxxxxxxx (the i5/OS user profiles password)

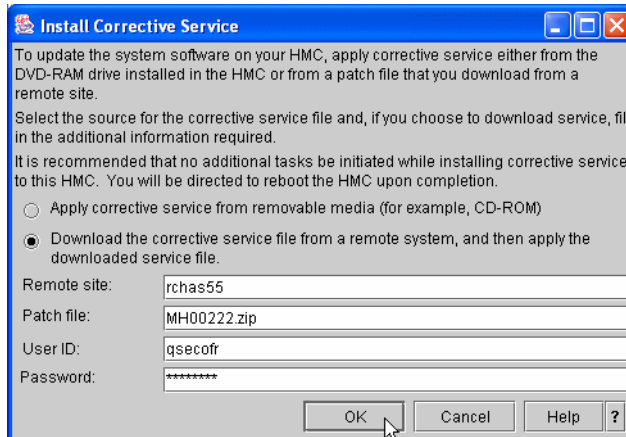


Figure 14-23 HMC fix install screen

When you have completed all the fields with the correct information, click **OK** to continue.

5. The HMC working panel appears showing the status of the install process (Figure 14-24). The HMC fix data is inflated and then installed on the HMC.

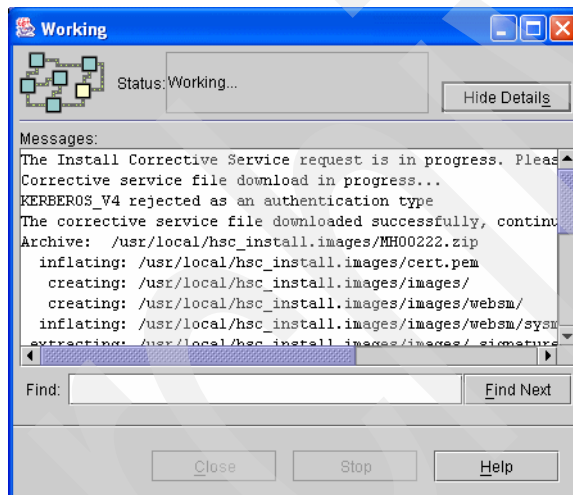


Figure 14-24 HMC fix install working screen

The size of the HMC fix zip file will depend on how long the install process takes. In our example, the install of the MH00222.zip package took around 5 minutes.

6. When the installation of HMC fix has completed you will receive a completion message (Figure 14-25)




Figure 14-25 HMC Install fix completion message

Click the selection box to automatically reboot the HMC when you select **OK**. You must reboot the HMC to activate any updates contained in the MHxxxx.zip fix.

7. This completes the HMC code installation.
8. To verify the new HMC software level, you should refer to section 14.1.1, “How to determine the HMC installed software level” on page 423

14.2 Licensed internal code updates

This section looks at the various methods used to manage and update an i5 managed systems firmware.

14.2.1 Firmware overview

Firmware refers to the underlying software running on an i5 system independently from any type of operating system (i5/OS, Linux or AIX). The firmware is physically located on the Service Processor hardware (similar to a processor card) on the i5 hardware.

Important: The server firmware can also be referred to as Licensed Internal Code. It is important to note that this is not the same as the Licensed Internal Code that resides in an i5/OS partition. OS/400 and i5/OS messages refer to the server firmware code as the Hypervisor.

The system firmware consists of the following main components:

- ▶ Service Processor
- ▶ Converged Power PC® Hypervisor (PHYP)
- ▶ SPCN/Power subsystem
- ▶ PFW (partition firmware Linux/AIX)

Depending on your system model and service environment, you can download, install, and manage your server firmware updates using different methods (policies), such as the HMC (out-of-band) or by using an i5/OS service partition (in-band).

Attention: 59x and 57x systems with redundant SPs can only use the HMC (out-of-band) to update the system firmware.

The default firmware update policy for a partitioned system is via the HMC. Both types of updates will be discussed in this section.

The SP maintains two copies of the server firmware. One is held in the t-side repository (temporary) and the other in the p-side repository (permanent). The idea being that new firmware updates can be applied to the t-side first and be tested before they are permanently applied. When you install server firmware updates on the t-side, the existing contents of the t-side are permanently installed on the p-side first.

If you decide that you do not want to keep the new level of server firmware, you can remove the current level of firmware. When you remove the current level of firmware, you copy the firmware level that is currently installed on the p-side from the p-side to the t-side.

We recommend that under normal operations the managed system runs on the t-side version of the system firmware. This concept of t-side/p-side is similar to the i5/OS system licensed internal code (SLIC), which has an a-side and b-side.

Starting with GA5 firmware, updates will be available in one of the following formats depending on the changes contained in the firmware update:

- ▶ **Concurrent install and activate:** Fixes can be applied without interrupting running partitions and restarting managed system.
- ▶ **Concurrent install with deferred disruptive activate:** Fixes can be applied as delayed and activated the next time the managed system is restarted.
- ▶ **Disruptive install with activate:** Fixes can only be applied by powering off the managed system.

Important: If your firmware update policy is set to operating system (in-band) then all your firmware updates will be disruptive - all partitions must be shutdown and the managed system restarted to apply the new firmware changes.

Figure 14-26 shows an overview of the firmware update process.

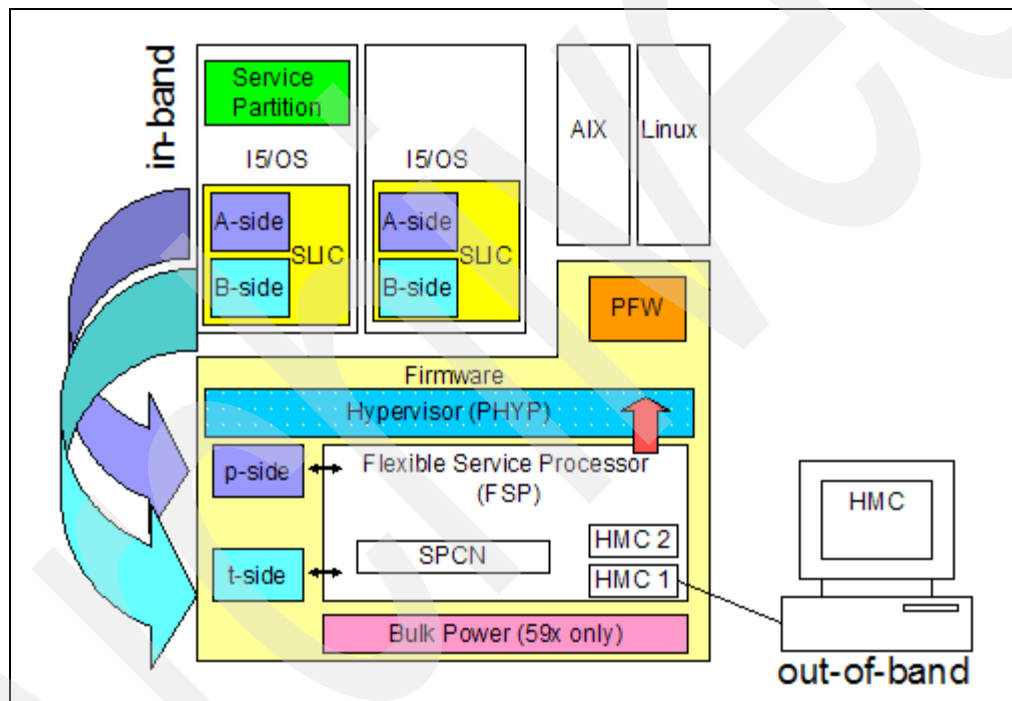


Figure 14-26 Firmware overview

In-band firmware overview

The firmware update policy must be set to operating system and an i5/OS partition must be defined as the service partition via the HMC.

New firmware updates/fixes are package as special MHxxxxx PTFs and can be ordered individually though Fix Central or via the i5/OS SNDPTFORD command. These MHxxxxx PTFs are also included in both Cumulative PTF packages and the HIPER Group package. You use the normal i5/OS PTF installation functions to apply the MHxxxxx PTFs on your service partition. When you apply these PTFs, the updates are placed in a repository on the load source disk. If the MHxxxxx PTF is applied temporarily, then the update is held on the (b-side) of the load source. If the MHxxxxx PTF is permanently applied, then the update is stored on the a-side of the load source disk.

When you shut down the i5/OS service partition, the b-side firmware is copied to the t-side on the SP. If the b and t-side are in sync, then the a-side will be copied to the p-side on the SP.

Important: Applying MHxxxx PTFs to a non-service i5/OS partition will have no impact on the firmware update process, as only the defined service partition has the authority to update the SP with any firmware updates.

The managed system must be restarted to activate any new firmware changes. A complete system reboot, where all partitions are ended is called a server IPL. When the managed system starts, the new firmware code is loaded in to memory from either the t or p-side, depending on the boot option chosen.

Out-of-band firmware overview

Managed system firmware updates can be made through the HMC. The firmware update policy should be set to Hardware Management Console (although you can update the managed system firmware regardless of this policy when the system is in a standby state).

New firmware updates can be downloaded directly to your HMC (if the HMC has access to the Internet). You can also download firmware updates to one of your company's internal FTP servers and install from there or burn to optical media. Lastly, you can order firmware updates on optical media, either from the IBM support Web site, through Fix Central or by using the i5/OS SNDPTFORD command.

Firmware updates follow the naming convention of SFXXX_YYY (where XXX = firmware release and YYY= version).

If you are updating the server firmware within the same version (for example, SF230_112 to SF230_113) then we use the HMC option **Change Licensed Internal code for the current release** (this option was called Change Internal Code prior to HMC V4R5)

If you are upgrading server firmware to a new version (for example, SF225_096 to SF230_113), then we use the HMC option **Upgrade Licensed Internal Code to a new release** (this option was called Manufacturing Equipment Specification Upgrade prior to HMC V4R5).

14.2.2 Setting the i5 system firmware update policy

The only way you can check or change the i5 system firmware policy is through the Advanced System Management Interface (ASMI) or the HMC command line.

Changing the firmware update policy with ASMI

The ASMI menus can be started through the physical HMC, Websm (as of HMC V4R5) or through a direct browser interface, if you managed systems SP is available on your ethernet network.

Important: If you change the update policy to allow firmware updates from the operating system, firmware updates from the HMC are not allowed unless the system is powered off.

When the managed system is powered off, firmware updates can be performed from the HMC regardless of the setting of this policy.

However, care should be taken when updating firmware from both the HMC and the operating system. If you change the update policy to operating system and an i5 service partition has a lower level of firmware installed than the firmware updated from the HMC, then the firmware will be overlaid with the older firmware from the i5/OS load source.

The following steps show how to start the ASMI through Websm:

1. Sign on to the HMC using the Websm client
1. In the HMC Navigation area, click **Service Applications**.
2. Then click **Service Focal Point**.
3. In the right-hand window click **Service Utilities**.
4. The Service Utilities panel is displayed (Figure 14-27). You must first select the Managed System you wish to connect to by clicking and highlighting that system **1**. Then click on the **Selected** tab on the drop down menu tool bar and click **Launch ASM Menu** **2**.

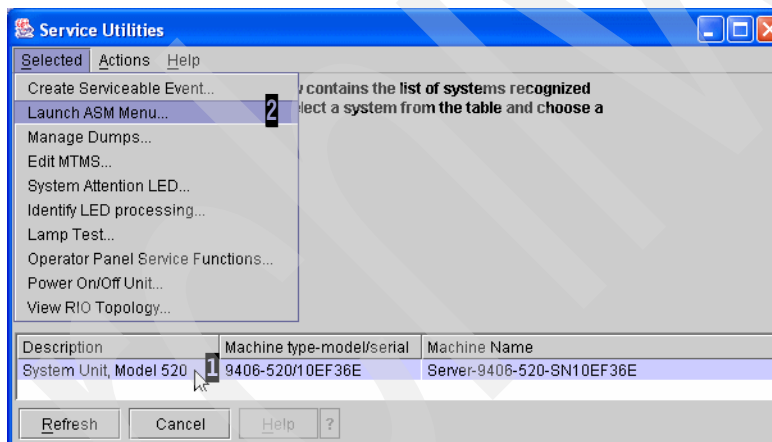


Figure 14-27 Service Utilities panel

5. The Launch ASM Interface screen is then shown (Figure 14-28).




Figure 14-28 Launch ASM Interface screen

Click **OK** to launch the ASM browser interface.

6. On the ASMI Welcome pane, specify your user ID and password, and click **Log In**. The default user ID is admin and password is also admin.

Important: For security reasons we recommend that the **admin** user ID password is changed from the default supplied password.

7. The main ASM screen is presented. In the navigation area, click **System Configuration** and select **Firmware Update Policy** (Figure 14-29).

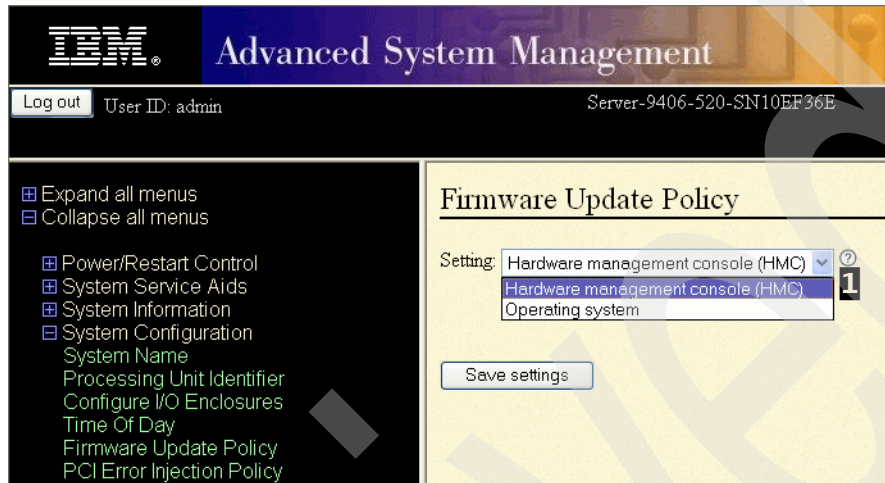


Figure 14-29 Firmware update policy screen

If you wish to change the firmware update policy, select the appropriate source from the drop down selection list **1** and click **Save settings** to complete the operation.

Changing the firmware update policy with HMC command line

You can change the firmware update policy from the HMC command line (HMC restricted shell or via `ssh`) by using the `updlic` command.

Use the following examples to change the firmware update policy using this command:

- To change the firmware update policy to Operating System, type:

```
updlic -m xxxx -o o
```

where `xxxx` is the name of the managed system.

- To change the firmware update policy to HMC, type:

```
updlic -m xxxx -o h
```

where `xxxx` is the name of the managed system.

Setting the service partition for i5 in-band firmware updates

The next steps show how to define an i5/OS service partition for the in-band firmware updates. These steps can be performed from the physical HMC or through the Websm client:

Attention: You cannot add or remove a service partition definition if that partition is in an active state.

1. In the HMC Navigation area, click **Server and Partition**.
2. Then click **Server Management**.
3. In the right-hand window select the managed system you wish to change, right click and select **Properties**.

- The Managed System Server Property Dialog window is shown (Figure 14-30). You can select any i5/OS partition to be the service partition, although it must be in an inactive state to be added or removed. Also, only one i5/OS partition can be the service partition at any given time.

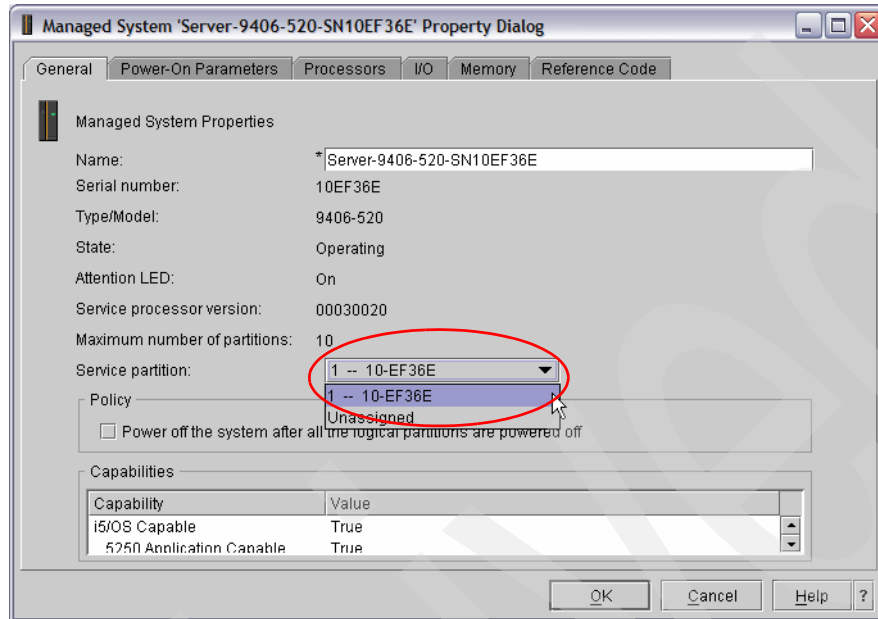


Figure 14-30 Set i5/OS service partition

- Select the new service partition from drop down menu and click **OK**. If the partition you are trying to add or remove is active you will receive message:
 HSCL0235 - Cannot change service partition from one partition to another if either one of them are not in the Not Activated states.

Verifying if an i5/OS partition is the service partition

You can check if an i5/OS partition is set to be the service partition from the output from DSPPTF i5/OS command. In Figure 14-31 we can see an additional information for the IPL source field. The ##SERV#T entry shows that this partition is the service partition and the system is running on the t-side of the managed system firmware. Only the service partition will show this information, all other i5/OS partitions will only show the normal ##MACH#B or ##MACH#A data.

```

Display PTF Status
System: XXXXXXXX
Product ID . . . . . : 5722999
IPL source . . . . . : ##MACH#B ##SERV#T
Release of base option . . . . . : V5R3M0 L00

Type options, press Enter.
5=Display PTF details 6=Print cover letter 8=Display cover letter

PTF                                IPL
Opt ID    Status                    Action
TL05032  Temporarily applied        None

```

Figure 14-31 DSPPTF command from an i5/OS service partition

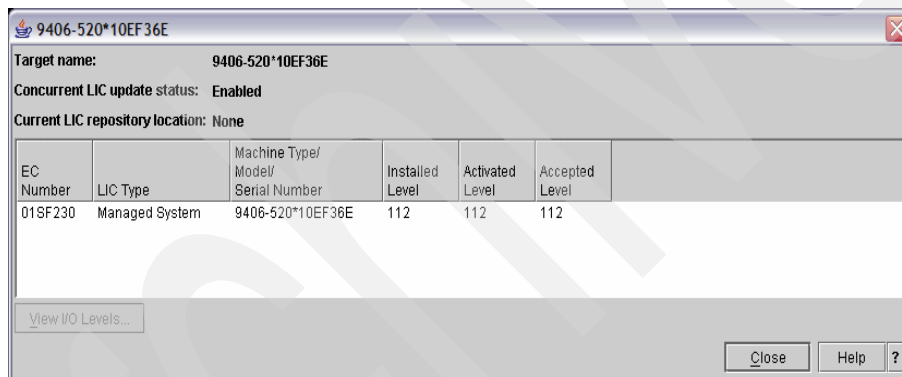
14.2.3 Displaying the current firmware levels

The installed firmware levels can be seen through both the HMC and i5/OS partitions. This section shows how you can use both methods to view the managed system firmware levels.

Using the HMC to display the current firmware levels

Use the following steps to display the managed system firmware levels:

1. In the HMC Navigation area, click **Licensed Internal Code Maintenance**.
2. Then click **Licensed Internal Code Updates**.
3. In the right-hand window click **Change Licensed Internal Code for the current release** (this option was called Change Internal Code prior to HMC V4R5).
4. In the Target Object Selection panel, highlight the relevant managed system and click **OK**.
5. From the Change Licensed Internal Code screen select the **View system information** radio button and click **OK**.
6. When Specify LIC repository panel is presented, select the **None** radio button and click **OK** to continue.
7. The firmware installed status panel is shown (Figure 14-32). This screen displays the different versions of firmware installed on the managed system.



EC Number	LIC Type	Machine Type/ Model/ Serial Number	Installed Level	Activated Level	Accepted Level
018F230	Managed System	9406-520*10EF36E	112	112	112

Figure 14-32 Current firmware levels installed from Websm

- ▶ **EC Number** - This is the engineering change that shows the system and GA level. It has the format of PPNNSSS where:
 - PP is the package identifier
 - NN is the model unique code for the type of system
 - SSS is the service pack code number, think of this as the release level of the firmware.
 - ▶ **Activated Level** - This is the version of firmware that is currently active (in memory) on the managed system.
 - ▶ **Installed Level** - This is the version of firmware that is currently installed on the *t-side* of the SP
 - ▶ **Accepted Level** - This is the version of firmware that is currently installed on the *p-side* of the SP.
8. Click **Close** return to the Licensed Internal Code main menu.

Using the i5/OS service partition to display firmware levels

You can use the following steps from an i5/OS 5250 screen to display the current installed firmware levels:

1. Enter the STRSST command on an i5/OS command line and enter a valid user ID and password.
2. Select option 1, Start a service tool, and press Enter.
3. Then select option 4, Display/Alter/Dump, and press Enter.
4. Next take option 1, Display/Alter storage, and press Enter.
5. Select option 2, Licensed Internal Code (LIC) data, and press Enter.
6. Then select option 14, Advanced analysis, and press Enter.
7. Page down and select FLASHLEVELS, and press Enter.
8. There are no special parameters to type, just press Enter.
9. The Display Formatted Data screen appears (Figure 14-33)

```

Display Formatted Data
Page/Line. . . . . 1 / 1
Columns. . . . . : 1 - 78
Find . . . . .
....+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....
DISPLAY/ALTER/DUMP
Running macro: FLASHLEVELS

LS Flash Sync Enabled.
HMC MANAGED
MiKeyword: NVRAM: SF225_096

Side          Date/Time      MI Keyword    PTFs
-----
Memory        20050415/0857  SF230_112   MH00261
Flash P       20050406/0941  SF230_112   MH00255
Flash T       20050415/0857  SF230_112   MH00261
Load Source A 20050311/1416  SF225_096    MH00248
Load Source B 20050311/1416  SF225_096    MH00248

Bottom

F2=Find  F3=Exit  F4=Top   F5=Bottom  F10=Right  F12=Cancel

```

Figure 14-33 Current firmware levels installed - shown from DST

Here we can see that our firmware policy is set to HMC managed. Also note that the firmware levels for the Load Source A/B and lower than that of the SP P/T. This is not a problem for a HMC managed firmware policy, as the load source levels are never used. However, if we decide to change our firmware update policy to operating system and set this i5/OS partition to be our service partition, then the firmware levels of the SP would be downgraded to SF225_096. In this situation we would recommend installing the relevant MHxxxxx PTF on the i5/OS partition to bring the firmware levels on the load source to the same or higher level before changing the firmware update policy.

14.2.4 Updating firmware through the HMC (out-of-band)

This section shows how to update the i5 firmware via the HMC.

Important: For all models except 59x model servers, we recommend that you install HMC fixes before you upgrade to a new server firmware release.

For 59x model servers, you must install HMC fixes before you upgrade to a new server or power subsystem firmware release.

Obtaining i5 out-of-band firmware updates

You should consult the IBM i5 recommended fixes Web site link below for information on firmware updates. The examples contained in this section are based on firmware levels available at the time of writing this redbook and may change with future releases:

http://www-912.ibm.com/s_dir/slkbase.nsf/recommendedfixes

Figure 14-34 shows an example of this fix support Web site:

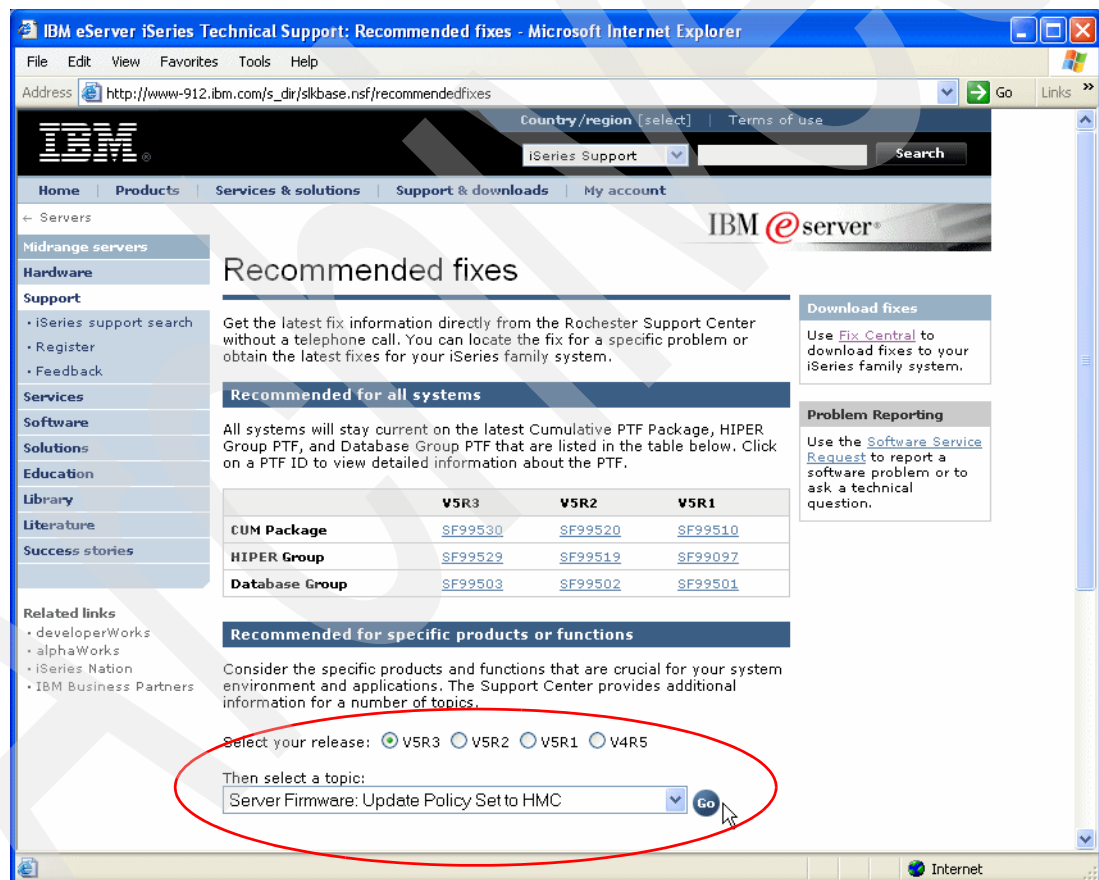


Figure 14-34 i5 recommended fixes Web site - out-of-band

You should select **Server Firmware: Update Policy Set to HMC** from the drop down topic window and click **Go**.

The next screen shown is the iSeries Recommended Fixes - Server Firmware: Update Policy Set to HMC Web page (Figure 14-35).

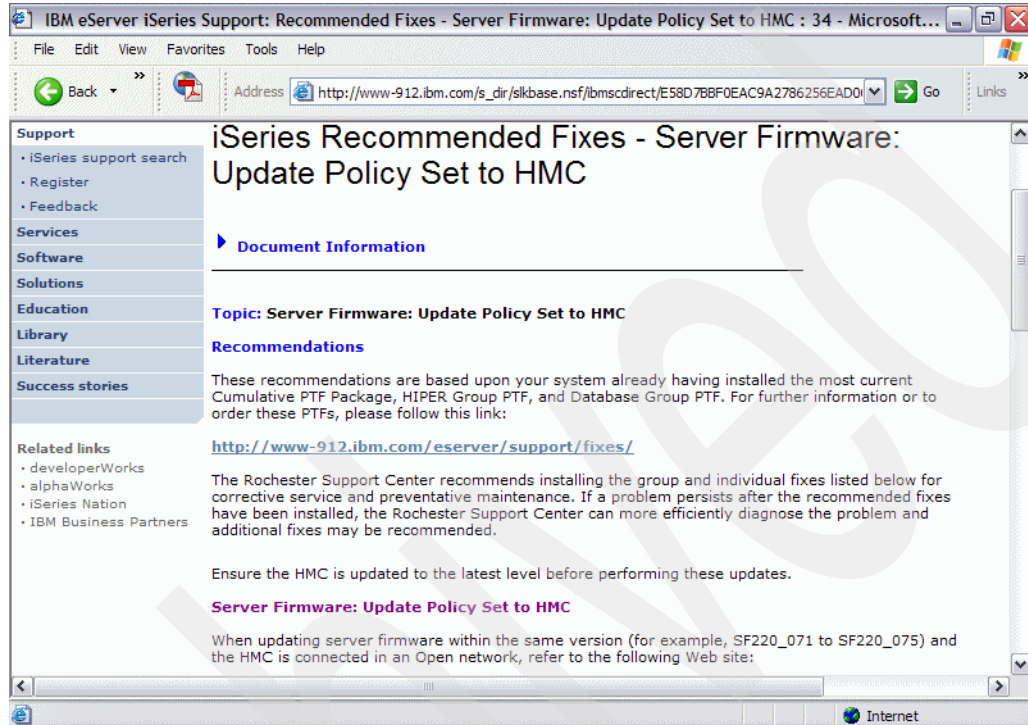


Figure 14-35 Server firmware - HMC update policy set to HMC

There are numerous ways of obtaining the i5 firmware which are explained in detail on this Web page. The main choices are:

- ▶ If the HMC has a VPN connection to the Internet, you may choose to download the new firmware directly to the HMC before installation.
- ▶ Download the firmware and burn to removable media (CD-R).
- ▶ Order the firmware on CD by using the SNDPTFORD i5/OS command or through Fix Central using the MHxxxx mark PTFs. For example, the SF225_096 firmware can be ordered on CD by using requesting PTF MH00254.

Installing i5 out-of-band firmware updates

In this section we show how to update the i5 firmware through the HMC, using a firmware update CD.

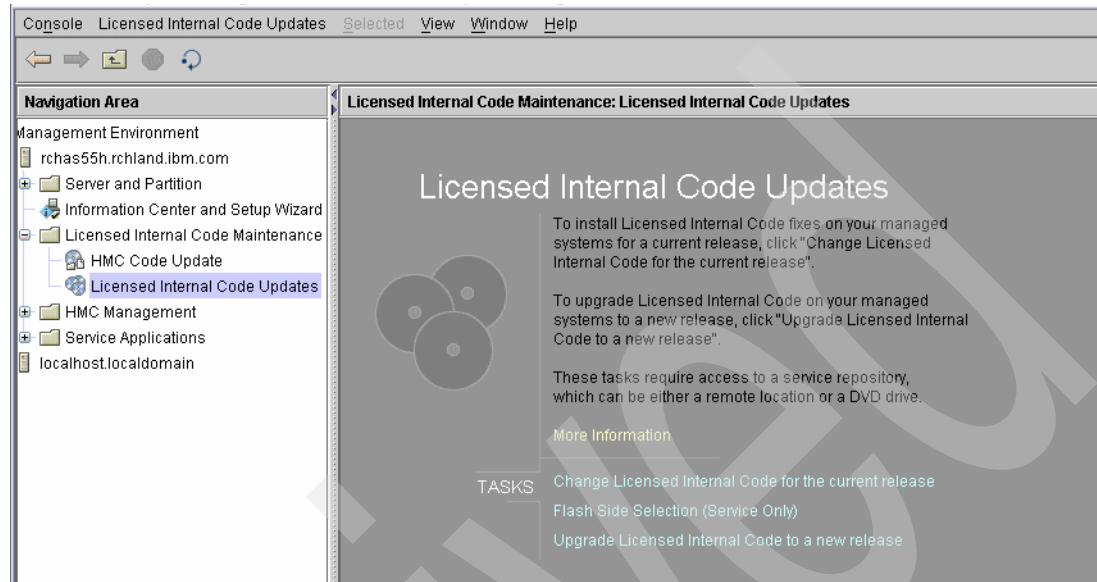


Figure 14-36 Licensed Internal Code Updates main menu screen

The method used to install a firmware update depends on the release level of firmware which is currently installed on your system and the release level you intend to install. The release level of the new firmware can be determined by the prefix of the new firmware levels filename.

Example: SFXXX_YYY (where XXX = firmware release and YYY= version)

If you are updating the server firmware within the same version (for example, SF230_112 to SF230_113) then select **Change Licensed Internal code for the current release** (this option was called Change Internal Code prior to HMC V4R5)

If you are upgrading server firmware to a new version (for example, SF225_096 to SF230_113), then select **Upgrade Licensed Internal Code to a new release** (this option was called Manufacturing Equipment Specification Upgrade prior to HMC V4R5).

Attention: All 'Upgrade Licensed Internal Code to a new release' changes are disruptive.

Change Licensed Internal code for the current release

This example shows how to update the manage system firmware from SF230_112 to SF230_113 using CD media:

1. Insert the firmware fix CD in to the DVD drive in the HMC.
2. In the HMC Navigation area, click **Licensed Internal Code Maintenance**.
3. Then click **Licensed Internal Code Updates**.
4. In the right-hand window click **Change Licensed Internal Code for the current release** (this option was called Change Internal Code prior to HMC V4R5).
5. In the Target Object Selection window, select the target system, and click **OK**.

6. In the Change Licensed Internal Code window, select **Start Change Licensed Internal Code wizard**, and click **OK to continue**.
7. In the Specify LIC Repository window, select the repository location from which you want to download/install the server firmware fixes, and click **OK**. The following options are available:
 - ▶ **IBM service Web site:** If your HMC has a VPN/modem connection to the Internet you can use this option to download the latest firmware fixes directly to the HMC.
 - ▶ **IBM support system:** If your HMC has a VPN/modem connection to the Internet you can use this option to download from all available levels of firmware directly to the HMC.
 - ▶ **DVD drive:** Use this option if your firmware updates are on CD media.
 - ▶ **FTP site:** Select this option if you previously downloaded the firmware fixes to an internal FTP server.
 - ▶ **Hard drive:** Use this option if you previously downloaded the firmware update the hard drive of the HMC.

In our example, we select **DVD drive** as our firmware update is on optical media and click **OK** to continue.
8. The Change Licensed Internal Code welcome wizard screen is presented. Click **Next** to continue with the update procedure.
9. The Change Licensed Internal Code select install updates is shown (Figure 14-37). This screen displays the type of concurrency installation to be performed.

Important: If the CD inserted contains a new release rather than a change for the current release, you will receive an information window stating “There is a new upgrade release available at the repository - HSCF0051”. Clicking the Cancel button on the window will start the firmware update process for a new release. These steps are listed in “Upgrade Licensed Internal Code to a new release” on page 454.

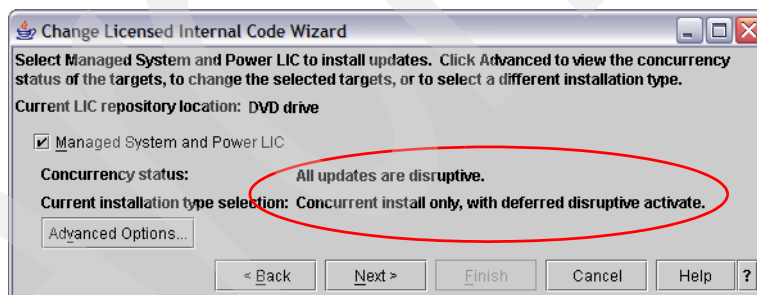


Figure 14-37 Change Licensed Internal Code Wizard screen

We can see that in our example, that the concurrency status shows all the updates are disruptive (this means that the managed system must be rebooted to activate the new changes). Also shown is the type of installation to be performed, we see that the firmware install will be concurrent with a deferred disruptive activate (this means that the new firmware can be installed immediately, with the managed system active, but the new changes will not be activated until the next system restart). This concept is similar to applying delayed i5/OS PTFs.

If you wish to change the type of firmware installation click the **Advanced Options** button. The Managed System and Power Licensed Internal Code (LIC) Concurrency window is presented (Figure 14-38).

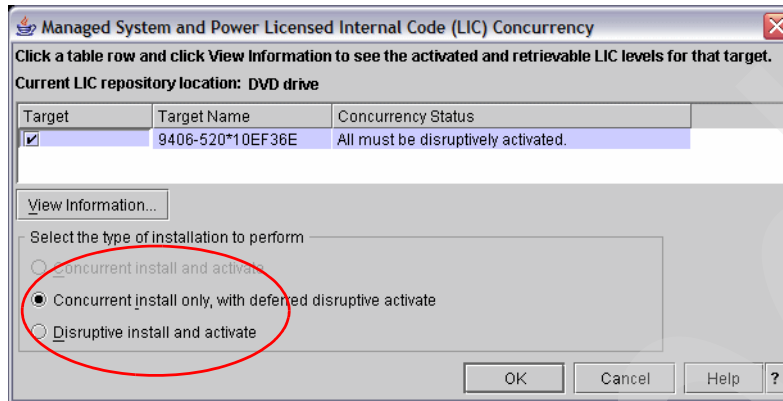


Figure 14-38 Licensed Internal Code Concurrency screen

The available options for the firmware installation are shown. We decide to leave the installation as Concurrent. If we changed the install to disruptive, then any running partitions would need to be ended and the managed system restarted when we start the update process.

If you wish to see the further details of the this installation click the **View Information** button. The Managed System and Power Internal Code (LIC) Level Details panel is opened (Figure 14-39).

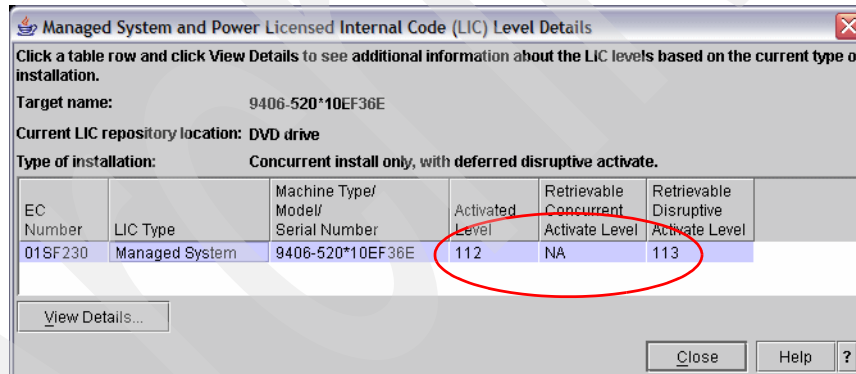


Figure 14-39 The Managed System and Power LIC Level Details panel

This panel shows the current running (activated level) firmware and the projected (retrievable level) new firmware and the type of installations available.

We click **close** and then **OK** to return back to the original screen shown Figure 14-37. Click **Next** to continue with the Change Licensed Internal Code wizard.

10. The Hardware Management Console License Agreement panel is then shown (Figure 14-40).

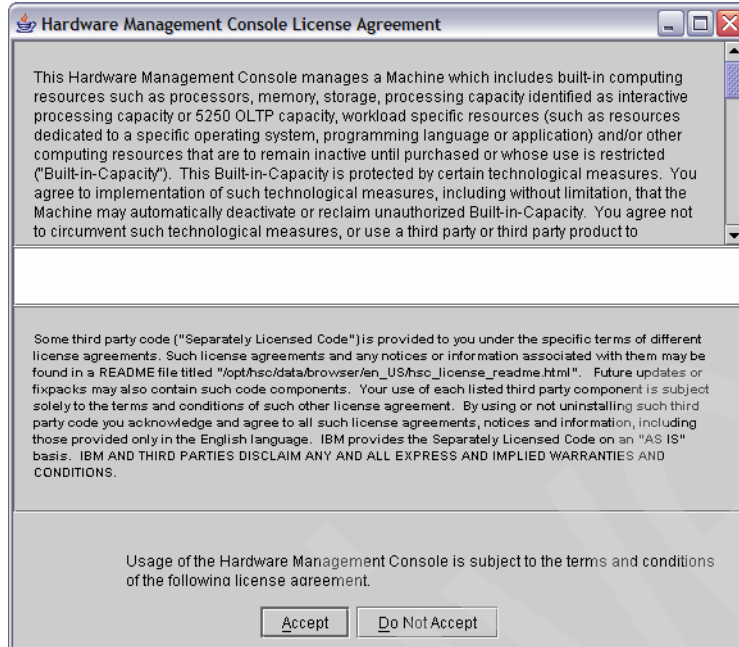


Figure 14-40 HMC License Agreement screen - update release

You should read the licensed agreement before clicking the **Accept** button.

11. The Change Licensed Internal Code Wizard Confirmation screen appears (Figure 14-41). This screen shows all the managed systems that will be updated and the type of update. You can use the View Levels button to see the level of firmware to be installed.

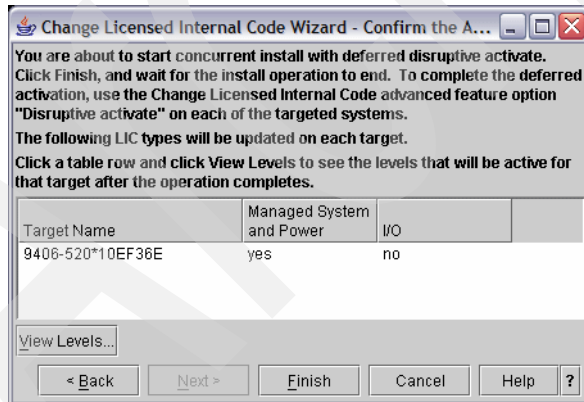


Figure 14-41 Change LIC Confirmation screen

In our example, we click **Finish** to start the installation of the firmware.

12. The Change Licensed Internal Code Wizard Progress window appears (Figure 14-42). When you install server firmware updates on the t-side, the existing contents of the t-side are permanently installed on the p-side first.

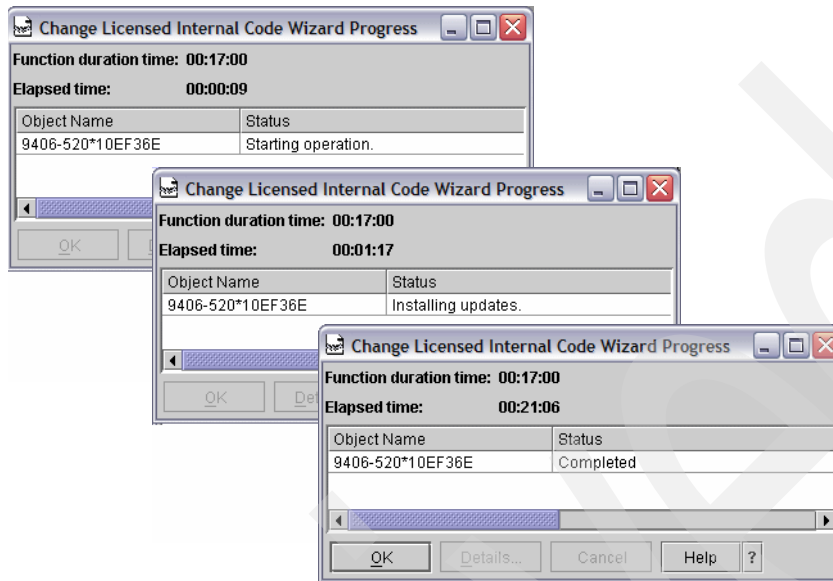


Figure 14-42 Change LIC Wizard - Starting Change LIC Wizard - status window

In our example the new firmware is installed after 20 minutes. To activate this new level of code, a complete restart of the managed system is required. After you have restarted the managed system, use one of the methods shown in Section 14.2.3, “Displaying the current firmware levels” on page 446 to see the new firmware levels.

This completes the installation of managed system firmware from the HMC.

Upgrade Licensed Internal Code to a new release

This example shows how to update the manage system firmware from SF230_096 to SF230_113 using CD media:

1. Insert the firmware release upgrade CD in to the DVD drive in the HMC.
2. In the HMC Navigation area, click **Licensed Internal Code Maintenance**.
3. Then click **Licensed Internal Code Updates**.
4. In the right-hand window click **Upgrade Licensed Internal Code to a new release** (this option was called Manufacturing Equipment Specification Upgrade prior to HMC V4R5).
5. In the Target Object Selection window, select the target system, and click **OK**.
6. In the Specify LIC Repository window, select the repository location from which you want to download/install the server firmware release upgrade, and click OK. The following options are available:
 - ▶ **IBM service Web site:** If your HMC has a VPN/modem connection to the Internet you can use this option to download the latest firmware fixes directly to the HMC.
 - ▶ **IBM support system:** If your HMC has a VPN/modem connection to the Internet you can use this option to download from all available levels of firmware directly to the HMC.
 - ▶ **DVD drive:** Use this option if your firmware updates are on CD media.
 - ▶ **FTP site:** Select this option if you previously downloaded the firmware fixes to an internal FTP server.

In our example, we select **DVD drive** as our firmware release upgrade is on optical media and click **OK** to continue.

- The Hardware Management Console License Agreement panel is then shown (Figure 14-43).

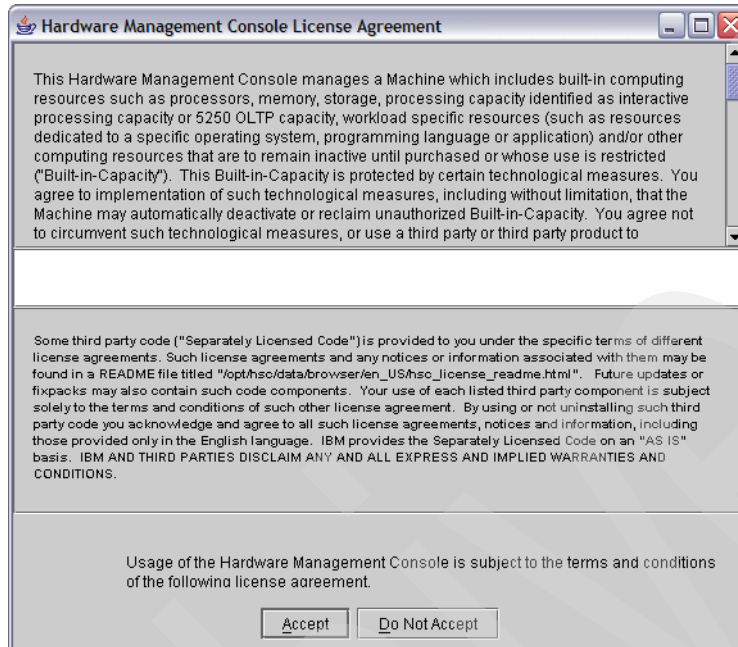


Figure 14-43 HMC License Agreement screen - upgrade release

You should read the licensed agreement before clicking the **Accept** button.

- The Upgrade LIC - Confirm the Action window appears (Figure 14-44). The current EC number (current installed firmware level) and the new EC number (new firmware level to be installed) are shown.

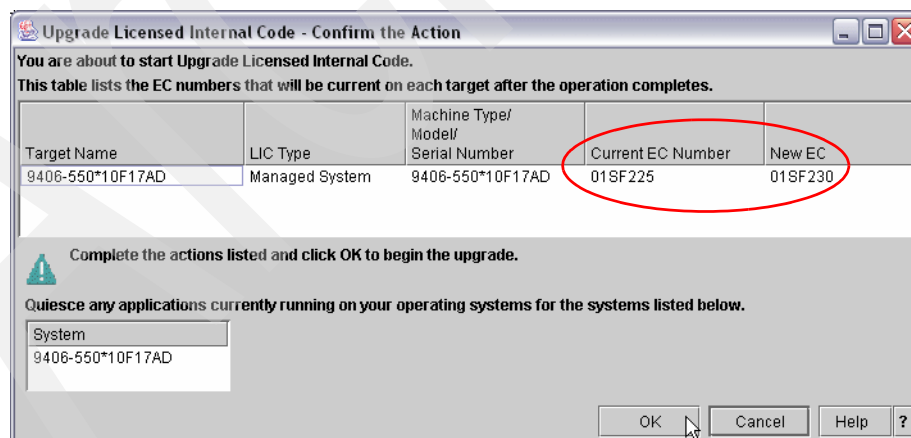


Figure 14-44 Upgrade LIC - Confirm the Action screen

In our example we see a current EC number of 01SF225 and a new EC number of 01SF230. We also see that the following message is displayed on the HMC window:

Quiesce any applications currently running on your operating systems for the systems listed below.

This message means that you will need to manually shut down all logical partitions on this managed system before continuing. If you do not power down these partitions, they will be shutdown abnormally during the firmware release upgrade.

Click **OK** to start the upgrade of the licensed internal code

9. The upgrade of the firmware starts, the status windows for the licensed internal code upgrade are shown below (Figure 14-45).

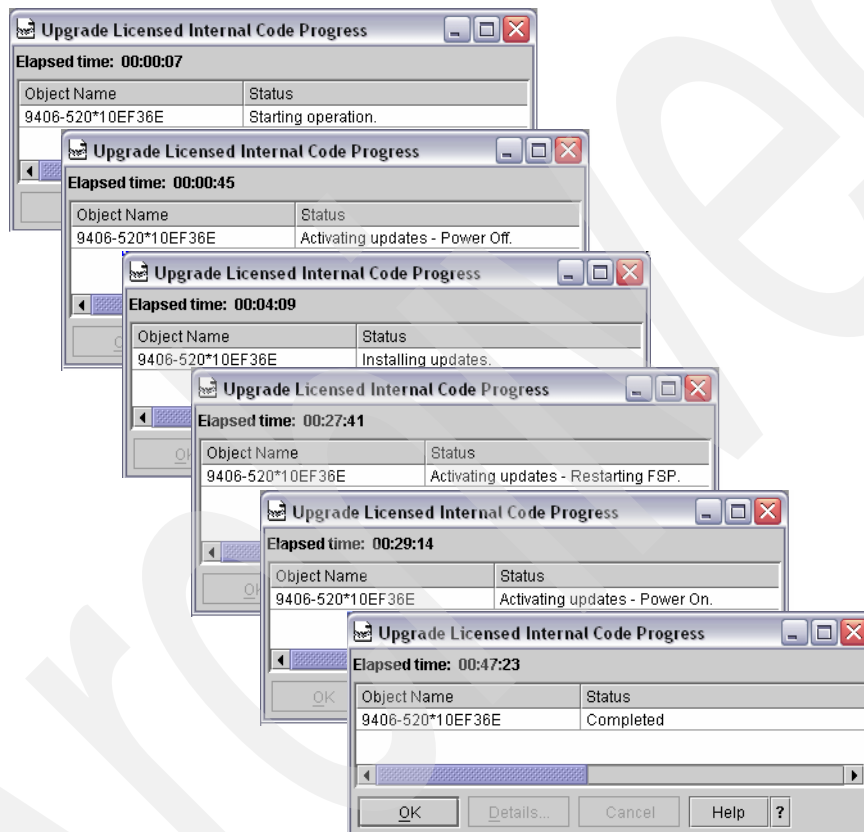


Figure 14-45 Upgrade LIC - progress windows

In our example the new firmware is installed after 47 minutes. You should verify the new firmware levels by using one of the methods shown in Section 14.2.3, “Displaying the current firmware levels” on page 446 to see the new firmware levels.

This completes the release upgrade of the managed system firmware from the HMC.

14.2.5 Updating firmware through an i5/OS service partition (in-band)

This section shows how to update the i5 firmware via an i5/OS service partition. The examples contained in this section are based on firmware levels available at the time of writing this redbook and may change with future releases.

Important: We recommend that you install HMC fixes before you upgrade to a new server firmware release.

Obtaining i5 in-band firmware updates

You should consult the IBM i5 recommended fixes Web site link below for information on firmware updates.

http://www-912.ibm.com/s_dir/slkbase.nsf/recommendedfixes

Figure 14-46 shows an example of this fix support Web site:

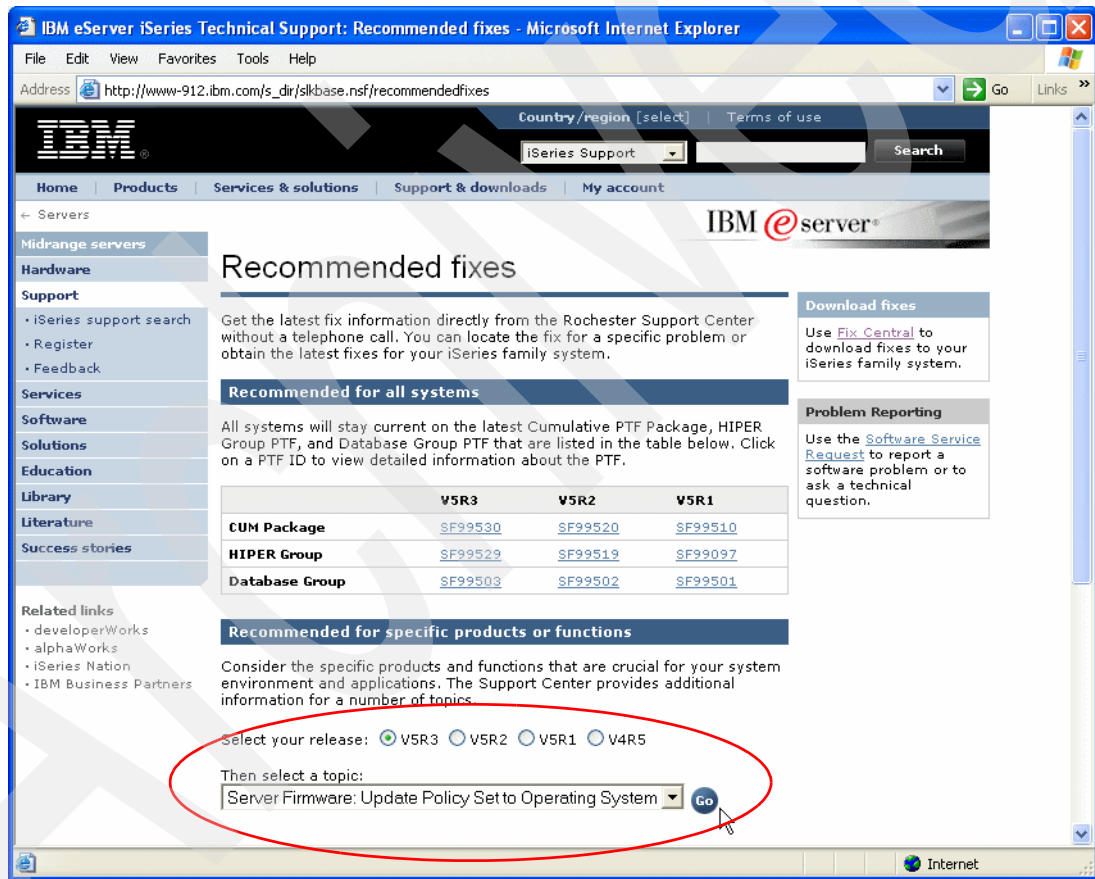


Figure 14-46 i5 recommended fixes Web site - in-band

You should select **Server Firmware: Update Policy Set to Operating System** from the drop down topic window and click **Go**.

The next screen shown is the iSeries Recommended Fixes - Server Firmware: Update Policy Set to Operating System Web page (Figure 14-47)

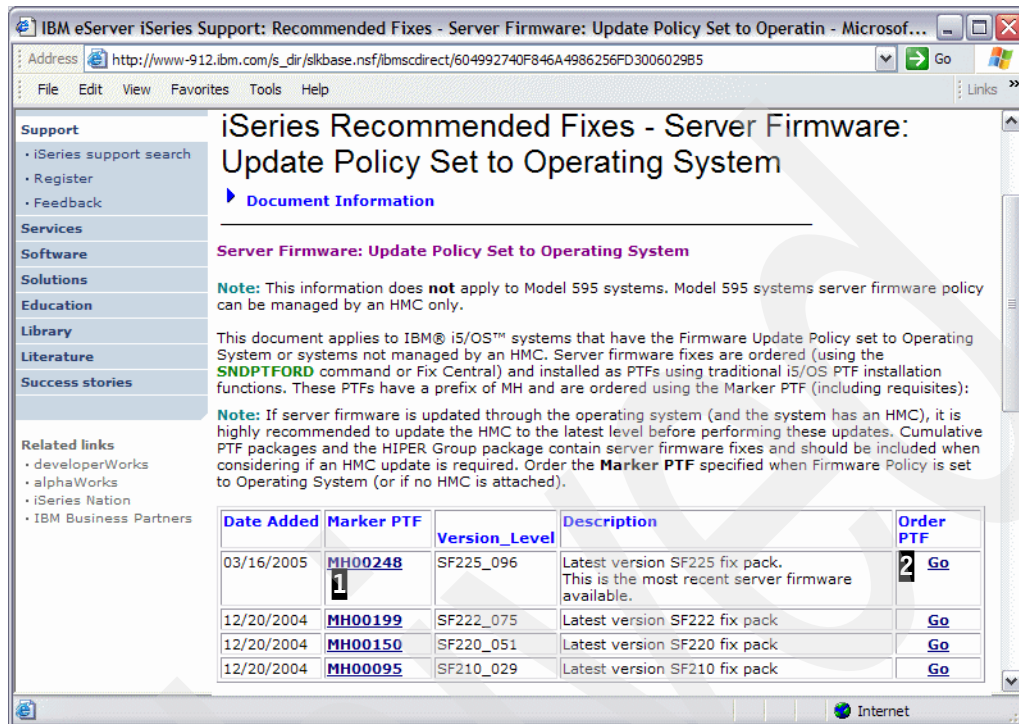


Figure 14-47 i5 recommended fixes Web site MHxxxxx- in-band

There are numerous ways of obtaining the i5 firmware. Normally the firmware is ordered by using a marker PTF MHxxxxx. This marker PTF may have several co-requisite PTFs which make up the firmware package.

- ▶ Order a firmware MHxxxxx [1](#) marker PTF electronically or on CD by using the SNDPTFORD i5/OS command or Fix Central. For example, the SF225_096 firmware can be ordered by using requesting PTF MH00248.
- ▶ Order the firmware by clicking the [Go](#) link [2](#) and ordering through the Internet. You must have a registered IBM ID to order the HMC fixes through the Internet.
- ▶ MHxxxxx PTFs are also contained in both Cumulative PTF packages and the HIPER Group package. You should check which firmware PTFs are included in these packages before installing.

Installing i5 in-band firmware updates

In this section we show how to update the i5 firmware through an i5/OS service partition using the MHxxxxx marker PTFs. After you have received the MHxxxxx PTFs they should be loaded and applied using the standard i5/OS PTF procedure. Our example shows how to update the managed system firmware from SF220_075 to SF225_096 using the marker PTF MH00248. The system used in this exercise is running on the b-side for i5/OS and the t-side for the server firmware - recommend normal system operations. Our example is based on firmware levels available at the time of writing this redbook and may change with future releases.

Important: You must have your firmware update policy set to operating system AND have one of your i5/OS partitions set to be the service partition to be able to update the managed systems firmware through the in-band method.

Figure 14-48 shows our current firmware levels before we install and apply our marker PTF. see “Using the i5/OS service partition to display firmware levels” on page 447 to view your own system firmware levels.

We can see that our system firmware update policy is set to operating system by the OS MANAGED key word show below.

```

Display Formatted Data
Page/Line. . . . . 1 / 4
Columns. . . . . : 1 - 78

Find . . . . .
...+...1...+...2...+...3...+...4...+...5...+...6...+...7...+...
LS Flash Sync Enabled.
OS MANAGED
MiKeyword: NVRAM: SF222_075

Side          Date/Time      MI Keyword    PTFs
-----
Memory        20041215/1425 SF222_075    MH00199
Flash P       20041215/1425 SF222_075    MH00199
Flash T       20041215/1425 SF222_075    MH00199
Load Source A 20041215/1425 SF222_075    MH00199
Load Source B 20041215/1425 SF222_075    MH00199

Bottom

```

Figure 14-48 i5/OS firmware levels before PTF installation

Important: We recommend that you install HMC fixes before you upgrade to a new server firmware release.

The following commands are issued from the i5/OS service partition 5250 command line:

1. Load the firmware PTFs on your service partition. In our example, we used Fix Central to download the marker PTF MH00248 and its co-requisite directly to the service partition.
 - a. LODPTF LICPGM(5722999) DEV(*SAVF) SELECT(MH00230) SAVF(*LIBL/QMH00230)
 - b. LODPTF LICPGM(5722999) DEV(*SAVF) SELECT(MH00251) SAVF(*LIBL/QMH00251)
 - c. LODPTF LICPGM(5722999) DEV(*SAVF) SELECT(MH00252) SAVF(*LIBL/QMH00252)
 - d. LODPTF LICPGM(5722999) DEV(*SAVF) SELECT(MH00248) SAVF(*LIBL/QMH00248)
2. Apply the marker PTF as a delayed PTF along with its co-requisite PTFs:


```
APYPTF LICPGM(5722999) SELECT(MH00248) DELAYED(*YES) APYREQ(*YES)
```

We receive message - CPI35A1 Licensed internal code fixes set to be temporarily applied.

- We use the i5/OS DSPPTF command to see the status of the applied PTFs (see Figure 14-49). There is a new PTF status indicator for the firmware PTFs. Notice in the status field that our firmware PTFs are set as 'Not applied - IPL'. This new status means that we must perform a system IPL (shutdown of all partitions and restart the managed system) to activate the new firmware PTFs.

```

                                Display PTF Status
                                System: XXXXXXXX
Product ID . . . . . : 5722999
IPL source . . . . . : ##MACH#B ##SERV#T
Release of base option . . . . . : V5R3M0 L00

Type options, press Enter.
  5=Display PTF details  6=Print cover letter  8=Display cover letter

PTF                               IPL
Opt ID      Status                 Action
MH00252    Not applied - IPL       Yes
MH00251    Not applied - IPL       Yes
MH00248    Not applied - IPL       Yes
MH00247    Superseded              None
MH00246    Superseded              None
MH00245    Superseded              None
MH00242    Superseded              None
MH00241    Superseded              None
MH00240    Superseded              None
More...

F3=Exit  F11=Display alternate view  F17=Position to  F12=Cancel

```

Figure 14-49 PTF status screen before system IPL

- Next we shut down ALL other partitions on the managed system (except the service partition).
- Finally, we power down the service partition using the i5/OS PWRDWNSYS command.

PWRDWNSYS OPTION(*IMMED) RESTART(*YES) IPLSRC(B)

When you shut down the i5/OS service partition the system reference code D6xx430B (copy firmware to the t-side) or D6xx430A (copy firmware to the p-side) might be displayed for an extended period of time. The xx should increment periodically and is a normal part of processing when server firmware fixes are being installed. When the service partition has finished installing the new server firmware, the managed system will be automatically restarted. You must allow the server to complete the processing; do not interrupt this process. In our exercise this process took around 40 mins to complete.

- When the service partition is active, check the PTF status with the `i5/OS DSPPTF` command and firmware levels again from the DST/SST environment. Figure 14-50 shows our firmware levels after the system IPL.

```

Display Formatted Data
Page/Line. . . . . 1 / 4
Columns. . . . . : 1 - 78


Find . . . . .
....+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....
LS Flash Sync Enabled.
OS MANAGED
MiKeyword: NVRAM: SF222_075
Side          Date/Time      MI Keyword  PTFs
-----
Memory        20050311/1416  SF225_096  MH00248
Flash P       20041215/1425  SF222_075  MH00199
Flash T       20050311/1416  SF225_096  MH00248
Load Source A 20041215/1425  SF222_075  MH00199
Load Source B 20050311/1416  SF225_096  MH00248
Bottom

```

Figure 14-50 *i5/OS firmware levels after PTF installation*

This completes the installation of i5 in-band firmware updates.

Archived



HMC Access Password Reset Using Advanced System Management Interface (ASMI)

The HMC Access password is a managed system password used to authenticate the HMC. It is one of three managed system passwords set when the system is first installed. Using Advanced System management Interface (ASMI), you can reset the HMC Access password if the password is lost. You can access ASMI via a Web browser, an ASCII console, or HMC. In this section, we will access ASMI via HMC or Web browser only.

15.1 Accessing the ASMI using the HMC

To access the Advanced System Management Interface using the Hardware Management Console (HMC), complete the following steps:

1. Ensure that the HMC is set up and configured properly.
2. In the HMC console navigation area, expand the managed system you want to work with.
3. Expand Service Applications and select Service Focal Point.
4. In the content area, select Service Utilities.
5. From the Service Utilities window, select the managed system you want to work with.
6. From the Selected menu, select Launch ASM menu to display ASM interface.

15.2 Accessing the ASMI using a Web browser

To access the Advanced System Management Interface using a Web browser, access the Service Processor using the following method:

https://ServiceProcessorIPAddress/

or

https://ServiceProcessorHostName/

When your Web browser is connected to ASMI, the Web browser will display ASMI main page as shown in Figure 15-1. Sign-in to ASMI as user admin and enter the password for user admin. The default password for user admin is admin.

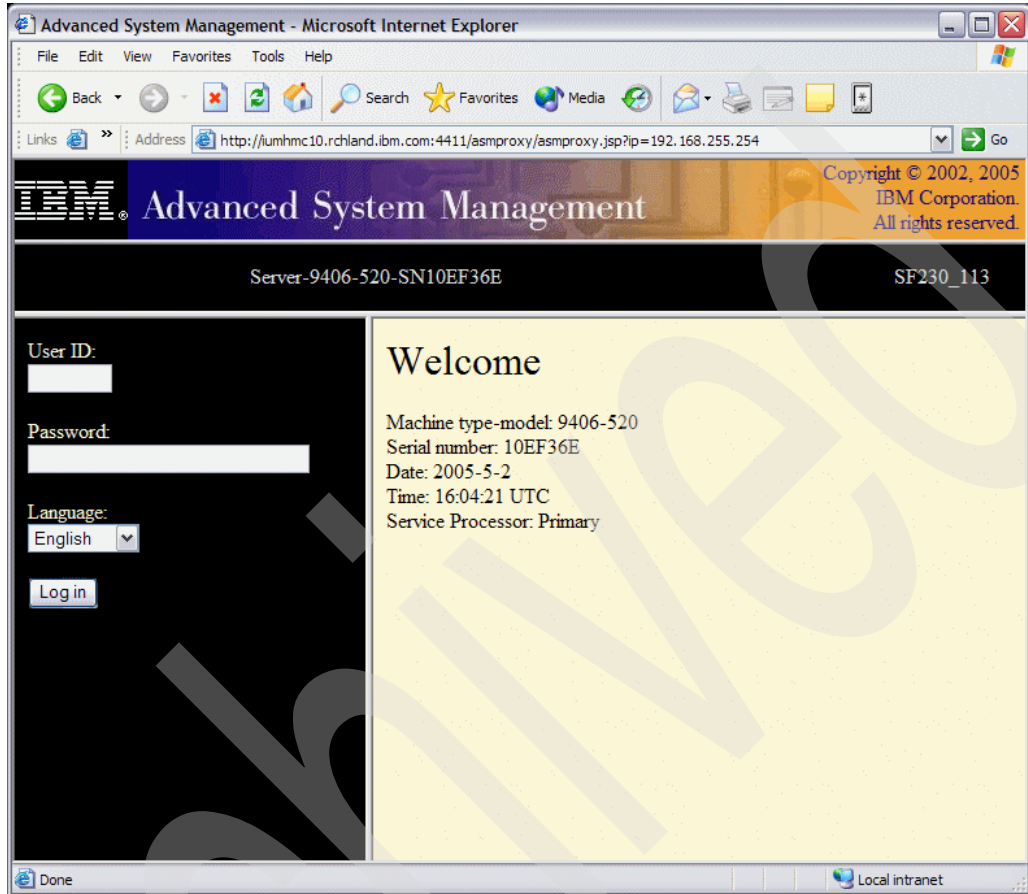


Figure 15-1 ASMI main page

After you have logged into ASMI, the Web browser will display the ASMI main menu as shown in Figure 15-2.

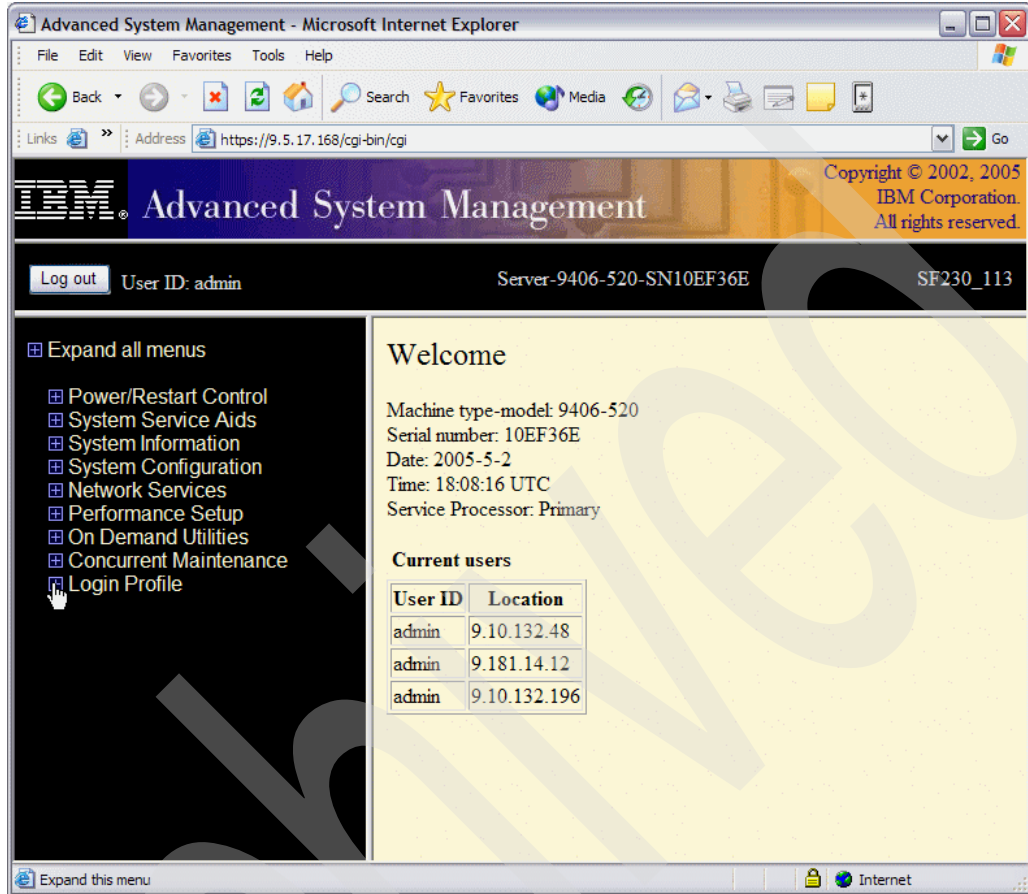


Figure 15-2 ASMI main menu

To change HMC Access password after you logged into ASMI, perform the following steps:

1. Select the Login Profile menu and expand it. It will display four sub-menus, they are Change Password, Retrieve Login Audits, Change Default Language, and Change Installed Language.
2. Select the Change Password menu (Figure 15-3).

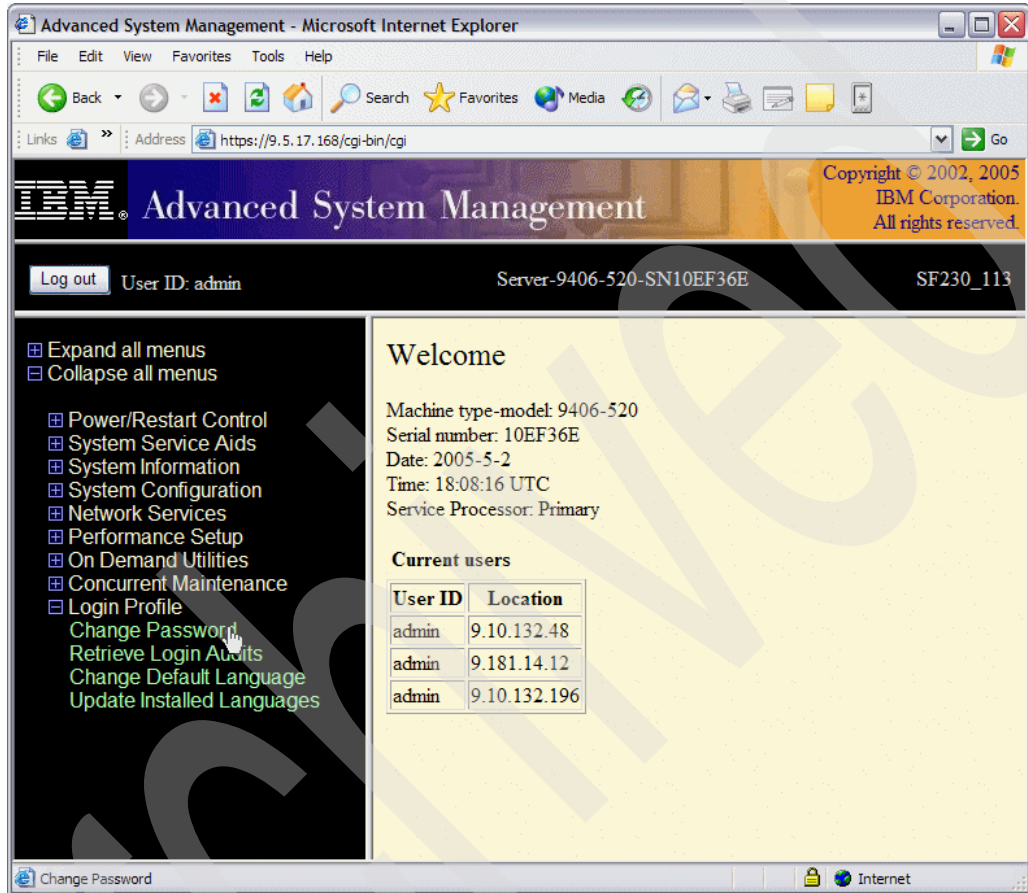


Figure 15-3 Select Change Password menu from Login Profile

3. Select the user ID of the user whose password you wish to change. In order to change the HMC Access password, select the HMC user ID from the drop down menu (Figure 15-4).

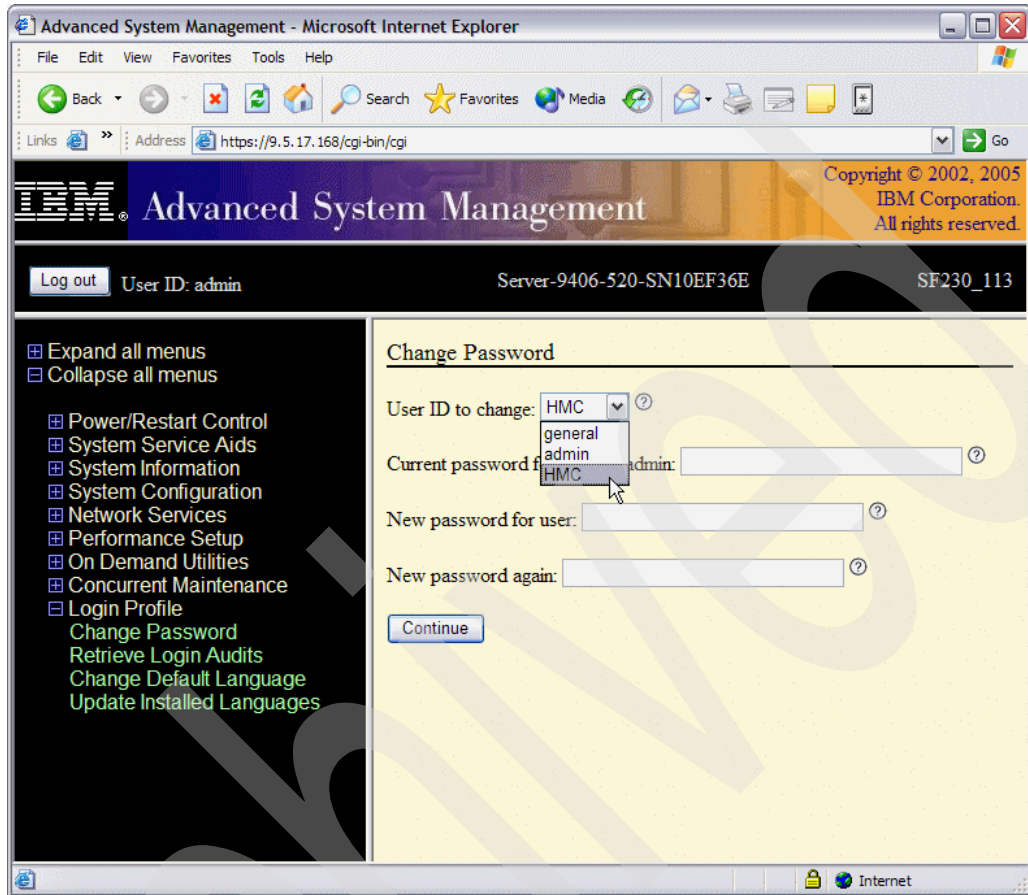


Figure 15-4 Select user ID HMC to change HMC access password

4. Enter the admin current password, insert the new password for the HMC Access user, and re-enter the new password as shown in Figure 15-5.

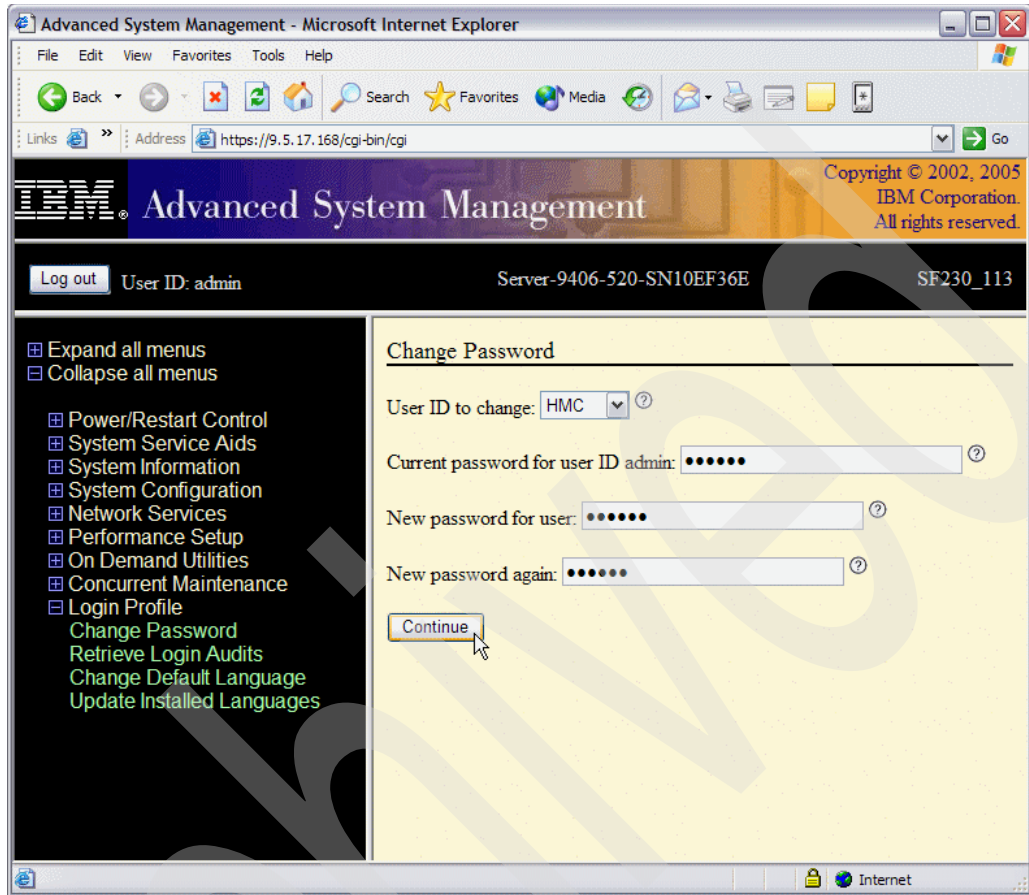


Figure 15-5 Enter password for admin and HMC

5. If the admin password was not entered correctly, ASMI will notify the failure and the HMC Access password is not changed as shown in Figure 15-6.

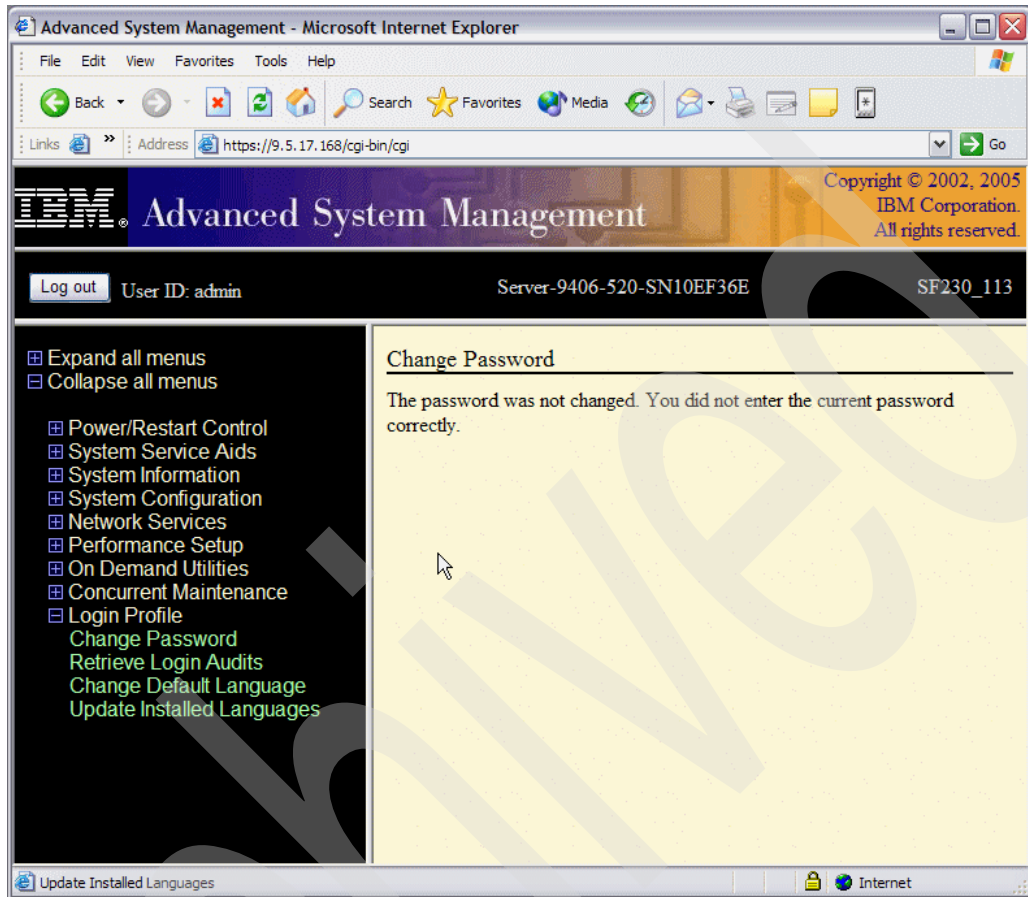


Figure 15-6 Wrong current password notification from ASMI

6. Click **Continue** to change the HMC Access password. ASMI will notify you if the password has been changed successfully (Figure 15-7).

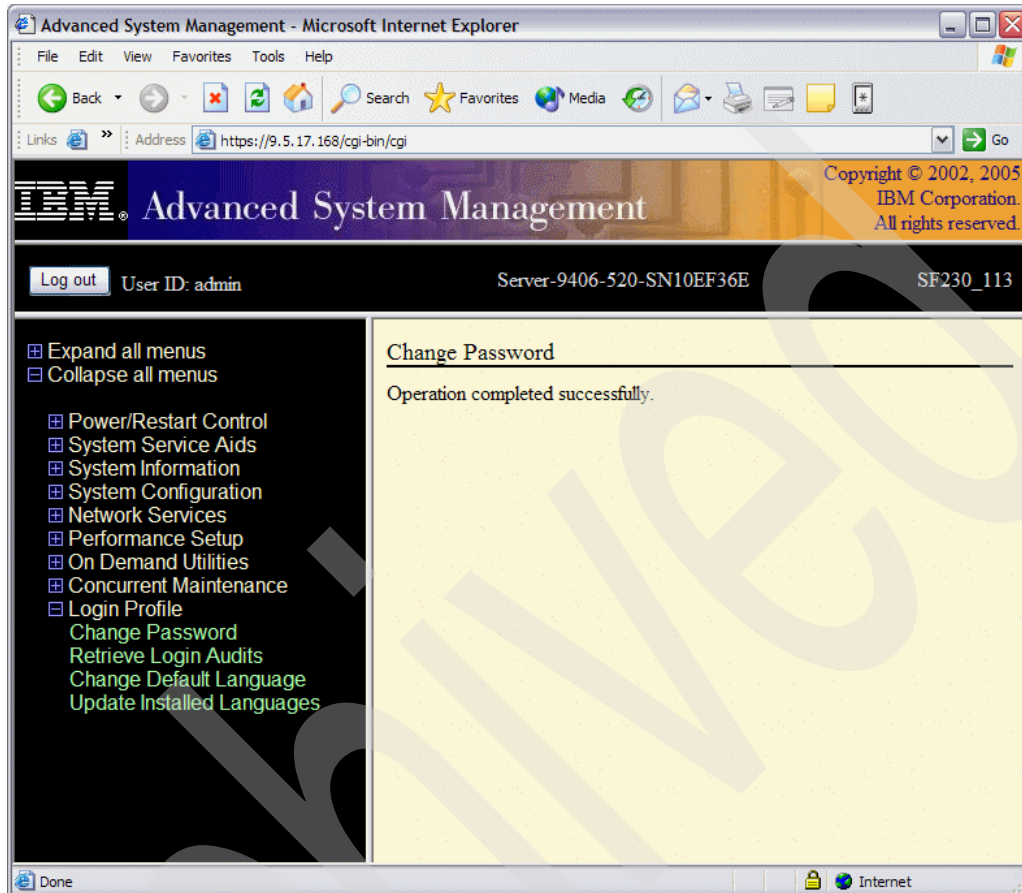


Figure 15-7 Change password completed

After resetting the HMC Access password, you can use it to authenticate HMC to the Managed System.

If you have lost your administrator password, you cannot reset the HMC Access password until you know your current administrator password. There are two methods to reset the administrator password:

1. Use the reset toggle jumpers on the service processors.
2. Contact your service providers.

To reset the toggle jumpers on the service processor, move both reset toggle switches to the opposite direction from current position. Please note that the last method of resetting toggle switches will cause loss of all configuration data, including passwords, partition information, profiles, and networking configurations. Please contact your local support representative for this activity.

Archived



HMC command list

This appendix contains the following topics:

- ▶ HMC CLI Introduction
- ▶ HMC CLI Commands listed by task
- ▶ HMC CLI Commands listed by name
- ▶ HMC CLI Commands listed by category
- ▶ HMC CLI Command usage
- ▶ HMC CLI Command attributes

HMC CLI introduction

The primary intention of the HMC CLI (Command Line Interface) is for the creation of scripts to automate the management of partition profiles. For example, a script could move processing resources into a partition for nightly batch processing and move those resources back before the start of daily operations in the morning.

HMC command naming convention

HMC commands are named following the UNIX naming convention for commands. In particular:

- ▶ **mk** is used for create/make actions.
- ▶ **ls** is used for list/query actions.
- ▶ **ch** is used for change/modify actions.
- ▶ **rm** is used for remove/delete actions.

Table A-1 shows the correlation between OS/400 command verbs, like DSP for display, and their UNIX counterparts.

Commands follow the form of verb plus an abbreviated name of the object on which the operation will be performed.

Table A-1 OS/400 command verbs and UNIX equivalent cross reference

OS/400 command verb	UNIX equivalent
DSP (Display)	ls (List)
CRT (Create)	mk (Make)
CHG (Change)	ch (Change)
DLT (Delete)	rm (Remove)

Note: Help for a command can be obtained by specifying the parameter `--help` on the command line. Any other parameters are ignored.

HMC CLI commands listed by task

Table A-2 lists the various partition and profile tasks along with the respective HMC CLI command.

Table A-2 HMC commands by task

Task	HMC CLI command
Activate partition	chsysstate
Activate system profile	chsysstate
Add a physical I/O slot to a partition	chhwres
Add memory to a partition	chhwres
Add processors to a partition	chhwres
Create LPAR	mksyscfg
Create LPAR profile	mksyscfg

Task	HMC CLI command
Create system profile	mksyscfg
Delete LPAR	rmsyscfg
Delete LPAR profile	rmsyscfg
Delete system profile	rmsyscfg
Determine DRC indexes for physical I/O slots	lshwres
Determine memory region size	lshwres
Fast power off the managed system	chsysstate
Get LPAR state	lssyscfg
Hard partition reset	chsysstate
List all partitions in a managed system	lssyscfg
List all systems managed by the HMC	lssyscfg
List I/O resources for a managed system	lshwres
List LPAR profile properties	lssyscfg
List managed system properties	lssyscfg
List memory resources	lshwres
List processor resources	lshwres
List system profile properties	lssyscfg
List virtual I/O resources for a managed system	lshwres
Modify LPAR profile properties	chsyscfg
Modify LPAR properties	chsyscfg
Modify managed system properties	chsyscfg
Modify system profile properties	chsyscfg
Move a physical I/O slot from one partition to another	chhwres
Move memory from one partition to another	chhwres
Move processors from one partition to another	chhwres
Power off the managed system	chsysstate
Power on the managed system	chsysstate
Re-IPL the managed system	chsysstate
Remove a physical I/O slot from a partition	chhwres
Remove memory from a partition	chhwres
Remove processors from a partition	chhwres
List LPAR properties	lssyscfg
Soft partition reset	chsysstate
Validate a system profile	chsysstate

HMC CLI commands by name

Table A-3 lists the HMC commands by name.

Table A-3 HMC commands listed by name

Command	Description	Associated tasks
chhwres	Change system memory and processor resources	<ul style="list-style-type: none"> Add a physical I/O slot to a partition Add memory to a partition Add processors to a partition Move a physical I/O slot from one partition to another Move memory from one partition to another Move processors from one partition to another Remove a physical I/O slot from a partition Remove memory from a partition Remove processors from a partition
chsyscfg	Change system configuration	<ul style="list-style-type: none"> Modify LPAR properties Modify LPAR profile properties Modify managed system properties Modify system profile properties Change system state - chsysstate Activate partition Activate system profile Fast power off the managed system Hard partition reset Power off the managed system Power on the managed system Re-IPL the managed system Soft partition reset Validate a system profile
chsysstate	Change system state	<ul style="list-style-type: none"> Activate partition Activate system profile Fast power off the managed system Hard partition reset Power off the managed system Power on the managed system Re-IPL the managed system Soft partition reset Validate a system profile
lshwres	List the hardware resources of a managed system	<ul style="list-style-type: none"> Determine DRC indexes for physical I/O slots Determine memory region size List I/O resources for a managed system List memory resources List processor resources List virtual I/O resources for a managed system
lssyscfg	List system configuration information	<ul style="list-style-type: none"> Get LPAR state List all partitions in a managed system List all systems managed by the HMC List LPAR profile properties List LPAR properties List managed system properties List system profile properties
mksyscfg	Create system configuration	<ul style="list-style-type: none"> Create LPAR profile Create system profile

Command	Description	Associated tasks
rmsyscfg	Remove system configuration	Delete LPAR Delete LPAR profile Delete system profile

HMC CLI commands by category

The following section discusses how to perform various functions using the HMC CLI. The functions are broken down into categories, managed system, DLPAR, etc.

Working with the managed system

Here we list the commands for working with the managed system.

Powering on the managed system

Use the **chsysstate** command to power on the managed system. This power on option is similar to a “White Button” power on. Enter the following command:

```
chsysstate -r sys -m <managed system> -o on
```

Powering off the managed system

Use the **chsysstate** command to power off the managed system. This can be done in two ways.

To power off a managed system, type the following command:

```
chsysstate -r sys -m <managed system> -o off
```

You can also perform a fast power off of a managed system. To do this, type the following command:

```
chsysstate -r sys -m <managed system> -o fastoff
```

Re-IPLing the managed system

To re-IPL the managed system, enter the following command:

```
chsysstate -r sys -m <managed system> -o reipl
```

Modifying the managed system properties

Use the **chsyscfg** command to change the properties of a managed system. The following example shows how to change the managed system's user-defined name and power off policy:

```
chsyscfg -r sys -m <managed system> -i "new_name=squadrons1,power_off_policy=0"
```

Valid attributes, specified with the **-i** flag, include:

- ▶ ipl_side
- ▶ ipl_speed
- ▶ ipl_speed_override
- ▶ new_name
- ▶ power_off_policy
- ▶ service_lpar_id
- ▶ service_lpar_name

Command attributes are discussed in “HMC CLI command attributes” on page 494.

Listing all systems managed by the HMC

Use the `lssyscfg` command to list system configuration and managed system MTMS information. To list the properties of all systems managed by this HMC, enter the following:

```
lssyscfg -r sys
```

To list only the user-defined names, model types, serial numbers, and states for all of the systems managed by this HMC, enter:

```
lssyscfg -r sys -F name,type_model,serial_num,state --header
```

Note that after executing this command, `<type_model>` concatenated with `*` then with `<serial_num>` represents the MODEL-TYPE*SERIAL, or MTMS, for the managed system.

Listing a managed system's properties

To list a specific managed system's properties, enter:

```
lssyscfg -r sys -m <managed system>
```

Working with logical partitions

This section describes the commands for working with logical partitions.

Creating LPARs

Use the `mksyscfg` command to create a partition.

The following example shows how to create an AIX/Linux partition:

```
mksyscfg -r lpar -m <managed system>
-i"lpar_id=2,name=aixlinux_lpar2,profile_name=prof1,lpar_type=aixlinux,boot_mode=norm,sfp_s
urveillance=1,desired_procs=1,min_procs=1,max_procs=1,min_proc_units=0.1,desired_proc_units
=0.5,max_proc_units=0.5,proc_type=shared,sharing_mode=cap,desired_mem=400,min_mem=400,
max_mem=400,auto IPL=1,power_ctrl_lpar_ids=0,io_slots=553713666/65535/1"
```

The following example shows how to create an OS/400 partition:

```
mksyscfg -r lpar -m <managed system>
-i"lpar_id=3,name=os400_lpar3,profile_name=prof1,lpar_type=os400,sfp_surveillance=1,desired
_procs=1,min_procs=1,max_procs=1,min_proc_units=0.1,desired_proc_units=0.5,max_proc_units=0
.5,proc_type=shared,sharing_mode=cap,desired_mem=400,
min_mem=400,max_mem=400,auto IPL=1,power_ctrl_lpar_ids=0,io_slots=553713699/65535/1,load_so
urce_slot=553713699,console_slot=553713699,min_interactive=0,desired_interactive=0,
max_interactive=0"
```

Valid attributes, specified with the `-i` flag, are listed below in Table A-4.

Table A-4 Command attributes for `mksyscfg -i` parameter

name	desired_proc_units	min_proc_units	lpar_id
max_proc_units	ecs_slot	profile_name	lpar_io_pool_ids
sni_windows	lpar_type	io_slots	alt_console_slot
cluster_id	boot_mode	sni_device_ids	sharing_mode
sfp_surveillance	console_slot	desired_procs	auto IPL
sni_config_mode	min_procs	power_ctrl_lpar_ids	alt_load_source_slot
max_procs	virtual_opti_pool_id	virtual_serial_adapters	desired_mem

name	desired_proc_units	min_proc_units	lpar_id
hsl_opti_pool_id	load_source_slot	min_mem	min_interactive
virtual_scsi_adapters	max_mem	desired_interactive	uncap_weight
proc_type	max_interactive	virtual_eth_adapters	max_virtual_slots

Tip: Instead of entering configuration information on the command line with the `-i` flag, the information can instead be placed in a file, and the filename specified with the `-f` flag.

Command attributes are discussed in “HMC CLI command attributes” on page 494.

Listing all partitions in a managed system

Use the `lssyscfg` command to list all partitions in a managed system.

To do this, enter:

```
lssyscfg -r lpar -m <managed system>
```

To list only the names, IDs, and states of all partitions in a managed system, enter:

```
lssyscfg -r lpar -m <managed system> -F name,lpar_id,state --header
```

Listing LPAR properties

Use the `lssyscfg` command to list the properties of a specific partition.

Type the following command:

```
lssyscfg -r lpar -m <managed system> --filter"lpar_ids=<partition ID>"
```

Note: The partition name can be specified instead of the partition ID by using the `lpar_names` filter in place of the `lpar_ids` filter. Also, more than one partition may be specified in the filter list.

Getting the LPAR state

Use the `lssyscfg` command to display the state of a partition. Type the following command:

```
lssyscfg -r lpar -m <managed system> --filter "lpar_names=<partition name>" -F state
```

Note: The partition ID can be specified instead of the partition name by using the `lpar_ids` filter in place of the `lpar_names` filter. Also, more than one partition may be specified in the filter list.

Modifying LPAR properties

Use the `chsyscfg` command to modify the properties of a partition. The following example shows how to change a partition’s cluster ID:

```
chsyscfg -r lpar -m <managed system> -i"lpar_id=1,cluster_id=3"
```

Valid attributes, specified with the `-i` flag, are:

- new_name name
- lpar_id
- default_profile
- cluster_id

Note: Instead of entering configuration information on the command line with the **-i** flag, the information can instead be placed in a file, and the filename specified with the **-f** flag.

Command attributes are discussed in “HMC CLI command attributes” on page 494.

Activating a partition

Use the **chsysstate** command to activate a partition. Type the following command:

```
chsysstate -r lpar -m <managed system> -o on -n <partition name> -f <partition profile name>
```

Note: The partition ID can be specified instead of the partition name by using the **--id** parameter instead of the **-n** parameter.

Using a soft partition reset

Use the **chsysstate** command to perform a soft reset of a partition. Type the following command:

```
chsysstate -r lpar -m <managed system> -o reset -n <partition name>
```

Note: The partition ID can be specified instead of the partition name by using the **--id** parameter instead of the **-n** parameter.

Using a hard partition reset

Use the **chsysstate** command to perform a hard reset of a partition. Type the following command:

```
chsysstate -r lpar -m <managed system> -o off --id<partition ID>
```

Note: The partition name can be specified instead of the partition ID by using the **-n** parameter instead of the **--id** parameter.

Deleting an LPAR

Use the **rmsyscfg** command to remove a partition. Type the following command:

```
rmsyscfg -r lpar -m <managed system> -n <partition name>
```

This command removes the specified partition and all of its associated partition profiles from the specified managed system. The partition's profiles are also removed from any system profiles that contain them.

Note: The partition ID can be specified instead of the partition name by using the **--id** parameter instead of the **-n** parameter.

Dynamic LPAR (DLPAR) operations

Use the **chhwres** command to perform dynamic logical partitioning (DLPAR) operations on running partitions. DLPAR operations can be performed for memory, physical I/O slots, and processor resources.

Memory

Memory can be dynamically added to a partition, removed from a partition, or moved from one partition to another. In the following commands, the quantity of memory to be added,

removed, or moved must be specified with the `-q` flag. This quantity is in *megabytes*, and *must* be a multiple of the memory region size for the managed system.

Determining memory region size

To see what the memory region size is for the managed system, enter this command:

```
lshwres -r mem -m <managed system> --level sys -F mem_region_size
```

The value returned is the memory region size in megabytes.

Adding memory to a partition

To add memory to a partition, enter this command:

```
chhwres -r mem -m <managed system> -o a -p <partition name> -q <quantity>
```

Removing memory from a partition

To remove memory from a partition, enter this command:

```
chhwres -r mem -m <managed system> -o r -p <partition name> -q <quantity>
```

Moving memory from one partition to another

To move memory from one partition to another partition, enter this command:

```
chhwres -r mem -m <managed system> -o m -p <source partition name> -t <target partition name> -q <quantity>
```

Physical I/O slots

A physical I/O slot can be dynamically added to a partition, removed from a partition, or moved from one partition to another. In the following commands, the DRC index of the I/O slot to be added, removed, or moved must be specified with the `-s` flag.

Note: Note that only one physical I/O slot can be added, removed, or moved at a time.

Determining DRC Indexes for physical I/O slots

To see the DRC indexes for all of the physical I/O slots that are on the managed system, enter this command:

```
lshwres -r io --rsubtype slot -m <managed system>
```

The DRC index for each slot is returned via the `drc_index` attribute.

Adding a physical I/O slot to a partition

To add a physical I/O slot to a partition, enter this command:

```
chhwres -r io -m <managed system> -o a -p <partition name> -s <DRC index>
```

Removing a physical I/O slot from a partition

To remove a physical I/O slot from a partition, enter this command:

```
chhwres -r io -m <managed system> -o r -p <partition name> -s <DRC index>
```

Moving a physical I/O slot from one partition to another

To move a physical I/O slot from one partition to another partition, enter this command:

```
chhwres -r io -m <managed system> -o m -p <source partition name> -t <target partition name> -s <DRC index>
```

Processors

Processing resources can be dynamically added to a partition, removed from a partition, or moved from one partition to another. These processing resources depend on the type of processors used by the partitions:

- ▶ For partitions using dedicated processors, processing resources are dedicated processors.
- ▶ For partitions using shared processors, processing resources include virtual processors and processing units.

Note: Currently, AIX/Linux partitions using shared processors do not support processor DLPAR operations.

In the following commands, for partitions using dedicated processors, the quantity of processors to be added, removed, or moved are specified with the `--procs` flag. For partitions using shared processors, the quantity of virtual processors to be added, removed, or moved are also specified with the `--procs` flag. The quantity of processing units to be added, removed, or moved are specified with the `--procunits` flag. Both of these flags can be specified, but only one is required.

Note: The quantity of processing units must be multiplied by 100 for the command. For example, to add, remove, or move .5 processing units, specify a quantity of 50.

Adding processors to a partition

To add processors to a partition using dedicated processors, enter this command:

```
chhwres -r proc -m <managed system> -o a -p <partition name> --procs <quantity>
```

To add processors to a partition using shared processors, enter this command:

```
chhwres -r proc -m <managed system> -o a -p <partition name> --procs <quantity> --procunits <quantity>
```

Removing processors from a partition

To remove processors from a partition using dedicated processors, enter this command:

```
chhwres -r proc -m <managed system> -o r -p <partition name> --procs <quantity>
```

To remove processors from a partition using shared processors, enter this command:

```
chhwres -r proc -m <managed system> -o r -p <partition name> --procs <quantity> --procunits <quantity>
```

Moving processors from one partition to another

To move processors from a partition using dedicated processors to another, enter this command:

```
chhwres -r proc -m <managed system> -o m -p <source partition name> -t <target partition name> --procs <quantity>
```

To move processors from a partition using shared processors to another, enter this command:

```
chhwres -r proc -m <managed system> -o m -p <source partition name> -t <target partition name> --procs <quantity> --procunits <quantity>
```


Processing resources can also be moved between partitions using dedicated processors and partitions using shared processors. To move processing resources from a partition using dedicated processors to a partition using shared processors, specify the quantity of processors using the `--procs` flag. This quantity is converted to processing units (by multiplying the quantity by 100) by the HMC for the target partition.

To move processing resources from a partition using shared processors to a partition using dedicated processors, specify the quantity of processing units (which must be a multiple of 100) using the `--procunits` flag. This quantity is converted to processors (by dividing the quantity by 100) by the HMC for the target partition. The `--procs` flag cannot be specified in this case.

Working with partition profiles

This section describes commands for working with partition profiles.

Creating an LPAR profile

Use the `mksyscfg` command to create a partition profile. The following example shows how to create a partition profile:

```
mksyscfg -r prof -m <managed system>
-i "name=prof3,lpar_id=2,boot_mode=norm,sfp_surveillance=1,desired_procs=2,min_procs=1,max_p
rocs=2,min_proc_units=0.1,desired_proc_units=0.5,max_proc_units=0.5,proc_type=shared,sharin
g_mode=cap,desired_mem=400,min_mem=400,max_mem=400,auto_ipl=1,power_ctrl_lpar_ids=0,io_slot
s=553713666/65535/1"
```

Valid attributes, specified with the `-i` flag, include those shown in Table A-5.

Table A-5 Valid attributes on the `mksyscfg` command

name	lpar_id lpar_name	power_ctrl_lpar_ids	desired_procs
min_procs	desired_mem	min_mem	max_mem
uncap_weight	sharing_mode	load_source_slot	alt_load_source_slot
console_slot	alt_console_slot	ecs_slot	min_proc_units
desired_proc_units	max_proc_units	lpar_io_pool_ids	io_slots
boot_mode	sfp_surveillance	sni_windows	virtual_opti_pool_id
hsl_opti_pool_id	min_interactive	desired_interactive	max_interactive
max_virtual_slots	virtual_eth_adapters	virtual_scsi_adapters	virtual_serial_adapters
sni_config_mode	sni_device_ids	auto_ipl	max_procs
proc_type			

The profile name (**name**) and the partition (**lpar_id** or **lpar_name**) must be specified.

Note: Instead of entering configuration information on the command line with the `-i` flag, the information can instead be placed in a file, and the filename specified with the `-f` flag.

Command attributes are discussed in “HMC CLI command attributes” on page 494.

Listing LPAR profile properties

Use the `lssyscfg` command to list a partition profile. Type the following command:

```
lssyscfg -r prof -m <managed system> --filter "lpar_names=<partition name>,  
profile_names=<profile name>"
```

Use the `--filter` parameter to specify the partition for which partition profiles are to be listed, and to specify which profile names to list.

While the filter can only specify a single partition, it can specify multiple profile names for that partition.

Note: The partition ID can be specified instead of the partition name by using the `lpar_ids` filter in place of the `lpar_names` filter.

Modifying LPAR profile properties

Use the `chsyscfg` command to modify a partition profile's properties. The following example shows how to change **prof1's** memory amounts:

```
chsyscfg -r prof -m <managed system> -i"name=prof1,lpar_name=lpar3,min_mem=256,  
max_mem=512,desired_mem=512"
```

Valid attributes, specified with the `-i` flag, include those shown in Table A-6.

Table A-6 Valid attributes on the `chsyscfg` command

name	lpar_name lpar_id	new_name	desired_procs
min_procs	max_procs	desired_mem	min_mem
max_mem	proc_type	uncap_weight	sharing_mode
load_source_slot	alt_load_source_slot	console_slot	alt_console_slot
ecs_slot	min_proc_units	desired_proc_units	max_proc_units
lpar_io_pool_ids	io_slots	boot_mode	sfp_surveillance
sni_windows	virtual_opti_pool_id	hsl_opti_pool_id	min_interactive
desired_interactive	max_interactive	max_virtual_slots	virtual_eth_adapters
virtual_scsi_adapters	virtual_serial_adapters	sni_config_mode	sni_device_ids
auto_ipl	power_ctrl_lpar_ids		

Command attributes are discussed in "HMC CLI command attributes" on page 494.

Note: Instead of entering configuration information on the command line with the `-i` flag, the information can instead be placed in a file, and the filename specified with the `-f` flag.

Deleting an LPAR profile

Use the `rmsyscfg` command to remove a partition profile. Type the following command:

```
rmsyscfg -r prof -m <managed system> -n <profile name> -p <partition name>
```

Note: The partition ID can be specified instead of the partition name by using the `--id` parameter in place of the `-p` parameter.

Working with system profiles

This section describes commands for working with system profiles.

Creating a system profile

Use the **mksyscfg** command to create a system profile. In the following example, the user is making a system profile named **sysprof1**, with partition profile **prof1** for partition **lpar1** and partition profile **prof1** for partition **lpar2**.

```
mksyscfg -r sysprof -m <managed system>
-i"name=sysprof1,\"lpar_names=lpar1,lpar2\", \"profile_names=prof1,prof1\""
```

Partition IDs can be specified instead of partition names when creating a system profile. This is done by using the **lpar_ids** attribute instead of the **lpar_names** attribute.

Note: Instead of entering configuration information on the command line with the **-i** flag, the information can instead be placed in a file, and the filename specified with the **-f** flag.

Activating a system profile

Use the **chsysstate** command to activate a system profile. Type the following command:

```
chsysstate -r sysprof -m <managed system> -o on -n <system profile name>
```

Validating a system profile

Use the **chsysstate** command to validate a system profile. Type the following command:

```
chsysstate -r sysprof -m <managed system> -n <system profile name> --test
```

To validate a system profile, then activate that system profile if the validation is successful, Type the following command:

```
chsysstate -r sysprof -m <managed system> -o on -n <system profile name> --test
```

Deleting a system profile

Use the **rmsyscfg** command to remove a system profile. Type the following command:

```
rmsyscfg -r sysprof -m <managed system> -n <system profile name>
```

Listing system profile properties

Use the **lssyscfg** command to list a system profile's properties. Type the following command:

```
lssyscfg -r sysprof -m <managed system> --filter "profile_names=<system profile name>"
```

To list all system profiles for the managed system, Type the following command:

```
lssyscfg -r sysprof -m <managed system>
```

Modifying system profile properties

Use the **chsyscfg** command to modify system profile properties. In the following example, the user is adding profiles **prof1** for partition **lpar3** and **prof2** for partition **lpar4** to system profile **sysprof1**:

```
chsyscfg -r sysprof -m <managed system>
-i"name=sysprof1,\"lpar_names+=lpar3,lpar4\", \"profile_names+=prof1,prof2\""
```

Valid attributes, specified with the **-i** flag, include:

- ▶ **new_name**

- ▶ profile_names
- ▶ lpar_names | lpar_ids
- ▶ name

Listing hardware resources

The `lshwres` command, which lists the hardware resources of a managed system, can be used to display I/O, virtual I/O, processor, and memory resources.

Listing I/O resources for a managed system

Use the following commands to list:

- ▶ I/O units on the managed system:


```
lshwres -m <managed system> -r io --subtype unit
```
- ▶ I/O buses on the managed system:


```
lshwres -m <managed system> -r io --subtype bus
```
- ▶ I/O slots on the managed system:


```
lshwres -m <managed system> -r io --subtype slot
```
- ▶ All partitions participating in an I/O pool and all slots assigned to an I/O pool:


```
lshwres -m <managed system> -r io --subtype iopool
```
- ▶ Tagged I/O for OS/400 partitions:


```
lshwres -m <managed system> -r io --subtype taggedio
```

Listing processor resources

Use the following commands to list processor information for:

- ▶ The managed system:


```
lshwres -m <managed system> -r proc --level sys
```
- ▶ Partitions:


```
lshwres -m <managed system> -r proc --level lpar
```
- ▶ The shared pool:


```
lshwres -m <managed system> -r proc --level sharedpool
```

Listing virtual I/O resources for a managed system

Use the following commands to list:

- ▶ Virtual Ethernet adapters:


```
lshwres -m <managed system> -r virtualio --subtype eth --level lpar
```
- ▶ System level virtual Ethernet information:


```
lshwres -m <managed system> -r virtualio --subtype eth --level sys
```
- ▶ Virtual OptiConnect pool information:


```
lshwres -m <managed system> -r virtualio --subtype virtualopti --level lpar
```
- ▶ HSL OptiConnect pool information:


```
lshwres -m <managed system> -r virtualio --subtype hsopti --level lpar
```
- ▶ Virtual serial adapters:


```
lshwres -m <managed system> -r virtualio --subtype serial --level lpar
```

- ▶ Virtual serial servers with open connections:
lshwres -m <managed system> -r virtualio --rsubtype serial --level openserial
- ▶ Virtual SCSI adapters:
lshwres -m <managed system> -r virtualio --rsubtype scsi --level lpar
- ▶ Partition-level virtual slot information:
lshwres -m <managed system> -r virtualio --rsubtype slot --level lpar
- ▶ Virtual slot information:
lshwres -m <managed system> -r virtualio --rsubtype slot --level slot

Listing memory resources

Use the following commands to list:

- ▶ Memory information for a managed system:
lshwres -m <managed system> -r mem --level sys
- ▶ Memory information for partitions:
lshwres -m <managed system> -r mem --level lpar

HMC CLI command usage

This section provides various examples of CLI command usage.

Change hardware resources (chhwres)

This command changes the hardware resource configuration of a managed system (Example A-1).

Example: A-1 Command usage for chhwres

```
Usage: chhwres -r io | mem | proc | virtualio
--rsubtype slot | taggedio | eth | scsi | serial
-m <managed system>
-o a | r | m | s
[-p <partition name> | --id <partition ID>]
[-t <target partition name> | --tid <target partition ID>]
[-s <DRC index>]
[-q <quantity>]
[--procs <quantity>]
[--procunits <quantity>]
[--interactive <percentage>]
[-w <wait time>]
[-d <detail level>]
[--force]
[-a "<attributes>"]
[--help]
```

Changes the hardware resource configuration of a managed system.

- | | |
|------------|--|
| -r | - the type of resource to be changed:
io - I/O
mem - memory
proc - processor
virtualio - virtual I/O |
| --rsubtype | - the subtype of resource to be changed: |

```

slot - I/O slot
taggedio - tagged I/O
eth - virtual ethernet
scsi - virtual SCSI
serial - virtual serial
-m <managed system> - the managed system's name
-o                 - the operation to perform:
                   a - add resources
                   r - remove resources
                   m - move resources
                   s - set attributes
-p <partition name> - the user defined name of the partition to add
                   resources to, to move or remove resources from,
                   or to set attributes for
--id <partition ID> - the ID of the partition to add resources to, to
                   move or remove resources from, or to set
                   attributes for
-t <target partition name> - the user defined name of the partition to move
                           resources to
--tid <target partition ID> - the ID of the partition to move resources to
-s <DRC index>           - the DRC index of the I/O slot to add, move,
                           or remove
-q <quantity>           - the quantity of memory, in megabytes, to add,
                           move, or remove
--procs <quantity>      - the quantity of processors to add, move, or
                           remove
--procunits <quantity> - the quantity of processing units to add, move,
                           or remove
--interactive <percentage> - the quantity of interactive processing
                           resources to add, move, or remove
-w <wait time>          - the elapsed time, in minutes, after which the
                           DLPAR operation will be aborted
-d <detail level>      - the level of detail to be displayed upon
                           return of the requested DLPAR operation -
                           values range from 0 (none) to 5 (highest)
-a "<attributes>"      - the attributes to be set with this command
                           the format is:
                           "attr_name1=value,attr_name2=value,..."
                           or
                           "'attr_name1=value1,value2,...',..."
--help                 - prints this help

```

The valid attribute names for this command are:

```

-r virtualio --rsubtype eth : mac_prefix
-r io --rsubtype taggedio : load_source_slot, alt_load_source_slot,
                             console_slot, alt_console_slot, ecs_slot

```

Change system configuration (chsyscfg)

This command changes the system configuration of a managed system (Example A-2).

Example: A-2 Command usage for chsyscfg

```

Usage: chsyscfg -r sys | lpar | prof | sysprof
              -m <managed system>
              -f <configuration file name> | -i "<configuration data>"
              [--help]

```

Changes the system configuration of a managed system.

- r - the type of system resource(s) to be changed:
 sys - managed system
 lpar - partition
 prof - partition profile
 sysprof - system profile
- m <managed system> - the managed system's name
- f <config file name> - the name of the file containing the configuration data for this command
 the format is:
 attr_name1=value,attr_name2=value,...
 or
 "attr_name1=value1,value2,...",...
- i "<config data>" - the configuration data for this command
 the format is:
 "attr_name1=value,attr_name2=value,..."
 or
 "attr_name1=value1,value2,...",..."
- help - prints this help

The valid attribute names for this command are:

- r sys optional: new_name, power_off_policy (0 | 1),
 service_lpar_id, service_lpar_name,
 ipl_side (perm | temp), ipl_speed (slow | fast),
 ipl_speed_override (none | slow | fast)
- r lpar required: name | lpar_id
 optional: new_name, default_profile, cluster_id,
 shared_pool_util_auth (0 | 1)
- r prof required: name, lpar_id | lpar_name
 optional: new_name
- r sysprof required: name
 optional: new_name, lpar_ids | lpar_names, profile_names

Change system state (chsysstate)

This command changes the system state (Example A-3).

Example: A-3 Command usage for chsysstate

```
Usage: chsysstate -r sys | lpar | sysprof
       -m <managed system>
       [-o on | off | fastoff | reipl ]
       [-n <object name>]
       [-f <profile name>]
       [--test]
       [--continue]
       [--help]
```

Changes the system state.

- m <managed system> - the managed system's name
- o on -r sys - power on the managed system
- o on -r lpar - activate a partition
- o on -r sysprof - activate a system profile
- o off - power off the managed system
 this parameter is only valid for -r sys
- o fastoff - perform a hard power off of the managed system
 this parameter is only valid for -r sys
- o reipl - perform a re-ipl of the managed system
 this parameter is only valid for -r sys
- n <object name> - the name of the partition or system profile

	to activate
-f <profile name>	- the name of the profile to use when activating a partition this parameter is only valid for -r lpar
--test	- validate the system profile this parameter is only valid for -r sysprof
--continue	- continue on error when activating a system profile this parameter is only valid for -r sysprof
--help	- prints this help

List hardware resources (lshwres)

This command lists the hardware resources of a managed system (Example A-4).

Example: A-4 Command usage for lshwres

```
Usage: lshwres -r io | mem | proc | virtualio
          [--rsubtype unit | bus | slot | iopool | taggedio | eth |
          hslopti | virtualopti | scsi | serial]
          -m <managed system>
          [--level sys | lpar | slot | sharedpool | openserial]
          [--filter "<filter data>"]
          [-F <attribute names>]
          [--header]
          [--help]
```

Lists the hardware resources of a managed system.

-r	- the type of resources to be listed: io - I/O mem - memory proc - processor virtualio - virtual I/O
--rsubtype	- the subtype of resource to be listed: unit - I/O unit bus - I/O bus slot - I/O slot or virtual I/O slot iopool - I/O pool taggedio - tagged I/O eth - virtual ethernet scsi - virtual SCSI serial - virtual serial
-m <managed system>	- the managed system's name
--filter "<filter data>"	- filters the data to be returned the format is: "filter_name1=value,filter_name2=value,..." or "filter_name1=value1,value2,...",..."
-F <attribute names>	- comma separated list of attributes to be output
--header	- prints a header of attribute names when -F is also specified
--help	- prints this help

Valid parameter combinations are as follows:

-r io	--rsubtype unit	[--filter units]
-r io	--rsubtype bus	[--filter units,buses]
-r io	--rsubtype slot	[--filter units,buses,slots,pools, lpar_ids lpar_names]


```

-r io      --subtype iopool  [--filter pools,lpar_ids | lpar_names]
-r io      --subtype taggedio [--filter lpar_ids | lpar_names]
-r mem     --level sys
-r mem     --level lpar      [--filter lpar_ids | lpar_names]
-r proc    --level sys
-r proc    --level lpar      [--filter lpar_ids | lpar_names]
-r proc    --level sharedpool
-r virtualio --subtype eth --level sys
-r virtualio --subtype eth --level lpar [--filter slots,vlans,lpar_ids |
                                         lpar_names]
-r virtualio --subtype hsopti [--filter pools,lpar_ids | lpar_names]
-r virtualio --subtype virtualopti [--filter pools,lpar_ids | lpar_names]
-r virtualio --subtype scsi      [--filter slots,lpar_ids | lpar_names]
-r virtualio --subtype serial --level lpar | openserial
                                         [--filter slots,lpar_ids | lpar_names]
-r virtualio --subtype slot --level lpar [--filter lpar_ids | lpar_names]
-r virtualio --subtype slot --level slot [--filter lpar_ids | lpar_names,
                                         slots]

```

List system configuration (Issyscfg)

This command lists the system configuration of a managed system (Example A-5).

Example: A-5 Command usage for Issyscfg.

```

Usage: Issyscfg -r sys | lpar | prof | sysprof
        [-m <managed system>]
        [--filter "<filter data>"]
        [-F <attribute names>]
        [--header]
        [--help]

```

Lists the system configuration of a managed system.

```

-r          - the type of objects to be listed:
             sys - managed systems
             lpar - partitions
             prof - partition profiles
             sysprof - system profiles
-m <managed system> - the managed system's name
--filter "<filter data>" - filters the data to be returned
                        the format is:
                        "filter_name1=value,filter_name2=value,..."
                        or
                        "'filter_name1=value1,value2,...',..."
-F <attribute names> - comma separated list of attributes to be
                        output
--header           - prints a header of attribute names when -F is
                        also specified
--help            - prints this help

```

Filters are as follows:

```

-r sys      none
-r lpar     lpar_names or lpar_ids, clusters
-r prof     lpar_names or lpar_ids, profile_names
-r sysprof  profile_names

```

Create (make) system configuration (mksyscfg)

This command creates partitions, partition profiles, or system profiles (Example A-6).

Example: A-6 Command usage for mksyscfg

```
Usage: mksyscfg -r lpar | prof | sysprof
        -m <managed system>
        -f <configuration file name> | -i "<configuration data>"
        [--help]
```

Creates partitions, partition profiles, or system profiles.

```
-r          - the type of system resource(s) to be created:
            lpar - partition
            prof - partition profile
            sysprof - system profile
-m <managed system> - the managed system's name
-f <config file name> - the name of the file containing the
                    configuration data for this command
                    the format is:
                    attr_name1=value,attr_name2=value,...
                    or
                    "attr_name1=value1,value2,...",...
-i "<config data>" - the configuration data for this command
                    the format is:
                    "attr_name1=value,attr_name2=value,..."
                    or
                    ""attr_name1=value1,value2,...",..."
--help     - prints this help
```

The valid attribute names for this command are:

```
-r lpar    required: name, profile_name,
                lpar_type (os400 | aixlinux)
            optional: lpar_id, min_mem, desired_mem, max_mem,
                desired_procs, min_procs, max_procs,
                io_slots (drc_index/slot_io_pool_id/is_required),
                boot_mode (norm | dd | sms | of | ds),
                proc_type (ded | shared),
                sharing_mode (norm | keep | cap | uncap),
                uncap_weight, load_source_slot,
                alt_load_source_slot, console_slot,
                alt_console_slot, ecs_slot,
                desired_proc_units, min_proc_units, max_proc_units,
                sfp_surveillance, auto_ipl, max_virtual_slots,
                virtual_opti_pool_id, hsl_opti_pool_id,
                desired_interactive, min_interactive, max_interactive,
                virtual_eth_adapters(slot_num/ieee_virtual_eth/
                    port_vlan_id/addl_vlan_ids/is_trunk/is_required)
                virtual_scsi_adapters(slot_num/device_attr/
                    server_lpar_id/server_lpar_name/
                    server_slot_num/is_required)
                virtual_serial_adapters(slot_num/device_attr/
                    supports_hmc/server_lpar_id/server_lpar_name/
                    server_slot_num/is_required)
                sni_device_ids, sni_windows,
                power_ctrl_lpar_ids, lpar_io_pool_ids,
                shared_pool_util_auth (0 | 1)
-r prof    required: name, lpar_id | lpar_name
            optional: all of the optional parameters supported for -r lpar
                    are supported, except cluster_id.
```

-r sysprof required: name, lpar_ids | lpar_names, profile_names

Remove system configuration (rmsyscfg)

This command removes a partition, a partition profile, or a system profile (Example A-7).

Example: A-7 Command usage for rmsyscfg

Usage: rmsyscfg -r lpar | prof | sysprof
-m <managed system>
[-n <object name>]
[-p <partition name>]
[--help]

Removes a partition, a partition profile, or a system profile.

-r	- specifies the type of object to be removed: lpar - partition prof - partition profile sysprof - system profile
-m <managed system>	- the managed system's name
-n <object name>	- the user defined name of the object to be removed
-p <partition name>	- the user defined name of the partition that has the profile to be removed
--help	- prints this help

List dumps (lsdump)

This command queries the managed system for available dumps (Example A-8).

Example: A-8 Command usage for lsdump

Usage: lsdump -m <managed system>
[--help]

Queries the managed system for available dumps.

-m <managed system>	- the managed system's name
--help	- prints this help

HMC CLI command attributes

Table A-7 lists the command attributes that are available, along with the commands in which they are valid, and provides a description of each command.

Table A-7 HMC CLI command attributes

Command attributes		
Attribute	Used in command	Description
activated_profile	lssyscfg	User defined name of the profile that was used when the partition was activated.
addl_vlan_ids	chhwres and lshwres	List of additional VLAN IDs assigned to the virtual ethernet adapter. Valid values are 2 - 4094, and the number of additional VLAN IDs is limited by the maximum number of VLAN IDs allowed per port (<i>max_vlans_per_port</i>). This attribute is only valid for IEEE-compliant virtual ethernet adapters.
alt_console_slot	chhwres, chsyscfg, lshwres lssyscfg, and mksyscfg	Complete physical location code of the slot that has the alternate console device. A value of none indicates that there is no alternate console device. <i>This attribute is only valid for OS/400 partitions.</i>
alt_load_source_slot	chhwres chsyscfg lshwres lssyscfg mksyscfg	Complete physical location code of the slot that has the alternate load source I/O device. A value of none indicates that there is no alternate load source I/O device. <i>This attribute is only valid for OS/400 partitions.</i>
auto_ipl	chsyscfg lssyscfg mksyscfg	Indicates whether or not the partition should be activated automatically when the managed system is powered on. Valid values are 0 (no) and 1 (yes).
Attribute	Used in command	Description
backplane_phys_loc	lshwres	Physical location code of the backplane on which the bus resides.
boot_mode	chsyscfg lssyscfg mksyscfg	The boot mode for the partition. <i>This attribute is only valid for AIX/Linux partitions.</i> Valid values are norm (normal), dd (diagnostic default boot path), sms (boot to SMS menus), of (boot to Open Firmware prompt), and ds (diagnostic stored boot path).
bus_class	lshwres	Possible values are RIO or integrated .
bus_grouping	lshwres	
bus_id	lshwres	
bus_num	chhwres lshwres	RIO bus (loop) number to be represented by the SMC virtual I/O adapter
client_id	lshwres	Virtual serial and SCSI attribute.
client_type	lshwres	Virtual serial and SCSI attribute.
cluster_id	chsyscfg, lssyscfg, and mksyscfg	ID of the workload management (eWLM) cluster to which the partition is assigned. A valid cluster ID is a number between 0 and 32767, inclusive, or the value none , which indicates that the partition is not assigned to any eWLM cluster.

Command attributes		
cod_capable	Issyscfg	Indicates whether the managed system supports Capacity on Demand (CoD). Possible values are 0 (no) or 1 (yes).
config	lshwres	Virtual slot configuration state. Possible values are empty, ethernet, SCSI, serial, or SMC .
config_proc_units	lshwres	Total number of processing units assigned to the shared processor pool.
configurable_sys_mem	lshwres	Total amount, in megabytes, of configurable memory on the managed system.
configurable_sys_proc_units	lshwres	Total number of configurable processing units on the managed system.
console_slot	chhwres, chsyscfg, lshwres, Issyscfg, and mksyscfg	Complete physical location code of the slot that has the console device. <i>This attribute is only valid for OS/400 partitions.</i>
curr_avail_interactive	lshwres	Current percentage of interactive performance that is available on the managed system.
curr_avail_proc_units	lshwres	Current number of processing units in the shared pool which are available to be assigned to partitions.
curr_avail_sys_mem	lshwres	Current amount, in megabytes, of memory on the managed system that is not assigned to partitions
curr_avail_sys_proc_units	lshwres	Current number of processing units on the managed system that are not assigned to partitions.
curr_cycles	lshwres	
Attribute	Used in command	Description
curr_hyp_ipl_state	Issyscfg	State the hypervisor was last IPLed to. Possible values are run or standby .
curr_hyp_mem	lshwres	
curr_interactive	lshwres	A percentage. <i>This attribute is only valid for OS/400 partitions.</i>
curr_ipl_side	Issyscfg	Platform IPL side for the previous/current IPL. Possible values are perm (permanent) or temp (temporary).
curr_ipl_speed	Issyscfg	Platform IPL speed for the previous/current IPL. Possible values are slow or fast .
curr_ipl_speed_override	Issyscfg	Platform IPL speed override for the previous/current IPL. Possible values are none, slow, or fast .
curr_max_interactive	lshwres	A percentage. <i>This attribute is only valid for OS/400 partitions.</i>
curr_max_mem	lshwres	Maximum amount of memory, in megabytes, that the partition will support when running.
curr_max_procs	lshwres	
curr_max_proc_units	lshwres	
curr_max_virtual_slots	lshwres	Current maximum number of virtual slots for the partition.

Command attributes		
curr_mem	lshwres	Current amount of memory, in megabytes, which are owned by the partition.
curr_mem_region_size	lshwres	The current memory region size in megabytes.
curr_min_interactive	lshwres	A percentage. <i>This attribute is only valid for OS/400 partitions.</i>
curr_min_mem	lshwres	Minimum amount of memory, in megabytes, that the partition will support when running.
curr_min_procs	lshwres	
curr_min_proc_units	lshwres	
curr_procs	lshwres	
curr_proc_type	lshwres	
curr_proc_units	lshwres	
curr_shared_procs	lshwres	Current number of shared processors in the shared pool.
curr_sharing_mode	lshwres	
curr_sys_keylock	lssyscfg	Keylock position for the managed system for the previous/current power on. Possible values are norm (normal IPL) or manual (IPL to maintenance/service screen).
curr_total_avail_proc_units	lshwres	Current number of processing units on the managed system that are not assigned to partitions.
Attribute	Used in command	Description
curr_uncap_weight	lshwres	
cycles_per_sec	lshwres	Number of CPU cycles that execute in one second.
default_profile	chsyscfg and lssyscfg	User defined name of the current default partition profile. This name can be up to 31 characters in length.
definition_state	lssyscfg	Definition state of the partition. Possible values are undefined , not bootable , and valid .
desired_interactive	chsyscfg, lssyscfg, and mksyscfg	Desired interactive performance percentage for the partition. The value for this attribute must be a number between 0 and 100. <i>This attribute is only valid for OS/400 partitions.</i>
desired_mem	chsyscfg, lssyscfg, and mksyscfg	Desired amount of memory for the partition. The amount of memory must be specified in megabytes, and <i>must be a multiple of the memory region size</i> . The amount must not be smaller than the minimum memory (<i>min_mem</i>) for the partition, nor larger than the maximum memory (<i>max_mem</i>) for the partition.
desired_procs	chsyscfg, lssyscfg, and mksyscfg	When the partition is using dedicated processors, this is the desired number of dedicated processors for the partition. When the partition is using shared processors, this is the desired number of virtual processors for the partition.

Command attributes		
desired_proc_units	chsyscfg, lssyscfg, and mksyscfg	Desired number of processing units for the partition. This attribute is only valid when the processing mode is shared.
device_attr	chhwres and lshwres	Indicates whether the virtual SCSI or serial device is a client or server device. Valid values are client or server .
drc_name	lshwres	The DRC name of the I/O slot.
dump_type	lsdump	Type of hardware dump. Possible values are eventlog (platform event entry log dump), sp (service processor dump), sys (system hardware and hypervisor dump), pss (power subsystem dump), and sma (SMA adapter dump).
ecs_slot	chhwres, chsyscfg, lshwres, lssyscfg, and mksyscfg	Complete physical location code of the slot that has the ECS modem. A value of none indicates that there is no ECS modem. <i>This attribute is only valid for OS/400 partitions.</i>
full_aix_linux_boot_mode	lssyscfg	The boot mode used for the previous/current/next full system <i>AIX/Linux partition</i> activation. Possible values are norm (normal), dd (diagnostic default boot path), sms (boot to SMS menus), of (boot to Open Firmware prompt), and ds (diagnostic stored boot path).
full_os400_ipl_source	lssyscfg	The IPL source used for the previous/current/next full system <i>OS/400 partition</i> activation. Possible values are a , b , c , and d .
has_os400_vpd	lshwres	
hsl_capable	lssyscfg	Indicates whether the managed system supports HSL OptiConnect. Possible values are 0 (no) or 1 (yes).
Attribute	Used in command	Description
hsl_opti_pool_id	chhwres, chsyscfg, lshwres, lssyscfg, and mksyscfg	ID of the HSL OptiConnect pool to which the partition is assigned. A value of 0 indicates HSL OptiConnect is disabled for the partition.
hyp_capabilities	lssyscfg	Hypervisor capabilities.
ieee_virtual_eth	chhwres and lshwres	Indicates whether the virtual ethernet adapter is IEEE-compliant. Possible values are 0 (no) or 1 (yes).
iop_info	lshwres	
io_pool_id	lshwres	I/O pool ID.

Command attributes		
io_slots	chsyscfg, lssyscfg, and mksyscfg	<p>List of I/O slots for the partition. Each item in this list has the format: phys_loc/ slot_io_pool_id/is_required</p> <p>Note that the attribute names are not present in the list, just their values are present. For example, U47070041076RX5L1-P2-C3/1/2/1 specifies an I/O slot with a physical location code of U47070041076RX5L1-P2-C3, it is assigned to I/O pool 2, and it is a required slot.</p> <p>If an attribute is optional and is not to be included, then no value would be specified for that attribute. For example, U47070041076RX5L1-P2-C3//1 specifies an I/O slot that is not assigned to an I/O pool.</p>
ip_addr	lssyscfg	Managed system's IP address.
ipl_side	chsyscfg	Platform IPL side for the next IPL. Valid values are perm (permanent) or temp (temporary).
ipl_source	lssyscfg	IPL source for the partition. Valid values are a , b , c , and d . <i>This attribute is only valid for OS/400 partitions.</i>
ipl_speed	chsyscfg	Platform IPL speed for the next IPL. Valid values are slow or fast .
ipl_speed_override	chsyscfg	Platform IPL speed override for the next IPL. Valid values are none , slow , or fast .
is_required	chhwres and lshwres	Indicates whether the I/O slot or virtual I/O adapter is required for the partition. Valid values are 0 (no) and 1 (yes).
is_trunk	chhwres and lshwres	Indicates whether the virtual ethernet adapter is a trunk adapter. Valid values are 0 (no) and 1 (yes).
local_slot_num	lssyscfg	Local virtual slot number for the SMC virtual I/O adapter.
local_hub_drc_index	lssyscfg	DRC index for the local hub for the SMC virtual I/O adapter.
logical_serial	lssyscfg	Logical serial number of the partition.
lpar_config	lssyscfg	Current LPAR configuration state of the managed system. Possible values are full (full system partition configuration) or norm (normal lpar configuration).
Attribute	Used in command	Description
lpar_id	chsyscfg, lshwres, lssyscfg, and mksyscfg	Partition ID. Valid values are 1 through the maximum number of partitions supported on the managed system (<i>max_lpars</i>).
lpar_ids	chsyscfg, lshwres, lssyscfg, mksyscfg	List of partition IDs.

Command attributes		
lpar_io_pool_ids	chsyscfg, lshwres, and mksyscfg	List of IDs of the I/O pools in which the partition is participating. A valid I/O pool ID is a number between 0 and the maximum number of I/O pools supported on the managed system (<i>max_io_pools</i>) - 1. A value of none , which indicates that the partition is not participating in any I/O pools, is also valid.
lpar_keylock	lssyscfg	Partition keylock position. Possible values are norm (normal) or manual (manual).
lpar_name	chsyscfg, lshwres, and mksyscfg	Partition name.
lpar_names	chsyscfg, lshwres, lssyscfg, and mksyscfg	List of partition names.
lpar_type	mksyscfg	Partition type. Valid values are aixlinux and os400 .
load_source_slot	chsyscfg, lshwres, lssyscfg, and mksyscfg	Complete physical location code of the slot that has the load source I/O device. This attribute is only valid for OS/400 partitions
mac_addr	lshwres	MAC address for the virtual ethernet adapter.
mac_prefix	chhwres and lshwres	The first 3 bytes of the MAC address to be assigned to all virtual ethernet adapters for this managed system. This must be specified as a 3 byte hex value.
max_interactive	chsyscfg, lssyscfg, and mksyscfg	Maximum interactive performance percentage for the partition. The value for this attribute must be a number between 0 and 100. This attribute is only valid for OS/400 partitions.
max_io_pools	lshwres	Maximum number of I/O pools supported on the managed system.
max_lpars	lssyscfg	Maximum number of partitions supported on the managed system.
max_mem	chsyscfg, lssyscfg, and mksyscfg	Maximum amount of memory that the partition will support. The amount of memory must be specified in megabytes, and <i>must be a multiple of the memory region size</i> . The amount must also be greater than or equal to the desired memory (<i>desired_mem</i>) for the partition.
max_optimal_interactive	lshwres	Maximum optimal interactive performance percentage for the specified processing units value. This attribute is only valid for OS/400 partitions.
max_power_ctrl_lpars	lssyscfg	Maximum number of power controlling partitions.
Attribute	Used in command	Description
Attribute	Used in command	Description
max_procs	chsyscfg, lssyscfg, and mksyscfg	
max_proc_units	chsyscfg, lssyscfg, and mksyscfg	

Command attributes		
max_shared_pools	lshwres	Maximum number of shared processing pools which are supported on the managed system.
max_virtual_slots	chsyscfg, lssyscfg, and mksyscfg	Maximum number of virtual slots for the partition. Valid input values are 2 - 65535. The default value is 4.
max_vlans_per_port	lshwres	Maximum number of supported VLAN IDs per virtual ethernet port.
mem_region_size	chhwres	The memory region size, in megabytes, for the managed system.
min_interactive	chsyscfg, lssyscfg, and mksyscfg	Minimum interactive performance percentage for the partition. The value for this attribute must be a number between 0 and 100. <i>This attribute is only valid for OS/400 partitions.</i>
min_mem	chsyscfg, lssyscfg, and mksyscfg	Minimum amount of memory that the partition will support. The amount of memory must be specified in megabytes, and <i>must be a multiple of the memory region size</i> . The amount cannot exceed the desired memory (<i>desired_mem</i>) for the partition, and must be greater than or equal to the required minimum memory (<i>required_min_mem</i>) for the partition.
min_optimal_interactive	lshwres	Minimum optimal interactive performance percentage for the specified processing units value. <i>This attribute is only valid for OS/400 partitions.</i>
min_procs	chsyscfg, lssyscfg, and mksyscfg	
min_proc_units	chsyscfg, lssyscfg, and mksyscfg	
min_proc_units_per_processor	lshwres	Smallest value of processing units which is valid per virtual processor.
msg_passing_capable	lssyscfg	Possible values are 0 (no) or 1 (yes).
name	chsyscfg, lssyscfg, and mksyscfg	The name of an object. Maximum name lengths are: <ul style="list-style-type: none"> • managed system names - 63 characters • partition names - 47 characters • profile names - 31 characters
new_name	chsyscfg	The new name to be assigned to an object. Maximum name lengths are: <ul style="list-style-type: none"> • managed system names - 63 characters • partition names - 47 characters • profile names - 31 characters
op_panel_value	lssyscfg	The op panel data for the partition.
Attribute	Used in command	Description
os_type	lssyscfg	The partition's operating system. Possible values are aixlinuxRPA and os400 .
os_version	lssyscfg	Version of the operating system currently running on the partition.

Command attributes		
os400_capable	Issyscfg	Indicates whether the managed system supports OS/400 partitions. Possible values are 0 (no) or 1 (yes).
parent_slot	Ishwres	Complete physical location code of the parent slot.
pci_class	Ishwres	
pci_device_id	Ishwres	
pci_revision_id	Ishwres	
pci_subs_device_id	Ishwres	
pci_subs_vendor_id	Ishwres	
pci_vendor_id	Ishwres	
pend_avail_interactive	Ishwres	A percentage. <i>This attribute is only valid for OS/400 partitions.</i>
pend_avail_proc_units	Ishwres	Pending number of processing units in the shared pool which are available to be assigned to partitions.
pend_avail_sys_mem	Ishwres	
pend_avail_sys_proc_units	Ishwres	
pend_hyp_ipl_state	Issyscfg	State to IPL the hypervisor to next. Possible values are run or standby .
pend_hyp_mem	Ishwres	
pend_interactive	Ishwres	A percentage. <i>This attribute is only valid for OS/400 partitions.</i>
pend_ipl_side	Issyscfg	Platform IPL side for the next IPL. Possible values are perm (permanent) or temp (temporary).
pend_ipl_speed	Issyscfg	Platform IPL speed for the next IPL. Possible values are slow or fast .
pend_ipl_speed_override	Issyscfg	Platform IPL speed override for the next IPL. Possible values are none , slow , or fast .
pend_max_interactive	Ishwres	A percentage. <i>This attribute is only valid for OS/400 partitions</i>
pend_max_procs	Ishwres	
pend_max_mem	Ishwres	
pend_max_proc_units	Ishwres	
pend_max_virtual_slots	Ishwres	Virtual slots
Attribute	Used in command	Description
pend_mem	Ishwres	
pend_mem_region_size	Ishwres	The memory region size to be used after the next power on or IPL of the managed system.

Command attributes		
pend_min_interactive	lshwres	A percentage. <i>This attribute is only valid for OS/400 partitions.</i>
pend_min_mem	lshwres	
pend_min_procs	lshwres	
pend_min_proc_units	lshwres	
pend_procs	lshwres	
pend_proc_type	lshwres	
pend_proc_units	lshwres	
pend_shared_procs	lshwres	Pending number of shared processors in the shared pool.
pend_sharing_mode	lshwres	
pend_sys_keylock	lssyscfg	Keylock position for the managed system to be used for the next power on. Possible values are norm (normal IPL) or manual (IPL to maintenance/service screen).
pend_total_avail_proc_units	lshwres	
pend_uncap_weight	lshwres	
phys_loc	lshwres	Complete physical location code of the slot.
port_vlan_id	chhwres lshwres	Port virtual LAN ID for the virtual ethernet adapter. <i>Valid values are 1 - 4094.</i>
power_ctrl_lpar_ids	chsyscfg lssyscfg mksyscfg	List of IDs of power controlling partitions for the partition. A power controlling partition has the authority to power the partition on and off
power_ctrl_lpar_names	chsyscfg lssyscfg mksyscfg	List of user defined names of power controlling partitions for the partition. A power controlling partition has the authority to power the partition on and off.
power_off_policy	chsyscfg lssyscfg	Power off policy for the managed system. Valid values are 0 (stay running after the last partition powers off) and 1 (power off when the last partition powers off).
proc_type	chsyscfg lssyscfg mksyscfg	Type of processors the partition will use. Valid values are ded (dedicated) or shared (shared).
profile_name	mksyscfg	The name of the default profile to be created. <i>This name can have a maximum length of 31 characters.</i>
profile_names	chsyscfg lssyscfg mksyscfg	
Attribute	Used in command	Description
recent_alt_load_source_slot	lshwres	Complete physical location code of the slot. This attribute is only valid for OS/400 partitions.

Command attributes		
remote_lpar_id	chhwres lshwres	For client adapters, this specifies the ID of the partition which has the hosting (server) virtual serial/SCSI adapter for this adapter. For server adapters, this specifies the ID of the partition which has the only client virtual serial/SCSI adapter allowed to connect to this adapter. A value of any indicates that any client virtual serial/SCSI adapter should be allowed to connect to this adapter.
remote_lpar_name	chhwres lshwres	For client adapters, this specifies the name of the partition which has the hosting (server) virtual serial/SCSI adapter for this adapter. For server adapters, this specifies the name of the partition which has the only client virtual serial/SCSI adapter allowed to connect to this adapter.
remote_slot_num	chhwres and lshwres	For client adapters, this specifies the virtual slot number of the virtual serial/SCSI adapter that is the hosting (server) adapter for this adapter. For server adapters, this specifies the virtual slot number of the only client virtual serial/SCSI adapter allowed to connect to this adapter. A value of any indicates that any client virtual serial/SCSI adapter should be allowed to connect to this adapter.
remote_smc_lpar_id		ID of the partition on the remote system the SMC virtual I/O adapter will pair with.
remote_smc_lpar_name		Name of the partition on the remote system the SMC virtual I/O adapter will pair with.
remote_smc_slot_num		Virtual slot number on the remote system the SMC virtual I/O adapter will pair with.
required_min_mem	lshwres	Required minimum memory amount, in megabytes, for the specified maximum memory amount for the partition.
rmc_active	lssyscfg	Indicates whether or not the partition has an active connection to RMC. <i>This attribute is only valid for AIX/Linux partitions.</i>
run_hyp_min_mem	lshwres	
run_interactive	lshwres	A percentage. <i>This attribute is only valid for OS/400 partitions.</i>
run_mem	lshwres	Current amount of memory, in megabytes, that the partition has varied on.
run_min_mem	lshwres	Minimum memory value returned by a running partition's operating system.
run_procs	lshwres	Number of processors (or virtual processors if the processor type is shared) that are varied on for the partition.
run_proc_units	lshwres	Number of processing units that are varied on for the partition.
run_uncap_weight	lshwres	
Attribute	Used in command	Description
serial_num	lssyscfg	Serial number of the managed system.

Command attributes		
service_lpar_id	chsyscfg and lssyscfg	For chsyscfg , this specifies the ID of the partition to be given service authority immediately. For lssyscfg , this shows the ID of the partition that currently has service authority
service_lpar_name	chsyscfg and lssyscfg	For chsyscfg , this specifies the name of the partition to be given service authority immediately. For lssyscfg , this shows the name of the partition that currently has service authority.
sfp_surveillance	chsyscfg, lssyscfg, and mksyscfg	Service focal point surveillance.
shared_pool_id	lshwres	
shared_pool_util_auth	chsyscfg, lssyscfg, and mksyscfg	Indicates whether the partition has authority to access the shared processor pool utilization data pertaining to that partition. Valid values are 0 (no) and 1 (yes).
sharing_mode	chhwres, chsyscfg, lssyscfg, and mksyscfg	Processor sharing mode for the partition. Valid values for partitions using shared processors are cap (shared capped) and uncap (shared uncapped). Valid values for partitions using dedicated processors are norm (dedicated normal) and keep (dedicated keep). Transition from uncapped to capped only if uncapped weight is 0.
slot_io_pool_id	chhwres lshwres	ID of the I/O pool to which the slot is assigned. Valid values are 0 through the maximum number of I/O pools supported on the managed system (<i>max_io_pools</i>) - 1. A value of none , which indicates that the slot is not assigned to any I/O pools, is also valid.
slot_num	lshwres	Virtual slot number. Valid input values are 2 through the maximum number of virtual slots for the partition (<i>max_virtual_slots</i>) - 1.
slots	lshwres	List of slots (complete physical location codes).
sni_config_mode	chsyscfg lssyscfg mksyscfg	
sni_device_ids	chsyscfg lssyscfg mksyscfg	
sni_windows	chsyscfg lssyscfg mksyscfg	
sp_boot_attr	lssyscfg	Service processor boot attributes.
sp_boot_major_type	lssyscfg	Service processor boot major type.
sp_boot_minor_type	lssyscfg	
sp_capabilities	lssyscfg	
sp_version	lssyscfg	
Attribute	Used in command	Description
src	lssyscfg	

Command attributes		
state	lshwres and lssyscfg	
status	lshwres	
supports_hmc	chhwres and lshwres	
sys_ipl_attr	lssyscfg	
sys_ipl_major_type	lssyscfg	
sys_ipl_minor_type	lssyscfg	
time	lssyscfg	
total_cycles	lshwres	
total_proc_units	lshwres	
type_model	lssyscfg	
uncap_weight	chhwres, chsyscfg, lssyscfg, and mksyscfg	
unit_id	lshwres	
unit_model	lshwres	
unit_serial_num	lshwres	
utilized_cycles	lshwres	
virtual_eth_adapters	chsyscfg, lssyscfg, and mksyscfg	<p>List of virtual ethernet adapters. Each item in this list has the format:</p> <p>slot_num/ieee_virtual_eth/ port_vlan_id/addl_vlan_ids/ is_trunk/is_required</p> <p><i>Note that the attribute names are not present in the list, just their values are present.</i></p> <p>For example: 3/1/5/"6,7,8"/0/1 specifies a virtual ethernet adapter with a virtual slot number of 3, is IEEE compliant, has a port vlan ID of 5, additional vlan IDs of 6, 7, and 8, it is not a trunk adapter, and it is required.</p> <p>If an attribute is optional and is not to be included, then no value would be specified for that attribute. For example: 3/1/5//0/1 specifies a virtual ethernet adapter with no additional vlan IDs.</p>
virtual_opti_pool_id	chhwres, chsyscfg, lshwres, lssyscfg, and mksyscfg	ID of the virtual OptiConnect pool to which the partition is assigned. A value of 0 indicates virtual OptiConnect is disabled for the partition. <i>This attribute is only valid for OS/400 partitions.</i>
Attribute	Used in command	Description

Command attributes		
virtual_scsi_adapters	chsyscfg, lssyscfg, and mksyscfg	<p>List of virtual SCSI adapters. Each item in this list has the format:</p> <pre>slot_num/device_attr /remote_lpar_id/remote_lpar_name /remote_slot_num/is_required</pre> <p><i>Note that the attribute names are not present in the list; only their values are present. If an attribute is optional and is not to be included, then no value would be specified for that attribute.</i></p> <p>For example: 4/client//lpar2/3/0 specifies a virtual client SCSI adapter with a virtual slot number of 4, a server partition name of lpar2, and a server slot number of 3, and is not required. The server partition ID was omitted.</p>
virtual_serial_adapters	chsyscfg, lssyscfg, and mksyscfg	<p>List of virtual serial adapters. Each item in this list has the format:</p> <pre>slot_num/device_attr/supports_hmc /remote_lpar_id/remote_lpar_name /remote_slot_num/is_required</pre> <p><i>Note that the attribute names are not present in the list; only their values are present. If an attribute is optional and is not to be included, then no value would be specified for that attribute.</i></p> <p>For example: 8/server/1///1 specifies a virtual server serial device that supports HMC, and has a virtual slot number of 8, and is required. The server partition ID, name, and slot number were all omitted.</p>
vpd_model	lshwres	
vpd_serial_num	lshwres	
vpd_type	lshwres	

Glossary

CCIN_Custom Card Identification Number

CoD_Capacity Upgrade on Demand

CSU_Customer Set Up.

DHCP_Dynamic Host Configuration Protocol

DNS_Domain Name Server

FRU_Field Replaceable Unit

HMC_Hardware Management Console.

HSL_High Speed Link

MTMS_Machine Type Machine Serial

PTF_Program Temporary Fix

SMA_Switch Mode Adapter

Term3. Term3 definition.

VPD_Vital Product Data

Archived

Index

Numerics

5250 client 99
5250 console 46
5250 Emulator 57
5250 OLTP 155
5250 virtual terminal 99
570 node 52
7310-C03 47
7310-CR2 47
7316 48
7316-TF2 48

A

Abnormal IPL 87
Additional ethernet LAN 50
Administrator mailing address 128
Advanced Operator 142, 238
Advanced System Manager 55
AIX boot mode 64
AIX error log ID 101
AIX partition 2
Alternate Console 168
Alternate IPL 140
ASM interface 73
Autodetection 117
Automatic allocation 119
Automatic reboot function 96
Automatically boot Partition 170
Automatically boot with platform 183
Automatically starting a partition 174

B

Backup critical console data 212
Backup profile data to disk 73
Boot section 192
Breakout cable 50
Bus numbering 253

C

CA 236
Call home 103
Capacity on demand option 72
Capacity Upgrade on Demand 32, 65
Capped shared processor 149
CCIN 66
Certificate Authority 41, 43, 236–237
Certificate authority 235
Change current locale 99
Change HMC Date 110
Change HMC Time 110
Change root password 112
Change the HMC passwords 106

Changing 175
Changing passwords 238
Changing the default profile 139, 175
Changing the IPL side 191
CLI 141
Cloning HMC Configuration 248
Cloning HMC Configurations 245
CoD Click to Accept 62, 214
Collect VPD information 100
Configure connectivity 106
Configure Dial-up 130
Configure DNS 115
Configure HMC Connectivity 129
Configure network settings 106
Configure Service Focal Point 106
Configured (bootable) 180
Connecting SSH client to the HMC 206
Console 165
Console function 159
Contact Information 128
Contiguous blocks of memory 26
Control Panel 252
Control Panel function 21 211
Control Panel function 22 211
Control Panel Function 3 210
Control Panel function 8 210
Create additional HMC users 113
Create new partition option 72
Create new user ids 106
Create partition profile 74
Create partition wizard 142
Create serviceable event 103
Creating a partition 141
Creating a user 238
Creating additional partition profiles 139
Creating an OS/400 logical partition 139
Creating Virtual IO Adapters 158
Critical Console Back Up 248
Cryptographic keys 41
CSU 48
CUoD 32
CuOD 3
Custom Card Identification Number 66
Customer contact information 106
Customer information 104
Customer Set Up 48
Customize date and time 98
Customize network settings 98

D

D mode IPL 163
DBCS 108
DCHP server 55
Dedicated processor 149

- Dedicated processors 149, 185
- Dedicated Service Tools 199
- Default gateway address 98
- Default gateway device 120
- Default partition profile 140
- Default profile 174
- Delayed partition shut down 87
- Deleting a partition 201
- Deleting a partition profile 202
- Deleting a user 238
- Desired 3
- Desired memory 148
- Desired processing units 152
- Desired processors 150
- Desktop HMC 47
- DHCP 53, 98
- DHCP client 98
- DHCP server 50, 55, 98, 118
- Dial prefix values 130
- Dial-up from the local HMC 129
- Digital certificate 41
- Direct Operations Console 32
- Disable remote command execution 99
- Disable remote virtual terminal 99
- Disconnect another HMC 72
- Disk unit IOP reset/reload 203
- Display System I/O Resources 253
- DLPAR 140, 195
- DLPAR function 91
- DNS 53
- DNS server 55
- DNS server IP address 115
- Domain name 55, 98, 114
- Domain suffix 116
- Double white button push 210
- DST 199, 252
- Dump 62, 214
- Dumps 103
- DVD-RAM 98, 212

E

- e-Config 48
- Edit MTMS 104
- Editing user information 238
- Electronic Customer Support 164
- Electronic Service Agent 129
- Emulator 56
- Enable remote command execution 99
- Enable remote virtual terminal 99
- Enabling SSH 205
- Enclosure ID 104
- Error class 102
- ESC 164
- Example SSH command 207

F

- Failed authentication 62, 213
- Failing device system name 102
- Fast power-on 64

- Firewall 53, 98, 123
- Fluxbox desktop 56
- Format removable media 212
- Function 8 72

G

- GA1 xv
- Gateway device 120
- Gateway IP address 49
- General Availability xv
- Guided Setup checklist 106
- Guided Setup planning 105
- Guided Setup wizard 53, 107

H

- Hardware 181
- Hardware conditions 45
- Hardware IO
 - Advanced options button 90
- Hardware Management Console 2
- Hardware page table 26
- HMC 2
- HMC code level 106
- HMC commands 208
- HMC connecting through other systems 129
- HMC console 56
- HMC critical data 98
- HMC Domain, DNS, IP address range 106
- HMC failed 62, 214
- HMC Guided Setup 106
- HMC Guided Setup wizard checklist 106
- HMC ids/passwords 106
- HMC initial password 108
- HMC language locale 106
- HMC Management Environment 56
- HMC Navigation Area 142
- HMC network configuration 106
- HMC Networking 53
- HMC Networking setup 105
- HMC Open Networking 55
- HMC Password 49
- HMC physical setup 50
- HMC redundancy 46
- HMC roles 106, 142
- HMC Service setup 105
- HMC setup 105
- HMC support xv
- HMC system wide options 71
- HMC Terminal menu 56
- HMC User management 235
- HMC user profile 142
- HMC Users 238
- Hosted guest partition 159, 170
- Hosted Linux partition 97
- Hostname 98, 209
- HPT 26
- HSL Opticonnect 169
- Hub 54
- Hypervisor 2, 26, 66, 181

I

I/O resources 3, 155
I/O slot view 66
i5/OS hang 86
i5/OS partition 2
IBM Service 3
IBM Service and Support 54
ibm5250 57
Identify LED processing 104
Immediate partition shut down 87
Immediate reboot 193
iNav 94
Inbound connectivity settings 104
Initial tour of the HMC desktop 56
Initialize profile data 73
Initializing 62, 213
Install hardware 100
Install/Add/Remove/hardware 103
Installing the HMC 50
Interactive capacity 155
Inventory scout profile configuration 100
Inventory Scout Service 212
Inventory scout service 99
Inventory Scout Services 129
Inventory Scout services 100
IO pool option 90
IOA 60
IOP 60, 162
IPL from external media 64
IPL Source 192
IPL source 96
iSeries control panel functions vs HMC options 210
iSeries service tools 46
Isolated processors 185

J

Japanese locale 108

K

Keyboard 47
Keylock 96
Keylock position 192

L

Lamp test 104
LAN "health check" 104
LAN adapter speed 117
LAN Console 32
Language 99
Language locale setup 107
Launch ASM menu 103
Launch the eServer Information 49
Launch the Guided Setup Wizard 108
Licensed Internal Code 162
Linux boot mode 64
Linux partition 2
Linux system error log 101
LMB 25

Load source 140, 162
Load source IOP 162
Locale 99
Locale and time zone 106
Location Codes 156
Lock 58
Lock menu 58
Logging off the HMC 58
Logical Memory Block 25
Logical Partitioning Validation Tool 253
Logical path 156
Logical path location codes 157
LPAR components 2
LPAR Validation Tool 252
LVT 67, 252–253

M

Machine Type Machine Serial 102
Main storage dump 194
Main storage dumps 200
Manage Dumps 104
Manage problem data 103
Manage Serviceable Event 102
Manage serviceable events 100
Managed Server properties 26
Managed system 46
Managed systems 45
Manually starting a powered off partition 191
Maximum memory 148
Maximum memory parameter 26
Maximum number of partitions 62
Maximum processing units 152
Maximum processors 150
Memory 184, 195, 197
Memory granularity 26
Memory Region 25
Memory resources 3
MES 3
MES upgrade 62
MHxxxx 63
Migrating existing VLAN 159
Migrating Linux Partitions 251
Migration Examples 251
Migration Planning 251–252
Migration to HMC 251
Minimum memory 148
Minimum number of virtual processors 154
Minimum processing units 152
Minimum processors 150
Miscellaneous profile settings 170
Miscellaneous tab 96
Modem Configuration 130
Mouse 47
MSD 194, 200
MSD IPL 200
MTMS 102, 104
Multi-partitioned environment 45

N

- Native IO support 5
- Net menu 57
- Network configuration 114
- No connection 62, 214
- Nonroutable IP address ranges 119
- Non-volatile random access memory 2
- Not configured (not-bootable) 180
- Number of connected minutes 104
- Number of disconnected minutes 104
- Number of minutes between outages 104
- NVRAM 2, 248

O

- Object Manager Security 238
- Object manager security 235
- OEM display 48
- Open Terminal Window 75
- Opera Browser 57
- Operating 62, 213
- Operation console device 164
- Operations Console 2, 32
- Operator 238
- Opticonnect 169, 189
- Optional languages 99
- OS/400 error log ID 101
- Outbound connectivity settings 104

P

- Parallel ports 47
- Partition ID 144, 184
- Partition id 88
- Partition Name 144, 184
- Partition operating system 184
- Partition options 74
- Partition Profile 3
- Partition profile 139
- partition profile 140
- Partition profile name 147
- Partition profile properties 184
- Partition properties 88
 - General tab 88
 - Hardware tab 89
 - Miscellaneous 96
 - Processors and Memory 91
 - Settings 96
 - virtual devices 92
- Partition Type 144
- Partitions folder 142
- Pass-Through System 133
- Pending authentication 62, 213
- Performing service functions 198
- Permanently applied PTF 64
- pHyp 26
- Physical Adapters 195
- Physical adapters 195
- Physical location code 156
- Plain Socket 238
- Port Virtual LAN ID 158

- Post Guided Setup tasks 137
- Power control 189
- Power control partition 170
- Power Controlling partition 96
- Power Controlling Partitions 183
- Power Down System 87
- Power down System 72
- Power off option 72
- Power Off policy check box 63
- Power On parameters tab 64
- Powering down a partition 191
- Powering Off 62, 214
- PowerPC 5
- Predefined partition profile 61
- Primary partition 96, 252–253
- Printer support 47
- Private and open network 53
- Private direct network 53
- Private direct networking 53
- Private indirect network 53
- Private indirect networking 54
- Private key 41, 43, 236
- Problem Determination 99
- Problem determination
 - Service applications 99
- Problem Events display 134
- Processor resources 3
- Processors 195–196
- Processors and Memory 181
- Profile data 73
- Profile information 212
- Profile name 88
- Profile summary 171
- PTF listings 103
- Public key 41
- PuTTY 99, 206
- PuTTY SSH client. 206
- PWRDWN SYS 72, 194

Q

- QSECOFR 111

R

- Rack config 252
- Rack configuration 66
- Rack mounted HMC 48
- Rack mounted HMC system unit 50
- rcshell 56
- Real Memory Offset 26
- Rear view of desktop HMC ports 50
- Rearside USB port 50
- Reboot Partition 193
- Rebuild the managed system 73
- Rebuilding the managed system 73
- Recalculate memory 26
- Recover Configuration Data 252
- Recovery 62, 214
- Recovery indicator 73
- Red Hat 5

- Redbooks Web site
 - Contact us xiii
- Redundant HMC configurations 245
- Remote access 123
- Remote management 205
- Remote service 99
- Remote Service and Support Facility 164, 200
- Remote support 104
- Remote Support Facility 104, 129
- Remote Support Information panel 128
- Remote support request 104
- Remote Technical Support 100
- Remove managed system connection 71
- Remove profile data 73
- Repair serviceable event 100–102
- Replace parts 100, 103
- Required 3
- Reset HMC connection 71
- Resource allocation printout 253
- Resource configuration 180
- Resource conflict 3
- Restart inventory scout daemon 100
- Restart Partition 86
- Restarting an operational partition 191
- Restore profile data 73
- Restricted shell 56
- RJ-45 connector 48
- RMO 26
- Root Password 112
- Routing information 98
- RSF 104
- RSSF 164, 200

S

- SA 104
- SAR 254
- Save to DVD 103
- Save upgrade data 212
- Schedule operations 98
- Secure Sockets Layer 41, 236, 238
- Secure system manager 237
- Server processor 50
- Server Security 237
- Server security 235
- Service Agent 2, 45, 104, 134, 212
- Service Applications 99
- Service Authority 183
- Service connection type 106
- Service Focal Point 100, 129, 190
- Service focal point 45, 99
- Service Partition 2–3, 63
- Service Partition p5 servers 63
- Service Processor 2, 52, 55
- Service processor code level 106
- Service Processor code support xv
- Service processor level 62
- Service Processor static IP address 55
- Service processor version 62, 216
- Service Program Agreement 130
- Service provider information 106

- Service Representative 238
- Service strategy 45
- Service Tools 253
- Service utilities 100, 103
- Serviceable Event Details 101
- Set HMC date and time 106
- Shared partition indicator 91
- Shared processor utilization 182
- Shut Down Partition 87
- Shut Down Timer 72
- Shutdown Console 59
- SHUTDTIMO 72
- Shutting down or logging off the HMC 58
- Single virtual system bus 158
- SLIC 2
- Slow power-on 64
- Software Knowledge Base Documents 164
- Software requirements 5
- Solutions Assurance Review 254
- SPCN cables 52
- Specifying console resource 166
- SSH client 206
- SSH protocol 99
- SSL 41, 236, 238
- SST 252
- Standby 62, 214
- Start here for hardware 48
- Starting partitions 191
- Starting the Web-based client 208
- Storage resources 160
- STRSST 253
- Summary screen 135
- Supported languages 107
- SUSE 5
- Switch 54
- Switchable configuration 90
- System Administrator 142, 238
- System Administrator role 236
- System attention LED 104
- System Licensed Internal Code 2
- System Manager Security 236
- System Name 184
- System preferences 212
- System Profile 3, 139
- System profile 140
- System properties 62
- System Reference Code 199

T

- T prefix 157
- Tagging partition resources 161
- Temporary applied PTF 64
- Test network connectivity 98
- Time zone 110
- Total system processors 185
- Trusted central administrative entity 41
- Twinax 2
- Twinax consol 32
- Types of console 32

U

Unassigned cards/slots 66
Uncapped partition 185
Uncapped shared processor 149
Uncapped shared processor partition 154
Update managed system password 73
USB ports 47
User Administrator 238
Uses of partition profiles 140
Using the HMC as the console 165

V

Validated partition profile 3
View console events 98
View Guided Setup Wizard Log 136
Viewer 238
Viewing user information 238
Virtual Adapter 75
Virtual Adapters 94
Virtual console 46
Virtual Devices 182
Virtual Ethernet 158
Virtual ethernet 92
Virtual ethernet adapter 94
Virtual IO Adapters 195
virtual LAN ID 158
Virtual panel option 1 210
Virtual panel option 7 210
Virtual port 93
Virtual private network 129
Virtual SCSI 92, 158, 160
Virtual Serial 158–159
Virtual serial 92
Virtual serial client adapter 159
Virtual serial server adapter 160
Vital Product Data 100
VLAN 94
VPD 100
VPN 117, 129, 133

W

Web-based System Management Remote Client 46, 55
Web-based System Manager Remote Client 42
White button power on 210
WLM 96
Work Load Management 183
Work Load Management group 96
Work Load Management groups 96
Work with System Partitions 251
Working with IOP functions 203
Working with remote service 200
WRKCNTINF 127
WSMRC 55



Redbooks

Logical Partitions on System i5

A Guide to Planning and Configuring LPAR with HMC on System i

(1.0" spine)

0.875" x 1.498"

460 <-> 788 pages



Logical Partitions on System i5

A Guide to Planning and Configuring LPAR with HMC on System i



Understand the new Logical Partitions for IBM Power5 architecture

This IBM Redbook gives a broad understanding of the new System i5 architecture as it applies to logically partitioned System i5 systems. This functionality is delivered through a new configuration and management interface called the Hardware Management Console (HMC).

Learn how to install, configure, and manage LPAR with the latest HMC

Reading this book will help you design your server partition scheme from scratch. We also discuss the requirements to create a solution to migrate from existing iSeries servers with and without logical partitions.

Discover how to implement OS/400 logical partitions

This book will help you install, tailor, and configure logical partitions (LPARs) on System i5 systems. You will understand how to set up the server and HMC via the Guided Setup Wizard. We also cover the creation of multiple i5/OS partitions, where the OS could be IBM i5/OS, AIX, or Linux.

**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**

SG24-8000-01

ISBN 0738496251