


Implementing POWER Linux on IBM System i Platform

Planning and configuring Linux servers
on IBM System i platform

Linux distribution on IBM System i
Platform installation guide

Tips to run Linux servers on
IBM System i platform



Yessong Johng
Erwin Earley
Rico Franke
Vlatko Kosturjak

Redbooks



International Technical Support Organization

Implementing POWER Linux on IBM System i Platform

February 2007

Archived

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

Archived

Second Edition (February 2007)

This edition applies to i5/OS V5R4, SLES10 and RHEL4.

© Copyright International Business Machines Corporation 2005, 2007. 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	vii
Trademarks	viii
Preface	ix
The team that wrote this redbook	ix
Become a published author	xi
Comments welcome	xi
Chapter 1. Introduction to Linux on System i platform	1
1.1 Concepts and terminology	2
1.1.1 System i platform	2
1.1.2 Hardware management console	4
1.1.3 Virtual Partition Manager (VPM)	10
1.2 Brief introduction to Linux and Linux on System i platform	12
1.2.1 Linux on System i platform	12
1.3 Differences between existing Power5-based System i and previous System i models	13
1.3.1 Linux enhancements on Power5 / Power5+	14
1.4 Where to go for more information	15
Chapter 2. Configuration planning	17
2.1 Concepts and terminology	18
2.1.1 Processor concepts	18
2.2 Virtual I/O versus direct I/O	20
2.2.1 Virtualization and the virtualization engine	20
2.2.2 Virtualization engine (VE) platform technologies	20
2.2.3 Virtualization engine platform services	24
2.2.4 Direct I/O	24
2.2.5 Virtual I/O versus direct I/O	24
2.2.6 Typical implementation scenarios	25
2.2.7 Virtual I/O server	26
2.3 Planning for resource allocation	26
2.3.1 Processor	26
2.3.2 Memory	31
2.3.3 I/O adapter	32
2.3.4 Network configuration	41
2.4 Planning for virtual I/O	42
2.4.1 Considerations for virtual I/O scenarios	42
2.4.2 Hardware requirements for a Linux logical partition	43
2.5 Planning for direct I/O	44
2.5.1 Considerations for direct I/O scenarios	44
2.5.2 Hardware requirements for a Linux logical partition	44
2.6 Planning for hybrid scenarios	48
2.6.1 Considerations for hybrid scenarios	48
2.6.2 Hardware requirements for a Linux logical partition	48
2.7 Planning for VPM	49
2.8 Power control and console options	49
2.8.1 Power control	50
2.8.2 Linux console options	51

Chapter 3. Creating logical partitions	53
3.1 Concepts and terminology	54
3.1.1 Logical partitions	54
3.2 Linux partition with all virtual resources via HMC	55
3.2.1 Creating the logical partition	56
3.3 Linux partition with all virtual resources via Virtual Partition Manager	86
3.3.1 Removing resources from i5/OS	87
3.3.2 Defining new Linux partitions	95
3.3.3 Recovering configuration data	99
3.4 Linux partition with virtual disk and native network via HMC	99
3.5 Linux partition with all physical resources via HMC	101
3.5.1 Creating the logical partition	102
 Chapter 4. Creating the storage environment	 117
4.1 Supporting virtual I/O	118
4.1.1 Network server description (NWSD)	118
4.1.2 Network Storage Space	123
4.2 Native storage	126
4.2.1 Starting Rescue Mode	132
4.2.2 Formatting the disks	134
4.2.3 Building the RAID Set	137
 Chapter 5. Creating the network support	 143
5.1 Networking with physical interfaces	144
5.2 Networking with virtual interfaces	144
5.2.1 Adding virtual Ethernet adapter to i5/OS partition	144
5.2.2 Creating the Ethernet Line Description in i5/OS	151
5.2.3 Enabling datagram forwarding	155
5.2.4 Proxy ARP	155
5.2.5 Network Address Translation	164
5.2.6 Other approaches	176
 Chapter 6. Linux installation	 177
6.1 Console access	178
6.1.1 Console access via the Hardware Management Console	178
6.1.2 Console access via i5/OS	178
6.2 Basic Linux installation steps	186
6.2.1 Novell / SUSE SLES 10	187
6.2.2 Red Hat RHEL 4	209
6.2.3 Installation of Linux on POWER Utilities	232
6.3 Native disk install - installation notes	235
6.4 Updating SLES 9 to SLES 10	239
 Chapter 7. Administrating the Linux environment	 249
7.1 Checking Linux partition resources	250
7.1.1 From HMC	250
7.1.2 From Linux OS	252
7.2 Linux utilities	255
7.2.1 YaST for SUSE	256
7.2.2 Red Hat Setup	256
7.2.3 PuTTY	257
7.2.4 VNC	258
7.2.5 Webmin	258

7.2.6 WinSCP	263
7.3 Backup and recovery	263
7.3.1 Backup strategy for Linux partition	263
7.3.2 i5/OS oriented backup	264
7.3.3 Linux-oriented backup.....	278
7.3.4 i5/OS oriented restore.....	284
7.3.5 Linux-oriented restore.....	285
Chapter 8. Leveraging virtualization for specific requirements	287
8.1 Adding space	288
8.2 Increasing size of storage space.....	288
8.2.1 Optional - Rename Storage Space.....	296
8.3 Adding additional storage spaces	299
8.4 Setting up Logical Volume Manager	303
8.4.1 Creation and allocation of virtual disk	304
8.4.2 Setting up the initial LVM environment	305
Making the volume available to Linux	307
Increasing the size of LVM.....	308
i5/OS steps	308
Preparing the disk for LVM usage.....	308
Appendix A. Virtual CD library.....	311
Virtual CD device creation	312
Virtual CD library creation	313
Loading images into CD library	315
Working with CD library	316
Appendix B. Setting up Linux on iSeries models	319
Freeing up system resources.....	320
Appendix C. Performance considerations	331
Virtual I/O support	332
Multi-path I/O	332
Memory pool	334
Virtual network	335
Frame size	336
Multiple network segments to separate internal/external traffic	336
Proxy ARP versus Network Address Translation	337
Processor considerations.....	338
Virtual processors	338
Uncapped partitions	339
Dedicated processors	340
Summary	340
Appendix D. Troubleshooting	343
Virtual I/O - network server configuration issues	344
Attempt to vary on fails - partition not found.....	344
Attempt to vary on NWSD fails - SCSI server resource definition problem.....	346
Vary on NWS fails - host partition not defined	348
File system error accessing stream file.....	350
Inability to access the Linux console through i5/OS virtual console support.....	350
Virtual console access denied (through i5/OS) - console in use.....	352
Inability to access Linux console (via HMC) - console in use	353
Partition fails to start - resource conflict	353

Inability to open Linux console in WebSM	355
Linux does not shutdown when the network server is varied off.....	355
An RPA communication error occurs when a DLPAR operation is attempted	355
Related publications	357
IBM Redbooks	357
Other publications	357
Online resources	357
How to get IBM Redbooks	359
Help from IBM	359
Index	361

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:

Redbooks (logo) ™
iSeries™
i5/OS®
pSeries®
xSeries®
zSeries®
AIX 5L™
AIX®
AS/400®

Domino®
DB2®
IBM®
Lotus®
OS/400®
PowerPC®
POWER™
POWER Hypervisor™
POWER4™

POWER5™
POWER5+™
Redbooks™
Resource Link™
System i™
System i5™
System p™
Virtualization Engine™
WebSphere®

The following terms are trademarks of other companies:

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

Active Directory, Microsoft, Windows, Win32, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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.

RealVNC and the RealVNC logo are trademarks of RealVNC Ltd.

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

Preface

This IBM® Redbook is about POWER™ Linux® implementation on the IBM System i™ platform. With the brand-new technologies of POWER5™ and IBM Virtualization Engine™, POWER Linux on System i5 is even more powerful, scalable, and flexible. We also have new versions of Linux distributions that are available for POWER Linux on IBM System i platform. They are SUSE Linux Enterprise Server (SLES) 10 and Red Hat Enterprise Linux (RHEL) 4.

This book gets you ready to plan, configure, and install Linux on System i platform. It provides information about setting up network connection between the server i5/OS® partition and the client Linux partition over virtual local area network (LAN) using various methods. It also covers various topics of administration, operations, and tips and techniques.

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.



Yessong Johng is an IBM Certified IT Specialist at the IBM International Technical Support Organization, Rochester Center. He started his IT career at IBM as an S/38 Systems Engineer in 1982 and has been with S/38, AS/400®, and now iSeries™, for 20 years. He writes extensively and develops and teaches IBM classes worldwide in the areas of IT optimization whose topics include Linux, AIX®, and Windows® implementations on System i platform. He is also interested in the e-business area, especially WebSphere® implementations on iSeries.



Erwin Earley is an Advisory Software Engineer for IBM assigned to the System i Technology Center in Rochester, Minnesota. In his role in the Linux Center of Competency, he is responsible for technical enablement of field personnel for Linux on the System i platform. Erwin has written numerous articles about Linux on System i and is a certified professional with the Linux Professional Institute as well as a Red Hat Certified Technician. He can be reached by e-mail at: erwin.earley@us.ibm.com



Rico Franke is an iSeries Technical Support Specialist with IBM in Germany. He has five years of experience in iSeries. His areas of expertise include iSeries Access, EIM / Kerberos, integrated xSeries® solutions (IXA / IXS), and logical partition (LPAR). He provides assistance in problem determination, configuration, and usage for IBM customers. He holds a degree of engineering in Information Technology from the University of Cooperative Education, Mannheim.



Vlatko Kosturjak is an IT Security Specialist with IBM Global Services Security and Privacy Services in IBM Croatia. In his practice, Vlatko specializes in ethical hacking, IT audit, OS/Network security hardening, and security policy development according to the ISO/IEC 17799 standard. Vlatko is also an LPI certified Linux specialist with experience in deploying a wide range of secure Linux solutions. He has extensive experience in Linux on almost every platform (from PDAs to zSeries®), including developing, porting, and contributing to open source software.

This book was built on the basis of the previous version of the book contributed by the following authors:

Fredy Cruz is the Independent Software Vendor (ISV) Coordinator at IBM Colombia. He helps developers in infrastructure migrations over software including WebSphere and DB2® on both Linux on iSeries and OS/400®. He also teaches ISVs and customers about the utilities and technologies related to this task. His major responsibilities are demonstrating new utility implementations on the i5, Linux, AIX and Windows integration on i5, DB2 and WebSphere in i5/OS, AIX and Linux, and Lotus® in the i5/OS environment.

Stacey Johnson is an Advisory IT Specialist with IBM. She has 10 years of experience on the iSeries platform. She started her IBM career 10 years ago in the area of iSeries system support. She remains in that area working in the IT Data Center in Rochester, Minnesota. There she supports and manages numerous iSeries systems running a multitude of applications (HTTP, Apache, WebSphere, Domino®, Linux, OS/400, i5/OS, BRMS) for IBM internal customers. She holds a bachelor's degree in Computer Science from Winona State University. She can be reached by e-mail at: SPFAFF@US.IBM.COM

Vaclav Matousek is an IT Specialist in the Czech Republic. He has over 17 years experience in computing, with eleven years in the AS/400 and iSeries area. He is a member of the FTSS team in IBM CEMA focusing on LPAR, Linux and AIX on iSeries and i5. He is a Certified Technical Specialist on LPAR, Linux and Windows on iSeries. He teaches LPAR and Linux education and workshops in IBM CEMA countries and Solution Assurance on projects in the area of focus. You can reach him by e-mail at: vaclav_matousek@cz.ibm.com

Vaseem Ansari is Assistant General Manager-Infrastructure Services at Melstar Information Technologies Limited, India. Melstar is an IBM Business Partner since 1994. Vaseem joined Melstar in 1994 and has been involved in the areas of technical support services and network integration services. He has 14 years of experience in the design and implementation of multi vendor network solutions. He has extensive experience with Windows and Solaris™ operating systems and specializes in Linux. He is a Red Hat Certified Engineer (RHCE). He is currently focused on information system security.

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

David Boutcher
Pam Morris
Jay Bryant
Ray Anderson
Mike Ranweiler
Monza Lui
Dharmesh Patel
Edith Lueke
John Janosik
Mark Vanderwiel
Richard Johnston

Fant Steele
Will Schmidt
Craig Johnson
IBM Rochester

Bob Foster
Christine Wang
IBM Austin

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 e-mail to:

redbook@us.ibm.com

- ▶ Mail your comments to:

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

Archived

Introduction to Linux on System i platform

This chapter defines the terms you require to become familiar with as you partition your System i platform to support Linux. The terms provided here are used throughout this book. Additionally, this chapter discusses the support for both virtual (that is, resources owned by another operating system) as well as native (resources dedicated to and owned by Linux) resources. It presents the differences between the various generations of iSeries and System i platforms that support Linux. Finally, it provides a roadmap of typical installation or configuration scenarios to help guide your way through the book. It also provides a list of useful Web sites for further exploration of the Linux operating system.

When reading this book, as well as when working with Linux on the System i platform, remember that Linux distributions are not the property of IBM nor are they distributed by IBM. Linux distributions are the licensed property of the Linux distributors (such as Red Hat or Novell for SUSE, for example) and the documentation provided with their distributions should be used in conjunction with this book.

At the time of the writing of this book, both Red Hat and Novell / SUSE (the two Linux distributors that have released distributions for the System i platform) have product evaluation programs in place. You can request evaluation copies of the Linux distributions through the distributor Web sites.

1.1 Concepts and terminology

This section provides concepts and terms you will need to know for the partitioning of your System i platform for the installation of Linux into a logical partition (LPAR).

1.1.1 System i platform

The System i platform is a server based on the Power5 and Power5+ technologies. Along with i5/OS (V5R3 and V5R4), these systems support the AIX 5L™ and Linux operating systems running in LPARs, as well as Windows integration on the Integrated xSeries Server and Adapter.

i5/OS

This is the latest version of OS/400. Two releases of the operating system are supported on the System i platform:

- ▶ V5R3
- ▶ V5R4

Note: Screens and examples provided in this book will be based on the V5R4 version of the operating system.

POWER5 and POWER5+

POWER5 is the 9th generation of 64-bit POWER processor technology from IBM. Although the hardware is based on POWER4™, POWER5 is more than just an improvement in processor or chip design. Enhancements have been made to improve performance, allow more processors to be used in a system, and improve power efficiency.

Some of the new features included with POWER5 systems include simultaneous multi-threading (SMT), virtualization, and dynamic logical partitioning (DLPAR). AIX 5L can now run on POWER5 systems along with versions of Linux compiled (built) for the POWER architecture.

Note: A number of Linux distributions can be implemented on the System i platform including distributions from Red Hat and Novell SUSE. Linux distributions built from code trees maintained by Debian, Mandrake, and CentOS have been known to run on the platform, however, the only two supported distributors are Red Hat and Novell / SUSE.

Linux screens and information provided in this book will be based on Red Hat's RedHat Enterprise Linux version 4 (RHEL 4) and Novell SUSE's SUSE Linux Enterprise Server version 10 (SLES 10).

POWER Hypervisor

The POWER Hypervisor™ is the level of code that is above the POWER5 technology-based hardware. It is shipped with the System i platform and is part of the firmware. It resides in the flash memory of the service processor (SP) and performs the initialization and configuration of the System i hardware. It also provides the virtualization support needed to run up to 254 logical partitions concurrently on the System i platform.

The POWER Hypervisor performs the logical partitioning and is responsible for assigning resources to the partitions and ensuring the partitions remain isolated from each other. If automatic load balancing and uncapped processors are utilized, the POWER Hypervisor is

responsible for assigning the unused processor capacity and making maximum use of the processor.

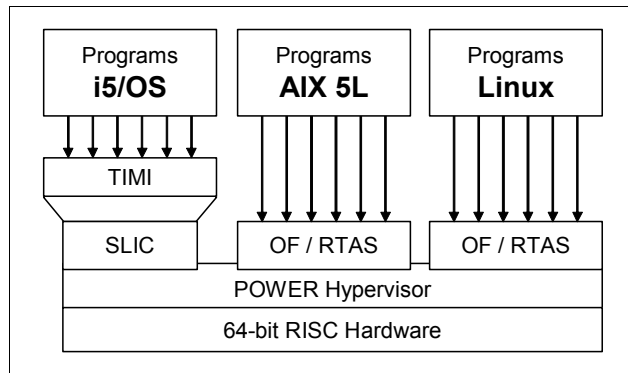


Figure 1-1 Power Hypervisor

The POWER Hypervisor allows for multiple operating systems to run on the System i platform. i5/OS, Linux, AIX 5L are supported. The layer above the POWER Hypervisor is different for each supported operating system.

For i5/OS, the layers above the POWER Hypervisor are still in place. SLIC (System Licensed Internal Code) has been changed and is enabled for interfacing with the POWER Hypervisor.

For Linux and AIX 5L operating systems, the layers above the POWER Hypervisor are similar but their content is characteristic for the operating system. These layers include system firmware and Run-Time Abstraction Services (RTAS).

System firmware

System firmware consists of Low Level Firmware and Open Firmware (OF). Low Level Firmware performs the server unique input/output configurations. Open Firmware contains the boot time drivers, the boot manager, and the device drivers required to initialize the PCI adapters and attached devices.

RTAS - Run-Time Abstraction Services

This consists of the code that supplies platform dependent accesses and can be called from the operating system.

System management services

System Management Services (SMS) is the user interface to Open Firmware. SMS is available in all non-i5/OS or non-OS/400 partitions. The SMS menus are displayed on the partition console.

Service processor

The Service Processor (SP) is a separate processor within the POWER5 system. It interacts with the hardware management console (HMC) via ethernet and with the user via the front panel of the system. It is active as long as the system is plugged in, regardless of the on/off status of the system.

Advanced system management

Advanced System Management (ASM) is a Web interface to the service processor (SP). This can be used to configure systems without HMCs. (Note that either an HMC or Virtual Partition Manager (VPM) is required for any partitioning of the System i platform.) ASM can perform a

number of service functions such as reading service processor error logs, reading vital product data, setting up the service processor and controlling the system power.

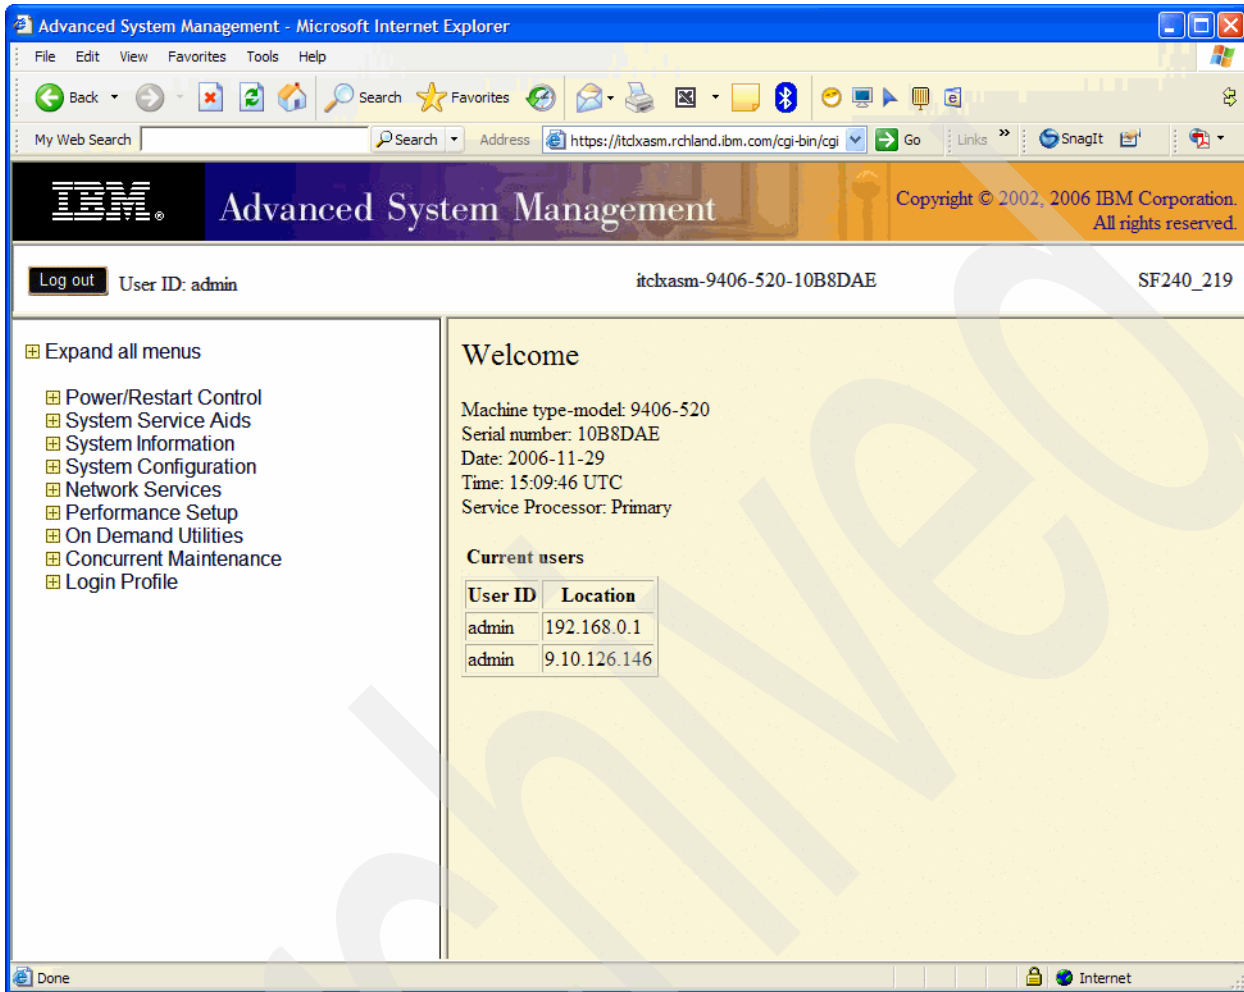


Figure 1-2 Advanced System Management

Automatic processor movement

The System i platform and the POWER Hypervisor can move processor resources based on the workload in the logical partitions. Resources can be removed from non-busy logical partitions and added to busy logical partitions automatically. This capability is often referred to as *uncapped* partitions and will be discussed in greater detail later in this book.

1.1.2 Hardware management console

The hardware management console (HMC) is used to manage the hardware in the System i platform as well as logical partitions and Capacity on Demand. It is a preinstalled Linux-based workstation with an ethernet connection and is available as a standalone desktop system or can be rack-mounted in the System i platform. It is a *closed system*, meaning no other software can be installed on it.

Note: Many of the logical partition functions described in this section are also supported in the Virtual Partition Manager. HMC is discussed exclusively here as it is the likely method that most customers will use to manage the partitioning of the System i platform. Information on Virtual Partition Manager will be presented later in this book.

The HMC is the typical method used for creating and managing logical partitions on the System i platform and is required for using Capacity on Demand as well as performing Dynamic LPAR (DLPAR) functions. The HMC can also be used to activate logical partitions are in a powered off state.

One HMC can support multiple System i platforms. Additionally, a single HMC can manage both System i as well as System p™ platforms.

There is no longer a primary partition on the System i platform. Some of the functions previously supported by the primary partition are now performed through the HMC. The HMC replaces the functions performed in service tools regarding partitions and also replaces the functions in DST for enabling Capacity Upgrade on Demand (CUoD).

The HMC window has two panes as seen in Figure 1-3. A Navigation Area is on the left side and this lists the functions that can be performed from the HMC. The right pane, Contents Area, will change as you make selections on the left. Figure 1-3 below shows the System i platform (sometimes referred to as the *Managed System*) that the HMC manages.

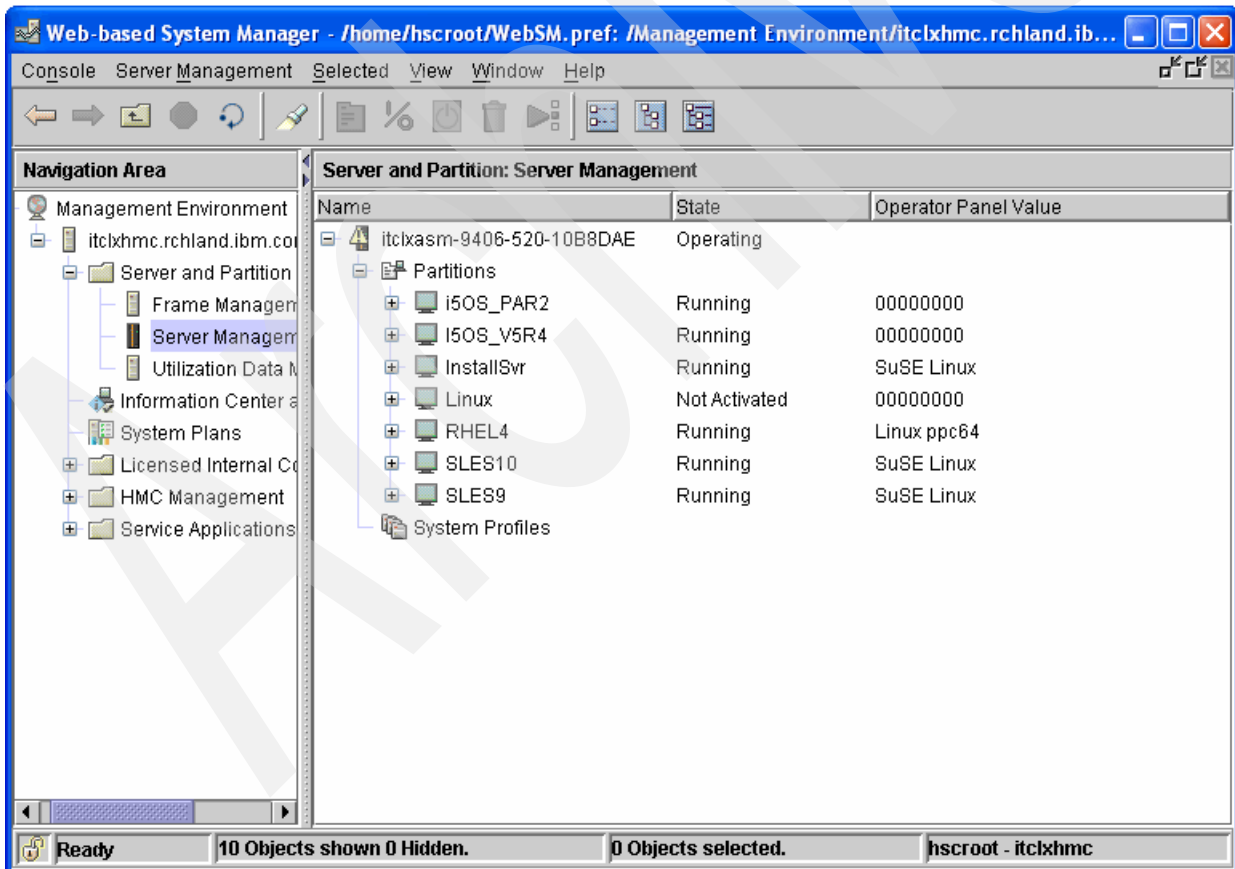


Figure 1-3 HMC Window

HMC scheduler

Memory and processor resource movements can now be done using the HMC scheduler. On the HMC, you go into HMC Management, HMC Configuration, Scheduled Operations, Select a Partition, and then use Dynamic Reconfiguration Option.

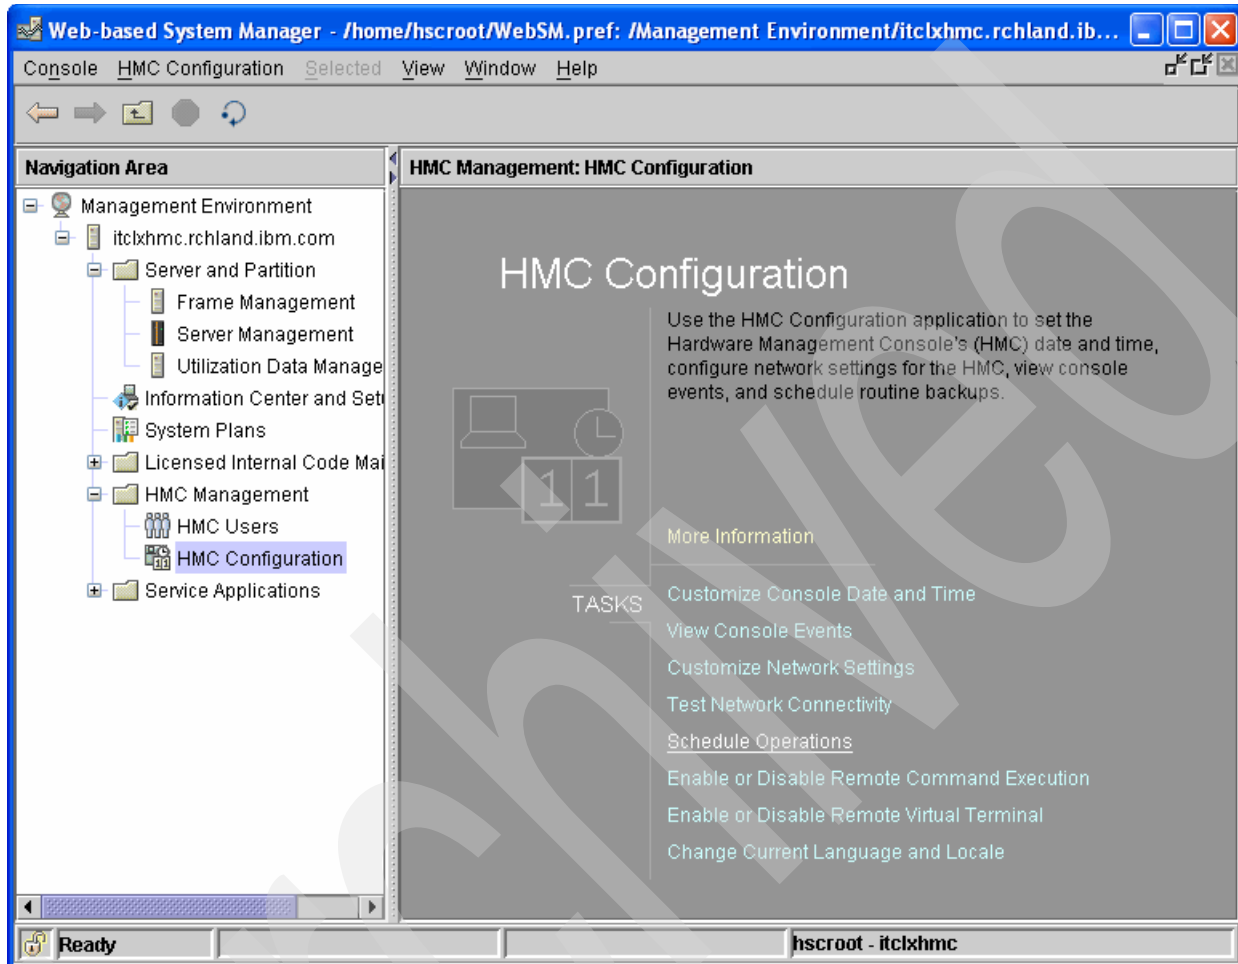


Figure 1-4 HMC Scheduler

Managed system

The managed system is the System i physical hardware that the HMC has been configured to manage.

Managed system properties

From the HMC, you can see where your system resources have been assigned. On your HMC window, select **Server and Partition** → **Server Management**. The left pane of your HMC will now show the managed systems. Below each managed system is the partitions and system profiles.

Right-click the managed system and then select **Properties**. A Managed System 520i5 Property Dialog window is displayed as seen in Figure 1-5. This first window will give you some general information regarding your i5. The tabs at the top will give you more detailed information and show you where your resources have been assigned.




Figure 1-5 Managed system properties - general information

Figure 1-6 shows the processor information for the managed system and the processor usage by each logical partition.

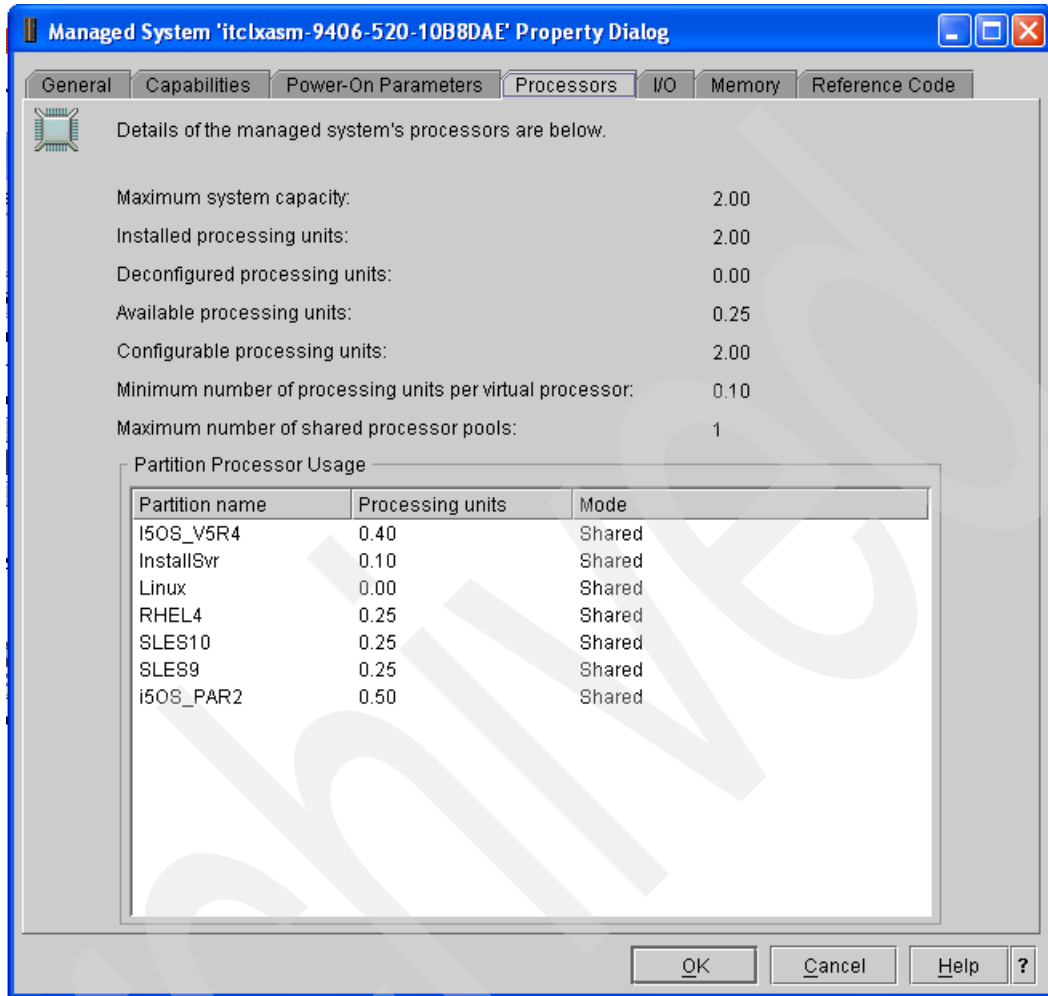


Figure 1-6 Managed system properties - processor information

Figure 1-7 lists the physical I/O resources for the managed system. Details on the I/O adapters and the logical partitions they are assigned to are displayed here.

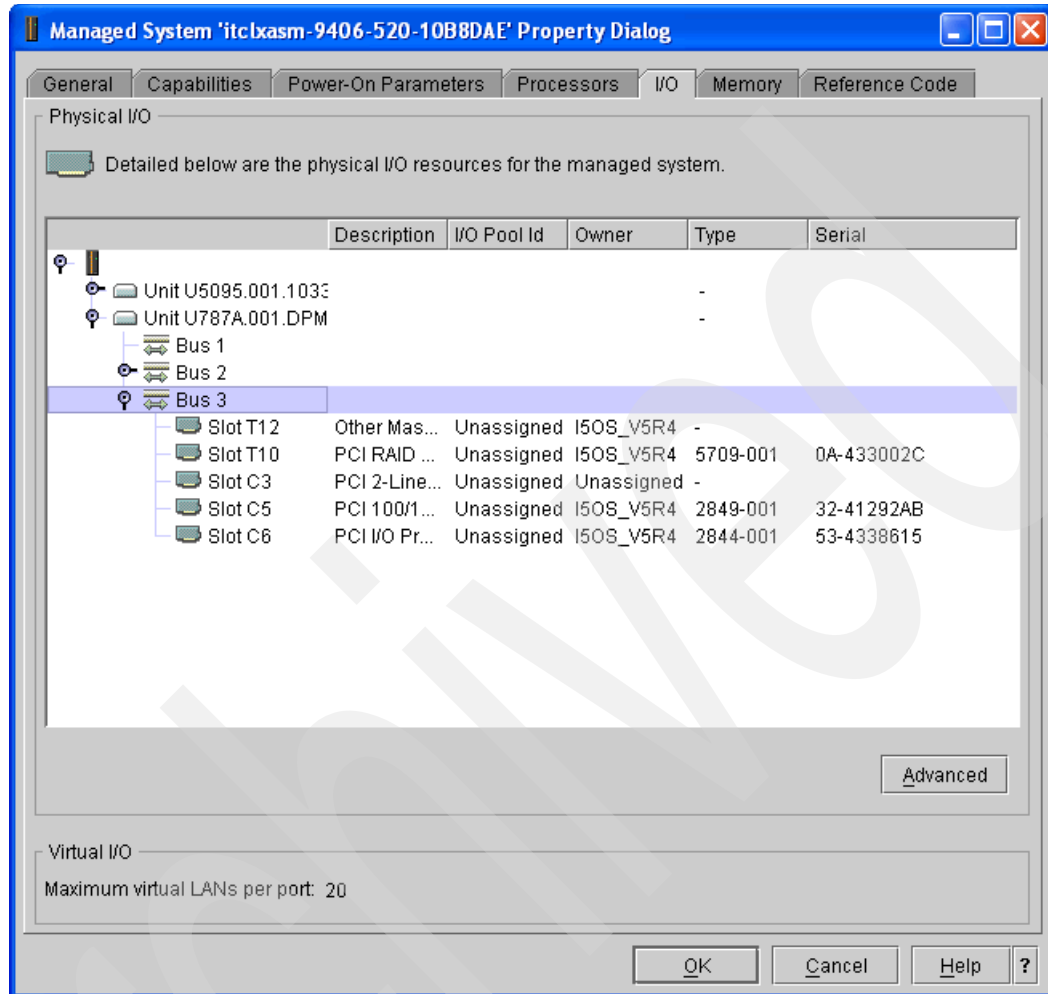


Figure 1-7 Managed system properties - I/O information

Figure 1-8 shows the details of the managed system's memory and the memory usage by logical partition.

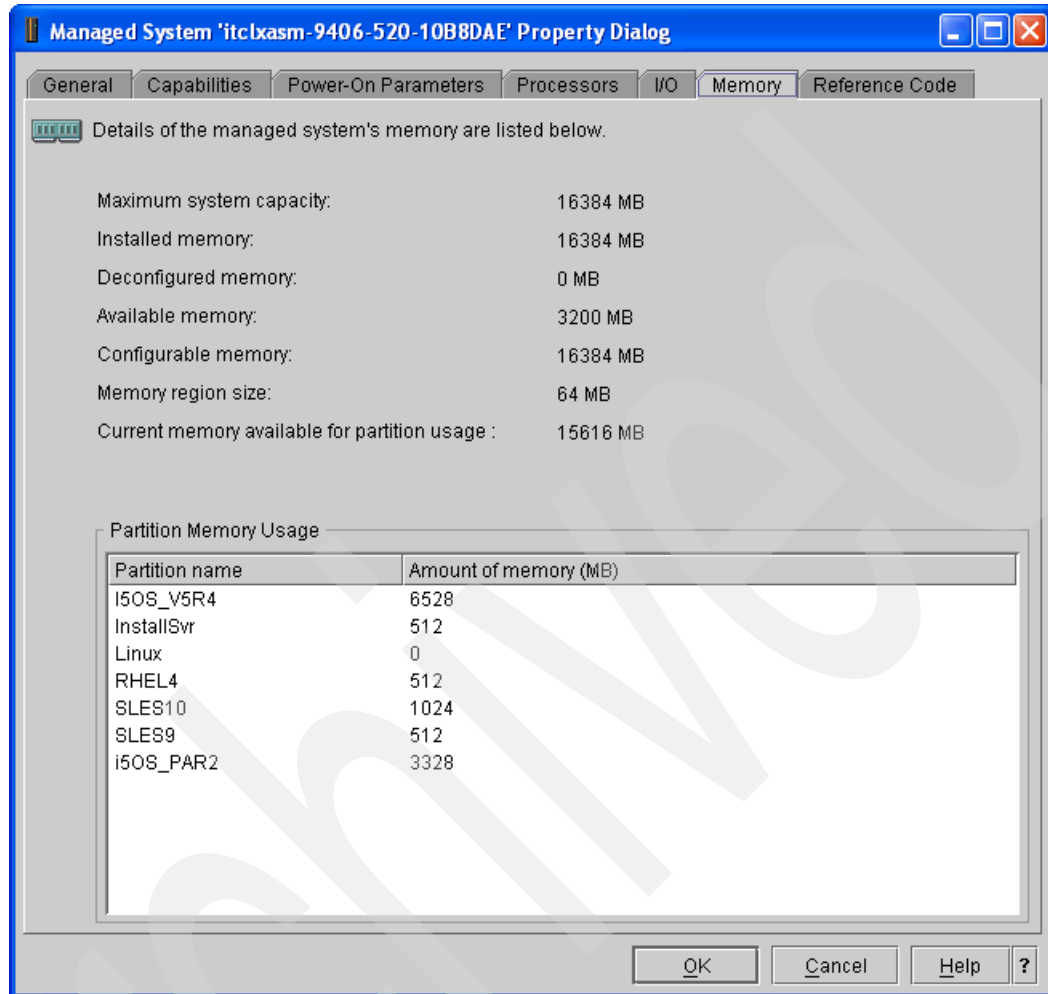


Figure 1-8 Managed system properties - memory information

Web-based System Manager

Web-based System Manager (WebSM) is a java-based application that allows you to perform HMC functions remotely from another PC. This tool is downloaded to the PC from the HMC via the following Web interface:

http://address_of_hmc/remote_client.html.

1.1.3 Virtual Partition Manager (VPM)

The Virtual Partition Manager (VPM) provides the capability to create and manage Linux partitions without the use of the HMC. The Virtual Partition Manager supports the needs of small and medium customers that want to add Linux workloads to their System i platform. The Virtual Partition Manager is included with i5/OS V5R3, as illustrated in Figure 1-9.




Figure 1-9 The Virtual Partition Manager

Note: Virtual Partition Manager can only be used on System i platforms that do not have a Hardware Management Console.

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.

The HMC remains the premier management tool for logical partitioning on the System i platform. 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 1-1.

Table 1-1 Differences between VPM and HMC

	Virtual Partition Manager	Hardware Management Console
Operating systems supported	i5/OS and Linux	i5/OS, Linux, and AIX
Maximum number of partitions	Five (One i5/OS and four 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	Four	4094
Maximum Virtual Disk per Partition	64 TB	64 TB

Note: A System i platform can only be managed by either Virtual Partition Manager of the Hardware Management Console. If you plan to switch from one solution to another (typically from Virtual Partition Manager to the Hardware Management Console), you should reset the service processor, which means deleting all the created partitions.

For more info about VPM, refer to *Virtual Partition Manager A Guide to Planning and Implementation*, REDP-4013.

1.2 Brief introduction to Linux and Linux on System i platform

Linux was first developed in 1991 by Linus Torvalds in Finland. As time has passed, others have added packages to the original kernel and began to distribute them. Hence, the concept of distributions was created. There are numerous distributions of all types now available from companies or individuals.

Linux's functionality, adaptability and robustness have made it an alternative for UNIX® and Microsoft®'s operating systems. It is being adopted world-wide as a server platform and its home and office desktop use is also on the rise.

Linux has an official mascot, Tux the penguin. He was selected by Linus Torvalds to represent the image he associates with the operating system.

1.2.1 Linux on System i platform

Supported distributions of Linux for the System i include SUSE LINUX Enterprise Server (SLES) and Red Hat Enterprise Linux (RHEL).

System i offers Linux integration and consolidation via logical partitioning and resource virtualization. The i5 can help reduce costs by server consolidation and improve investments with new applications.

System i can consolidate several servers with Linux partitions. It can also share resources among i5/OS and Linux partitions. Logical partitioning supports sharing and dynamically moving processor resources, thus maximizing processor utilization.

SUSE LINUX Enterprise Server 10 for POWER (SLES10)

The first distribution of SUSE Linux began in 1993. The company was acquired by Novell in 2004.

SLES10 is a secure, reliable platform for open-source computing supporting a range of hardware platforms and software packages. It also provides open application programming interfaces (APIs) and other development tools to help simplify Linux integration and customization.

SLES10 includes the Linux kernel version 2.6, an improvement over earlier systems in terms of scale, speed and power. Many of the improvements in SLES10 are a direct result of the new kernel.

Red Hat Enterprise Linux AS for POWER Version 4

Red Hat is a well-known Linux and open-source provider. It was founded in 1993. Red Hat Enterprise Linux runs on multiple architectures and is certified by enterprise software and hardware vendors.

Red Hat EL4 is based on the 2.6 kernel.

1.3 Differences between existing Power5-based System i and previous System i models

Table 1-2 shows some of the significant differences between the iSeries and i5 systems.

Table 1-2 Differences at a glance

	Pre-Power5	Power5 / Power5+
Documentation	IBM @server® iSeries Information Center	IBM @server Hardware Information Center and IBM @server iSeries Information Center
Architecture	POWER4	POWER5
Max # partitions	32	254
Max partitions per processor	10	10
Processor movement	<ul style="list-style-type: none"> ▶ Static ▶ Dynamic movement of processor units 	<ul style="list-style-type: none"> ▶ Static ▶ Dynamic movement of processor units and number of processors ▶ Automatic
Memory movement	<ul style="list-style-type: none"> ▶ Static ▶ Dynamic ▶ Assigned in 1 MB increments 	<ul style="list-style-type: none"> ▶ Static ▶ Dynamic ▶ Assigned in 16 MB increments
Max # of virtual ethernets	16	4094
Partition management	Primary partition	HMC, VPM
Types of partitions	<ul style="list-style-type: none"> ▶ Primary ▶ Secondary 	<ul style="list-style-type: none"> ▶ No primary or secondary ▶ i5/OS service partition ▶ Partition profiles ▶ System profiles
LPAR user interface	<ul style="list-style-type: none"> ▶ iSeries Navigator ▶ Dedicated Service Tools (DST) or System Service Tools (SST) ▶ LPAR API 	<ul style="list-style-type: none"> ▶ HMC required for partitioning ▶ LPAR DST/SST not available ▶ HMC remote command
Operating systems	<ul style="list-style-type: none"> ▶ V5R3 OS/400 ▶ Linux 	<ul style="list-style-type: none"> ▶ i5/OS ▶ Linux ▶ AIX 5L
Sharing I/O with other operating systems	Linux	<ul style="list-style-type: none"> ▶ Linux ▶ AIX 5L
Scheduling of I/O resources, CPU and memory	<ul style="list-style-type: none"> ▶ LPAR toolkit ▶ iSeries Navigator 	<ul style="list-style-type: none"> ▶ Memory and CPU movement through HMC ▶ Partial CPU and I/O movement planned for future release

	Pre-Power5	Power5 / Power5+
I/O allocations	<ul style="list-style-type: none"> ▶ Allocated at bus-level or IOP-level ▶ IOP may be switched among partitions dynamically ▶ Bus ownership or bus changes occur dynamically 	<ul style="list-style-type: none"> ▶ IOP may be switched among logical partitions ▶ I/O resources allocated at slot level ▶ Bus ownership cannot be configured. All buses are effectively shared
Uncapped processor partitioning	no	yes
LPAR configuration data	Contained in load source	Contained in service processor and HMC
Console support	<ul style="list-style-type: none"> ▶ iSeries Operations Console - LAN or direct attach ▶ Twinax console 	<ul style="list-style-type: none"> ▶ HMC ▶ iSeries Operations Console - LAN or direct attach
Fixes (PTFS) for Hypervisor	Primary partition	<ul style="list-style-type: none"> ▶ HMC ▶ Service partition
Main storage dumps	<ul style="list-style-type: none"> ▶ System level dump ▶ Logical partition dumps 	<ul style="list-style-type: none"> ▶ Service processor dumps ▶ Platform system dumps ▶ Logical partition dumps

1.3.1 Linux enhancements on Power5 / Power5+

Table 1-3 shows the enhancements for Linux on the i5 hardware.

Table 1-3 Linux on i5

Function	Linux on Pre-Power5	Linux on Power5
Supported Linux version	<ul style="list-style-type: none"> ▶ SUSE Linux Enterprise Server 8 (SLES8) ▶ Red Hat Enterprise Linux 3 	<ul style="list-style-type: none"> ▶ SUSE Linux Enterprise Server 9 for POWER (SLES9) ▶ Red Hat Enterprise Linux AS for POWER Version 3 ▶ Red Hat Enterprise Linux 4
Partition management		<p>Requires the installation from DRM (DLPAR Resource Manager) of service tools within the partition available on the Web at: http://techsupport.services.ibm.com/server/topdiags</p>

Function	Linux on Pre-Power5	Linux on Power5
Processors	<ul style="list-style-type: none"> ▶ Cannot be changed dynamically ▶ Can be shared among multiple partitions 	<ul style="list-style-type: none"> ▶ Can be changed without restarting the logical partition with Linux distributions running the 2.6 kernel and using DRM ▶ Can be shared among multiple partitions ▶ Sharing mode of capped and uncapped ▶ Power-off partition using dedicated processors will have its processors available to shared processing pool
Max size of a single virtual storage space	64 GB	1 TB
Max total virtual storage per Linux partition	2 TBs	64 TBs

More i5 enhancements

This section discusses more i5 enhancements.

iSeries Navigator

iSeries Navigator has been enhanced to support server and storage management for Linux logical partitions. An administration can start and stop Linux logical partitions in the same interlace they use to manage Integrated xSeries Servers. In addition, iSeries Navigator supports the creation of virtual storage spaces that can be linked/unlinked from Linux network server descriptions (NWSDs).

Backup, Recovery, and Media Services (BRMS)

Support has been added to BRMS to shut down hosted logical partitions before a backup is done and to restart the logical partitions after the backup completes.

1.4 Where to go for more information

The following Web sites have useful information and may be helpful to you as you migrate your Linux server to the i5:

- ▶ iSeries Information Center
 - V5R4
 - <http://publib.boulder.ibm.com/infocenter/series/v5r4/index.jsp>
- ▶ iSeries recommended reading
 - *IBM @server i5 and iSeries System Handbook i5/OS Version 5 Release 3 October 2005 -- Draft, GA19-5486*
 - *IBM System i5, @server i5, and iSeries Systems Builder IBM i5/OS Version 5 Release 4 - January 2006, SG24-2155*

- ▶ Linux-related information
 - Linux on iSeries
<http://www.ibm.com/servers/eserver/series/linux>
 - Linux Online
<http://www.linux.org>
 - SUSE
<http://www.suse.com>
 - Novell on Linux
<http://www.novell.com/linux>
 - Red Hat Linux
<http://www.redhat.com>
- ▶ Education and general information
 - Resource Link™
<http://www.ibm.com/servers/resourceLink>
 - Redbooks
<http://www.redbooks.ibm.com>
 - IBM Midrange Servers - i5 and iSeries
<http://www.ibm.com/servers/eserver/series>



Configuration planning

This chapter provides information to help you in planning for the Linux partition setup and installation.

2.1 Concepts and terminology

Here are some concepts and terms that are introduced in this chapter that you need to know as you work with configuration planning for Linux on the System i platform.

2.1.1 Processor concepts

There are varying options for processor assignments for logical partitions such as dedicated whole processors, shared processors, or fractions of processors. The System i platform introduces the concept of capped or uncapped partitions. Results of your capacity planning and your system hardware configuration will help you determine your optimal processor assignments for each of your logical partitions. One logical partition may need dedicated processors while another logical partition can use the shared processing pool.

Dedicated processors

When a system is logically partitioned using dedicated processors, each processor in its entirety is assigned to a partition. In this case, you would only be allowed two partitions at most with a two-processor system, for example. Each partition would then have only one processor assigned to it. Dedicated processors may get better performance than shared processors, but shared processors are more flexible and allow better use of the processors for multiple partitions.

Shared processors (micro-partitioning)

Shared processors allow you to assign fractions of a processor to a logical partition and are held in the shared processing pool. This allows multiple logical partitions to share a processor.

Virtual processors

A virtual processor represents a single physical processor to the operating system running in the logical partition. Selecting the optimal number of virtual processors depends on the workload of the partition. One partition may benefit from greater concurrent operations while another partition may benefit from greater power.

The default setting for virtual processors is based on the number of processing units specified for a logical partition. The number of processing units is rounded up to the next whole number to get the minimum number of virtual processors. If you have 1.5 processing units, two virtual processors will be assigned by default.

You may want to specify a greater number of virtual processors so that you do not limit the processing capacity of an uncapped partition.

Shared processor pool

The shared processor pool is a group of physical processors that provide processing power to multiple logical partitions. With this shared pool, you can assign fractions of processors (or shared processor units) to logical partitions. There is only one shared processor pool for the system.

When shared processors are used, a physical processor may be moved from one logical partition to another logical partition many times a second. This allows for flexible use of the processor and maximizes processing power across the logical partitions. However, there is a performance impact due to the cost of switching from one logical partition to another. In addition, memory caches are reloaded when processors are switched between logical partitions.

Shared processor units

Shared processor units are a unit of measure for shared processing power across one or more virtual processors. One shared processing unit on one virtual processor accomplishes approximately the same workload as one dedicated processor.

Partitions that use shared processors are assigned these shared processor units. The minimum units that can be assigned to a partition is 0.10 processor units per virtual processor. If two virtual processors are assigned to the logical partition, then 0.20 processor units must be assigned as a minimum.

Capped/uncapped logical partitions

A capped partition is limited to using the number of shared processor units that it has been assigned. It is never allowed to exceed that processing capacity. This allows for very predictable performance and should be used when doing any performance benchmarks.

Uncapped partitions are a feature of Power5 and are supported in the System i platform and System p. While an uncapped partition is guaranteed the number of shared processor units it has been allocated (this is sometimes referred to as a partition's *entitled capacity*), it is now allowed to automatically consume any unused processor capacity in the shared pool. An uncapped logical partition can consume unused processor capacity up to its maximum virtual processor setting.

Note: An important aspect of the uncapped setting to understand is that a partition is only granted access to additional or unused resources for the current system timing cycle (which is 10ms). At the start of each timing cycle, a partition reverts back to its *entitled capacity* which is the amount of processor resources allocated to the partition when it was activated (or the processor setting changed through a Dynamic logical partition (LPAR) function).

You can find more detailed information about the performance aspects of uncapped partitions in Appendix C, "Performance considerations" on page 331.

Uncapped partitions are given a relative weight to other uncapped partitions. This weight is a value from 0 to 255 with 255 being the highest weight, the default weight is 128. The System i platform and the POWER hypervisor will move processor resources to other partitions based on the activity in the partitions and the relative weight given to those partitions.

To illustrate this point, consider the following example: You have a multiple partitioned system. Partitions 1, 2, and 3 are uncapped. Partition 1 has two (2.0) processing units assigned to it, using 50% of its allocated processing resource. Partition 3 has one (1.0) processing unit assigned to it, but it is running at 100% utilization and is in need of more processing resource. Since partition 3 is uncapped and has two virtual processors, the unused processor units in Partition 1 can be used in Partition 3, thus increasing its processor capacity and allowing it to complete its workload.

If Partition 2 required additional resources at the same time, unused processing capacity could be distributed to both partitions. The distribution would be determined by the uncapped weight of each of the partitions. If Partition 2 had a weight of 100 and Partition 3 had a weight of 200, Partition 3 would get twice the unused processing capacity as Partition 2.

Remember, as indicated in the note box above, at the start of the 10ms timing cycle, each partition will revert back to its entitled capacity of processor resources.

Note: CUoD processors that have not yet been activated cannot be used by uncapped logical partitions.

2.2 Virtual I/O versus direct I/O

You have basically two options in terms of providing hardware resources to Linux partitions on System i platform:

- ▶ Virtual I/O
- ▶ Direct I/O

2.2.1 Virtualization and the virtualization engine

Virtualization enables resource sharing in an integrated environment on a single System i platform. Processors, network adapters, storage adapters are among the hardware resources that can be shared among logical partitions.

Virtualization is implemented in the firmware and allows the resource to appear similar to a physical adapter attached to the logical partition. This ability to share hardware allows multiple operating systems to run on a single physical system.

The Virtualization Engine (VE) is a combination of software and hardware technologies that enables computer systems to 'clone' themselves. One server can act like ten and multiply the performance capabilities of the single server. Logical partitioning, virtual storage and virtual ethernet are all components of the virtualization engine.

2.2.2 Virtualization engine (VE) platform technologies

These highly-integrated products allow you to streamline the management of heterogeneous partitions or servers. These can be used to simplify your IT infrastructure.

On/Off Capacity on Demand

Capacity on Demand is a fast, nondisruptive method for activating extra processor capacity already built into select System i platform models. These select models ship with a number of processors and memory units activated. There are additional processors and memory units in these select models that is unavailable until activated. These additional processor and memory units can be activated temporarily, permanently or for a free-trial period. Once activated (either temporarily or permanently), processors can be dynamically allocated to logical partitions as workload demands.

Dynamic logical partitioning

Dynamic Logical Partitioning (DLPAR) is designed to pool resources and optimize their use across logical partitions running multiple operating systems and application environments.

To make the best use of system resources, it sometimes is necessary to move resources between logical partitions while the System i platform is running and the logical partitions are active. DLPAR refers to the functions that moves these resources (processors, memory, I/O devices) between logical partitions while they are active.

Simultaneous multi-threading

Simultaneous Multi-Threading (SMT) allows a single processor to have two sets of registers and two threads of execution and swap between them. When one thread stalls, the other thread can take over. This switch occurs quickly since the processor has two sets of registers. This concurrent execution of multiple instruction streams optimizes processor utilization.

The System i platform servers fully automate SMT without requiring any application modifications or tuning. Depending on the workload, SMT may make the system more efficient. Performance increases have been seen while utilizing SMT function.

Due to SMT, Linux will think it has twice the number of processors than it is configured for. For example, if the Linux partition is assigned two processors, four processors will show in the `/proc/cpuinfo`.

Virtual I/O

Virtual I/O enables sharing of physical resources among partitions. These resources include storage, adapters and devices. Virtual I/O allows the creation of new logical partitions without requiring the addition of physical I/O adapters to the system. Multiple logical partitions can share the same physical I/O resources and each logical partition can simultaneously use both virtual and physical I/O devices. With virtual I/O, Linux can leverage the performance, availability and manageability of the System i platform architecture.

Virtual ethernet adapters

Virtual ethernet adapters are created using the HMC. To the operating system, this virtual ethernet adapter appears as a real ethernet adapter.

Virtual ethernet provides the same function as using a 1 Gb ethernet adapter without requiring additional hardware. It can be used by logical partitions to establish high speed connections to other logical partitions on the same iSeries or System i platform. These virtual ethernet segments can be created or removed dynamically and can be restricted for security or traffic requirements.

When a virtual ethernet port is configured, the i5/OS system creates a virtual ethernet communications port, CMNxx, with a resource type of 268C. This resource is used to configure the ethernet line descriptions and set up the virtual local area network (LAN).

Up to 4094 virtual LANs can be created using virtual ethernet adapters. Each LAN is identified by a virtual LAN ID number. To communicate with each other, ethernet adapters must be assigned the same LAN ID number.

Benefits of virtual ethernet

- ▶ Reliability

Unlike the physical ethernet adapter and its link which exists externally to the System i platform, virtual ethernet exists internally to the system. It runs over system memory bus. Thus, as long as the System i platform is up and running, virtual ethernet always exists providing the highest level of reliability.

- ▶ Security

For the same reason of internal existence to the system, the needs of securing externally existing links are gone. Virtual ethernet is the most secured communication link possible.

- ▶ Fast

Virtual ethernet emulates a 1Gb ethernet connection and provides a fast and convenient communication method between partitions.

- ▶ Economical

Logical partitions can be added and can communicate with external LAN with no extra networking hardware required.

- ▶ Flexible

Distinctive connections can be configured to enable selective communications paths between logical partitions. Logical partitions implement both virtual ethernet and a physical LAN connection.

- ▶ Versatile

The i5/OS and Linux can communicate with each other via TCP/IP over the virtual ethernet ports.

- ▶ Reduced congestion

Using virtual ethernet for communication between logical partitions reduces traffic on the external LAN. This will help prevent degradation of service for other external LAN users.

Connecting virtual ethernet and external LAN

There are three methods for connecting the virtual ethernet and external LAN. Refer to the IBM Information Center for more details on the implementation of these methods:

- ▶ Proxy ARP

This is a built-in function of TCP/IP and uses transparent subnetting to associate a logical partition's virtual interface with an external interface. This method is recommended if you have the necessary IP addresses available.

- ▶ Network address translation (NAT)

This method uses i5/OS packet filtering to route traffic between a logical partition and the external LAN.

- ▶ TCP/IP routing

This is used to route traffic to the virtual ethernet LANs in the same way you would define routing to any LAN. If using this method, updating routing information throughout the network will be required.

Trunk adapter

Each virtual LAN can have only one trunk adapter. This trunk adapter links the LAN to the external network. When data transmission on the virtual LAN is addressed to an adapter that does not exist on that virtual LAN, the transmission is sent to the trunk adapter which forwards it to the external network. Using the trunk adapter along with Layer 2 Bridging actually provides the easiest and most flexible method for connecting virtual local area networks (VLANs) to the external network. For this to work however, the operating system of the logical partition must be configured to forward information to the external network. Note that not all operating systems support this method of connecting virtual LANs to external networks. Consult the documentation for the operating system to determine if it supports trunk adapters.

IEEE 802.1 1Q compatible adapter

IEEE 802.1 1Q is a network standard. This is an option as you create your virtual ethernet LAN and adapters. It allows you to add your virtual ethernet adapter to multiple virtual LANs per the standard. Note that not all operating systems support this network standard. If this is not supported by your operating system, then your virtual ethernet can only connect to one port virtual LAN ID. Consult the documentation for the operating system to determine if it supports the IEEE 802.1 1Q network standard.

Virtual SCSI adapters

Virtual Small Computer System Interface (SCSI) adapters are created using the HMC and provide access to disk for the Linux server. To the operating system, they appear as any other SCSI adapter.

A virtual SCSI client adapter connects to a virtual SCSI server adapter in an i5/OS logical partition. This virtual SCSI server adapter is then associated with the network server description (NWSD) and NWSSTG objects in the i5/OS logical partition.

NWSD is an i5/OS configuration object. If using virtual disk, this NWSD is used to IPL the logical partition, boot Linux, and must be active for the server to function. In an environment using native disk (that is, direct I/O), the NWSD is needed to load the kernel for the initial installation of Linux.

NWSSTG is network server storage space. The Create Network Server Storage Space (CRTNWSSTG) command creates a storage space to be used by the network server, the network server being the Linux logical partition. The Linux installation program reformats this disk for use by the Linux server.

Virtual serial adapters

Virtual serial I/O provides the console function for the Linux server.

Virtual CD

A virtual CD may be used for the installation of Linux. By default, a Linux logical partition can see all the CD drives on the host logical partition.

Virtual tape

A virtual tape provides access to the i5/OS tape drives for the Linux logical partition. By default, a Linux logical partition can see all the tape drives attached to the host logical partition.

2.2.3 Virtualization engine platform services

With support from the POWER Hypervisor, these platform services compliment the platform technologies to help simplify your IT infrastructure.

IBM Director Multiplatform

IBM Director Multiplatform provides a common Web-based console for systems management and monitoring of multiple operating environments. Operating systems include i5/OS, AIX 5L, Linux and Windows. Director Multiplatform ships with all copies of i5/OS.

IBM Grid Toolbox V3 for Multiplatforms

IBM Grid Toolbox V3 for Multiplatforms helps build more secure, robust infrastructures using open standards by connecting systems and resources through a single management interface.

Systems provisioning

Systems provisioning features enable businesses to deploy and re-purpose virtualized IT resources to meet changing business needs. This is designed to maximize resource utilization and reduce data center complexity.

IBM Enterprise Workload Manager (EWLM)

EWLM is designed to monitor computing resources, response times and transaction flows to help businesses improve service level management and capacity planning. It gives you a graphical view of the performance of your enterprise applications. Self-managing capabilities can help isolate problems and optimize workload management.

2.2.4 Direct I/O

Linux on System i platform works directly with almost all I/O devices. Linux communicates directly with System i platform hardware, such as disk units, ethernet adapters, tapes, and so on. In such a scenario, I/O resources dedicated to Linux partition are owned by and are under the control of Linux. It does not depend on i5/OS for any of its resources. Consequently, i5/OS cannot use these resources since they are allocated to the Linux server. For directly attached hardware, all failure and diagnostic messages are displayed within the Linux server.

2.2.5 Virtual I/O versus direct I/O

Linux on System i platform supports virtual I/O and direct I/O (native I/O). These implementations of Linux are often referred to as hosted and non-hosted, respectively.

Virtual I/O: hosted Linux server

With virtual I/O, the resources are owned and managed by the i5/OS logical partition that is hosting the Linux logical partition. The i5/OS shares these resources (disk, tape, CD-ROM, and so on) with the hosted Linux server. The i5/OS provides the direct access storage device (DASD) protection and some backup/restore facilities for the Linux environment. Virtual ethernet provides 1Gb communications paths between logical partitions without requiring additional hardware resources.

Linux is started from the hosting i5/OS logical partition by varying in the network server description (NWSD). This Linux server can only be active when the i5/OS hosting logical partition is active. If the i5/OS hosting logical partition is in restricted state, then the NWSDs are in a varied off state.

Direct I/O: non-hosted Linux server

With direct I/O, the resources are owned by and are under the control of Linux. It does not depend on i5/OS for any of its resources, consequently, i5/OS cannot use these resources since they are allocated to the Linux server. For directly attached hardware, all failure and diagnostic messages are displayed within the Linux server.

Usually, direct I/O is used when performance is critical, because direct I/O is often faster to use than virtual I/O.

Also, with direct I/O, a Linux server can be active even when the i5/OS logical partition is not active, which is not the case in a virtual I/O scenario.

Comparison of virtual I/O and direct I/O

Therefore, when should you use virtual I/O or direct I/O? The answer depends on user requirements and available resources. To help you decide, refer to Table 2-1.

Linux can take advantage of both virtual and direct I/O resources at the same time. Also, it is possible to switch from virtual I/O to direct I/O or vice versa once we realize that other solution will fit better.

Table 2-1 Direct I/O versus Virtual I/O versus Hybrid I/O

	Requirements	Advantages	Disadvantages
Direct I/O	One LAN adapter 0.1 CPU 256 MB RAM One DASD with a minimum size of 512 MB Dedicated CD-ROM	The performance is not dependant on i5/OS, the downtime of i5/OS is <i>not</i> the downtime of the hosted partition, and no i5/OS resources are used.	You need more hardware, and there is administrative overhead (backup and hardware maintenance).
Virtual I/O	0.1 CPU 256 MB RAM One VLAN 512 MB Free on i5/OS Virtual CD-ROM	There is no need for physical hardware or sharing resources, and the VLAN is secure, fast, and can reduce congestion on the external network. There is a single backup procedure, and it is flexible.	The performance is dependant on i5/OS, you need some i5/OS resources, and the downtime of i5/OS is the downtime of the hosted partition as well.
Hybrid I/O	0.1 CPU 256 MB RAM One VLAN or One LAN adapter One DASD with minimum size of 512 MB or 512 MB Free on i5/OS Virtual or dedicated CD-ROM	You combine direct I/O and virtual I/O for the best cost/benefit ratio.	Complex

2.2.6 Typical implementation scenarios

As Linux can take advantage of virtual and direct I/O resources at the same time, some common configurations are hybrid, for example, combining direct and virtual I/O at the same time.

A very popular scenario is what we call a typical hybrid scenario: You have all the I/O devices served by the server i5/OS partition (virtual I/O), except for the network connections (direct I/O for network adapters). You can easily set up a firewall or even a demilitarized zone (DMZ) scenario that way.

2.2.7 Virtual I/O server

The Virtual I/O server is an appliance partition that can be added to a System i or System p to provide virtualized I/O resources. When the Virtual I/O server feature code has been installed on the Managed System an AIX or Linux-based Virtual I/O server partition can be installed on the system. Virtual I/O server can be used to provide virtualized I/O resources (that is, disk) as well as bridging of network traffic between physical and virtual networks.

Note: The Virtual I/O server is an appliance partition that is not intended to run applications or for general use.

Typically, the Virtual I/O server feature code will not be enabled on System i as equivalent functions are integrated into i5/OS itself.

2.3 Planning for resource allocation

Planning for the Linux implementation on the System i platform consists of determining resource allocation in the following areas:

- ▶ Processor
- ▶ Memory
- ▶ I/O Adapter
- ▶ Network

This section discusses each of these items, presents steps for determining current resource allocation on your System i platform and discusses considerations for allocating resources for the Logical Partition that will contain the resources for the Linux instances.

Note: The example shown in the following section assumes a system that already has an i5/OS and Linux partition configured. The following sections walk through collecting the resource definitions to plan for a third logical partition.

2.3.1 Processor

Allocation of processor to the Logical Partition for Linux requires the following considerations:

- ▶ Whole or partial processor allocation
- ▶ Number of virtual processors
- ▶ Capped or uncapped setting

Whole or partial processor allocation

The System i platform supports the allocation of processor resources to a logical partition as either whole/dedicated processors or as partial or shared processor units. A dedicated processor is a processor that is allocated in totality to a single logical partition and used exclusively by that partition. Partial processor allows for the sharing of a processor across a number of logical partitions (up to 10 partitions on a single processor). Additionally, the use of shared processor units allows for the automatic movement of processor units between partitions as determined by the Hypervisor.

The allocation of dedicated processors is the most efficient use of a processor on the System i platform as there is no processor migration or task switching events that need to be managed by the Hypervisor. The use of shared processors, on the other hand, allows for a fuller exploitation of the overall resources of the system as many workloads that can be implemented on the System i platform require less than a full processor allocation.

Note: Most Linux implementations on the System i platform make use of the shared processor feature.

One step of planning for the Linux instance will be a determination of the current processor allocation on the system. For an HMC managed system, the current resource allocation can be reviewed through the properties of the Managed System.

Note: As you review the system resource settings for the overall system and the i5/OS partition it is helpful to complete an LPAR planning sheet similar to the following:

Hosting i5/OS Partition

Name: _____ Partition ID: _____
 Profile Name: _____
 Memory: Min: _____ Desired: _____ Max: _____
 Processor: Min: _____ Desired: _____ Max: _____
 Uncapped: Weight: _____
 Virtual Proc: Min: _____ Desired: _____ Max: _____

Virtual Slots:

Slot	Adapter	Connection/LAN #
0	Server Serial	
1	Server Serial	
2	_____	_____
3	_____	_____
4	_____	_____
5	_____	_____
6	_____	_____
7	_____	_____
8	_____	_____
9	_____	_____
10	_____	_____

For whole or partial processor allocation, perform the following steps:

1. From the HMC, select **Management Environment** → **Server and Partition** → **Server Management** in the navigation pane. Right-click the Managed System in the contents pane and select **Properties** (as shown in Figure 2-1).

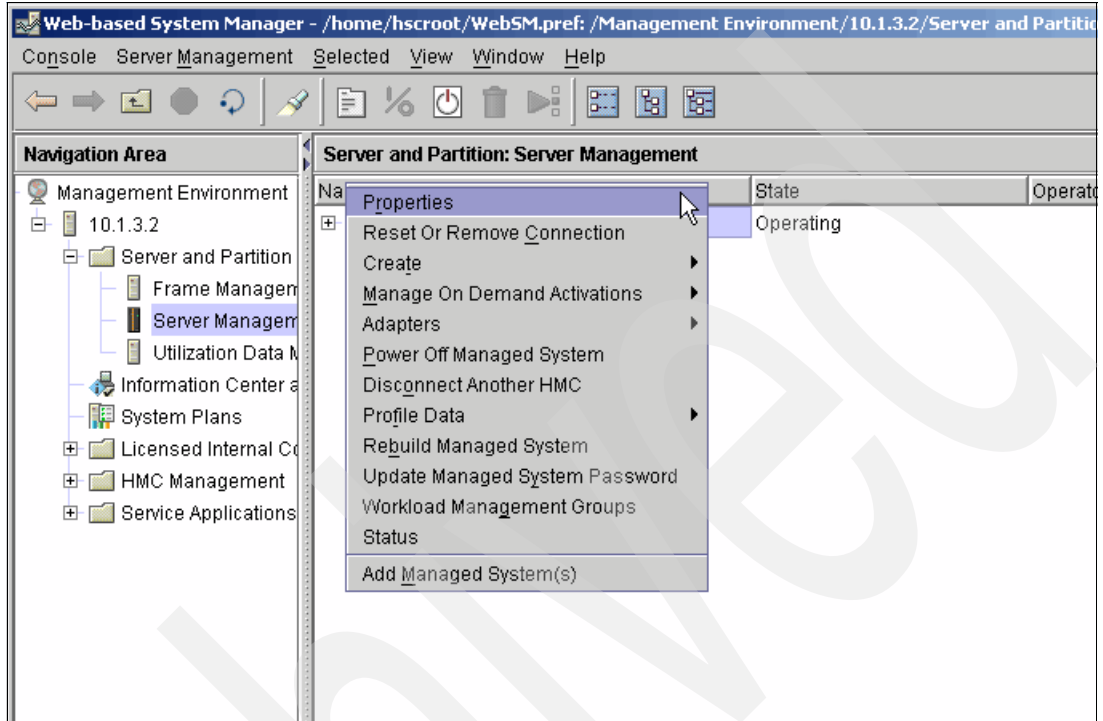


Figure 2-1 Managed System Properties Selection

2. The Managed System Property dialog box is displayed as shown in Figure 2-2. To review the current processor allocation, click the **Processors** tab.

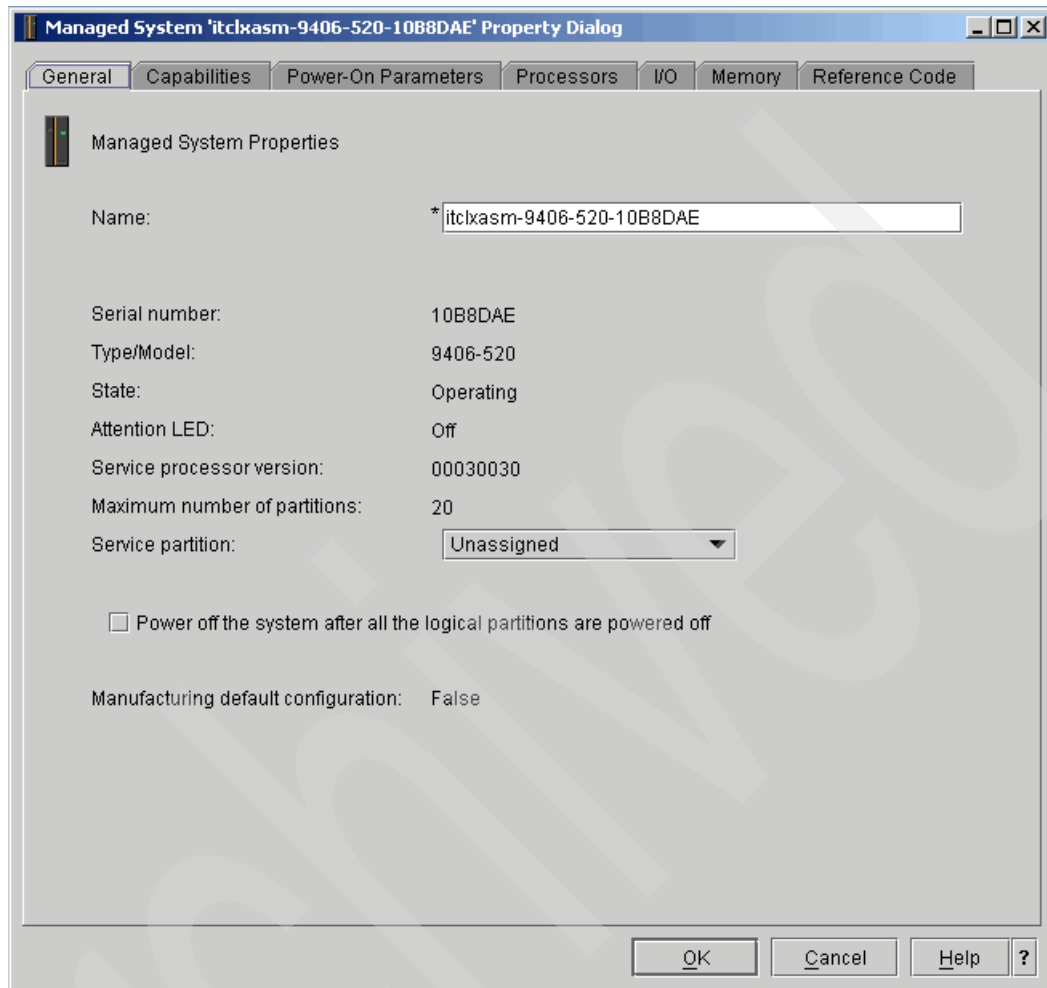


Figure 2-2 Managed System Property dialog box

- As shown in Figure 2-3, the Property dialog box indicates the total processor allocation across all of the partitions that are currently active on the Managed System.

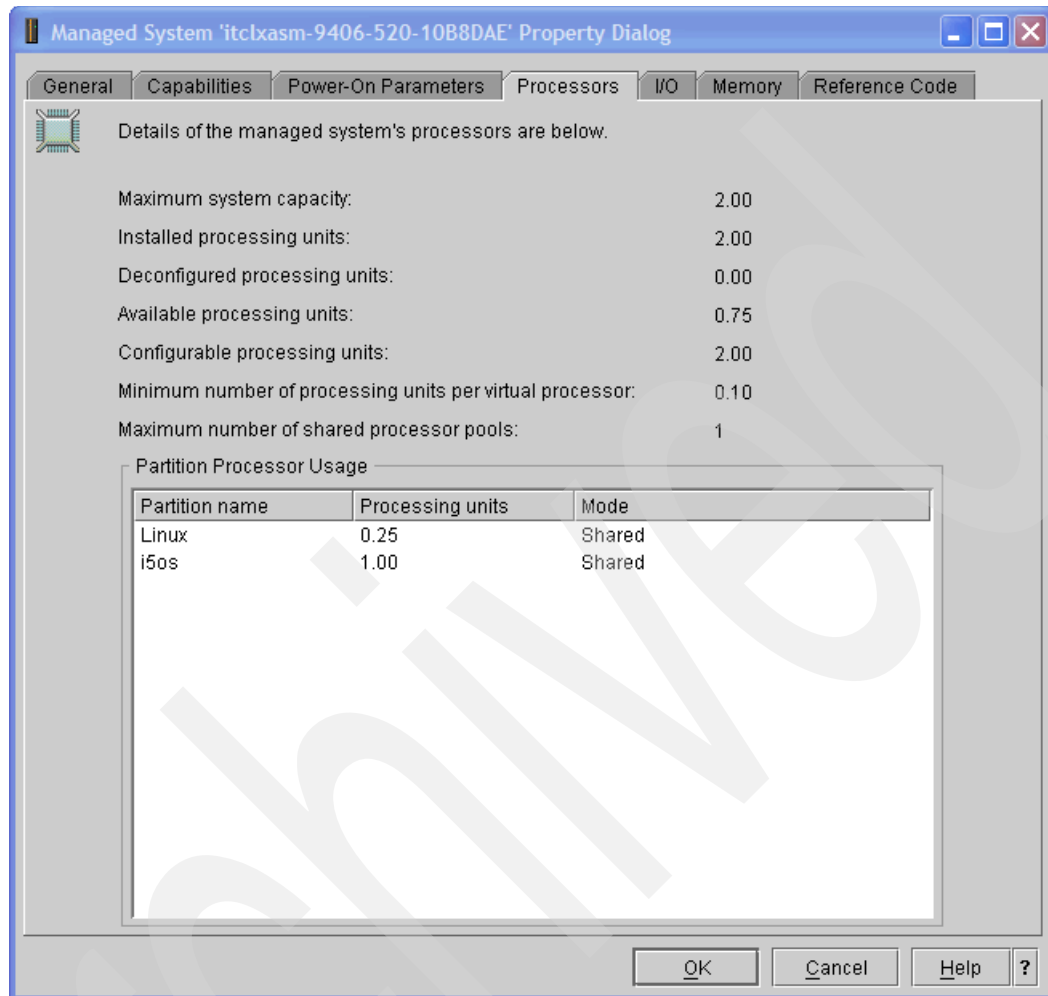


Figure 2-3 Managed System Properties - Processors

As Figure 2-3 shows, there are two processors on the Managed System, with 1.25 processor units currently allocated across two logical partitions leaving 0.75 processor units that can be allocated to additional logical partitions.

Note: On VPM managed systems all process resources will initially be allocated to the single i5/OS partition on the system and the SST screens of VPM will be used to free up processor resources and allocate them to the Linux partition.

Number of virtual processors

Once the amount of processor resource has been determined, the number of virtual processors to spread the workload across needs to be determined. A virtual processor can be thought of as a manifestation of a processor and is represented to the operating system as a processor thread. The number of virtual processors to allocate to a logical partition can be affected by a number of factors including:

- ▶ Type of workload
- ▶ Amount of processor allocation
- ▶ Number of physical processors in the system

There are certain workloads, such as database, that can benefit from a large number of processor threads, however, most workloads that are implemented on the System i platform do not require a large number of processor threads. Each processor thread that will be allocated to a logical partition requires at least 1/10th of a processor unit allocated to the partition. Additionally, no more than a full processor (that is, 1.00 processor units) can be allocated to a single processor thread. Put another way, if 4.20 processor units are allocated to a partition then the minimum number of virtual processors that can be allocated is 5 while the maximum number of virtual processors that can be allocated is 42.

As a general rule for Linux partitions the number of virtual processors allocated will be the least amount required by the allocation of processor units to the platform.

Note: On the System i platform, for each virtual processor allocated to the Logical partition, the Linux operating system will actually see two processors. This is due to the SMP support on the System i platform and ensures that the Linux operating system will benefit from multi-threading even across a single virtual processor.

Capped/Uncapped setting

The system i platform provides the ability for the system to balance the allocation of processor resources across the system based upon the active workloads. This is referred to as *uncapped* partitions.

Note: Most Linux partitions on the System i platform are configured as uncapped partitions.

2.3.2 Memory

Memory is allocated to the logical partitions on the System i platform from the overall memory installed in the system. The amount of memory to allocate to a Linux partition is directly dependent on the workload that will be implemented in the Linux partition as well as the type of I/O (that is, virtual or native) that will be allocated to the Linux partition.

For each partition on the Managed System, the hypervisor will set aside memory resources to manage the memory addressing for the partition. This memory is referred to as the Hardware Paging Table (HPT). The size of the HPT is based on the maximum memory definition for the partition and provides a set of offsets from the partitions memory address to the physical memory of the system. The size of the HPT can be calculated by taking the maximum memory definition for the partition, dividing that figure by 64 and then rounding to the next power of 2.

When a partition leverages I/O resources hosted by another partition (that is, an i5/OS partition hosting resources for a Linux partition), memory resources are required in the hosting partition to process the I/O requests from the guest partition. In other words, for each guest partition hosted by i5/OS, consideration for the amount of memory allocated to the i5/OS partition has to be made. The amount of additional memory required in the hosting partition will vary based on the I/O footprint of the guest partitions.

The current memory allocation on the Managed System can be reviewed through the properties of the Managed System.

1. From the Managed System Property dialog box (refer Figure 2-2, “Managed System Property dialog box” on page 29), select the **Memory** tab to display the memory allocation on the Managed system.

2. Available memory (as shown in Figure 2-4) reflects the overall memory installed in the Managed System less the amount of memory allocated to each active partition as well as memory set aside for the Hardware Paging Tables and overhead for Hypervisor overhead.

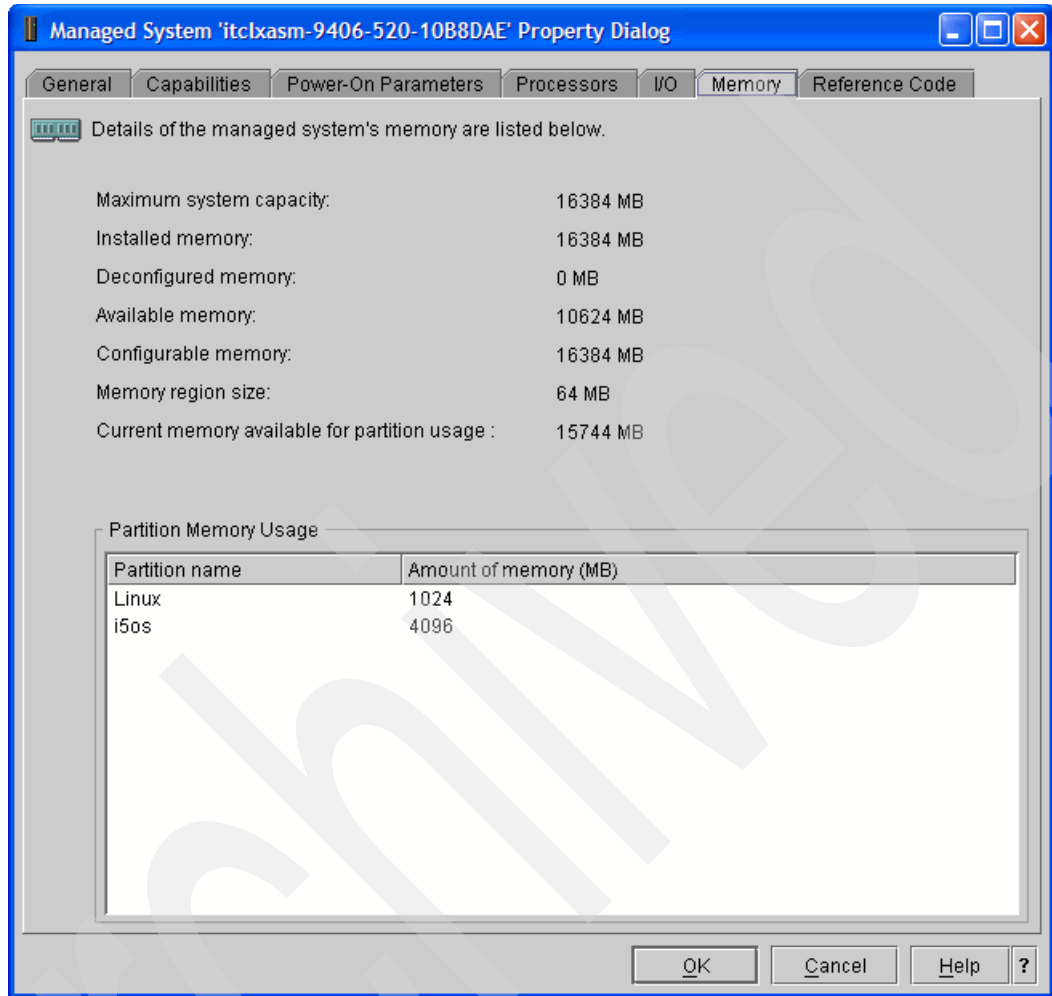


Figure 2-4 Managed System Properties - Memory Allocation

2.3.3 I/O adapter

Each logical partition on the System i platform can have both virtual and native I/O adapters allocated to it. For Linux partitions, I/O adapters are typically used for accessing storage devices (either virtual or native).

Virtual I/O adapter

Virtual I/O is storage devices (direct access storage device (DASD), CD/DVD, tape, and so on) that are physically *owned* by one partition and accessed by the operating system running in another partition. Virtual I/O is the predominate method of storage allocation for a Linux partition on the System i platform. For planning purposes, access to virtual I/O resources requires that a virtual Small Computer System Interface (SCSI) pairing be established between the Linux partition and the i5/OS partition that will *host* the I/O resources. Therefore, it is important that the current virtual I/O configuration of the partition that will host the I/O resources be reviewed and documented.

1. From the HMC, right-click the partition profile of the i5/OS partition that will host the I/O resources for Linux (as shown in Figure 2-5) and then select **Properties**.

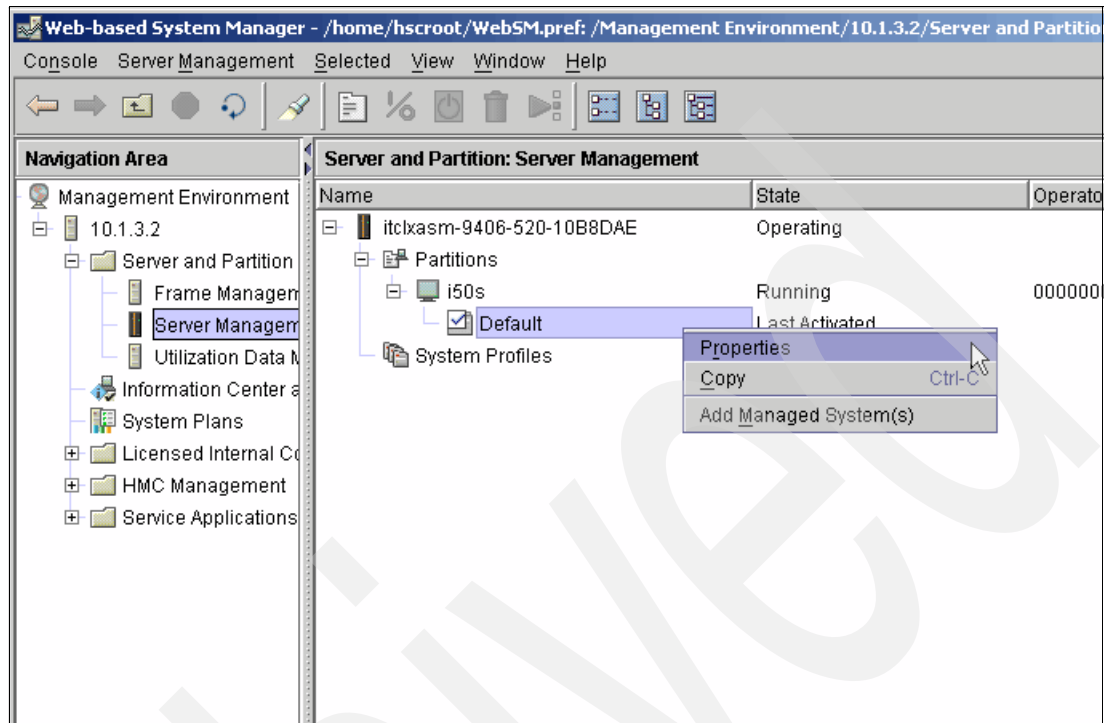


Figure 2-5 Hosting i5/OS Partition - Properties Selection

2. Once the Logical Partition Profile Properties dialog box is displayed (as shown in Figure 2-6), select the **Virtual I/O Adapters** tab.

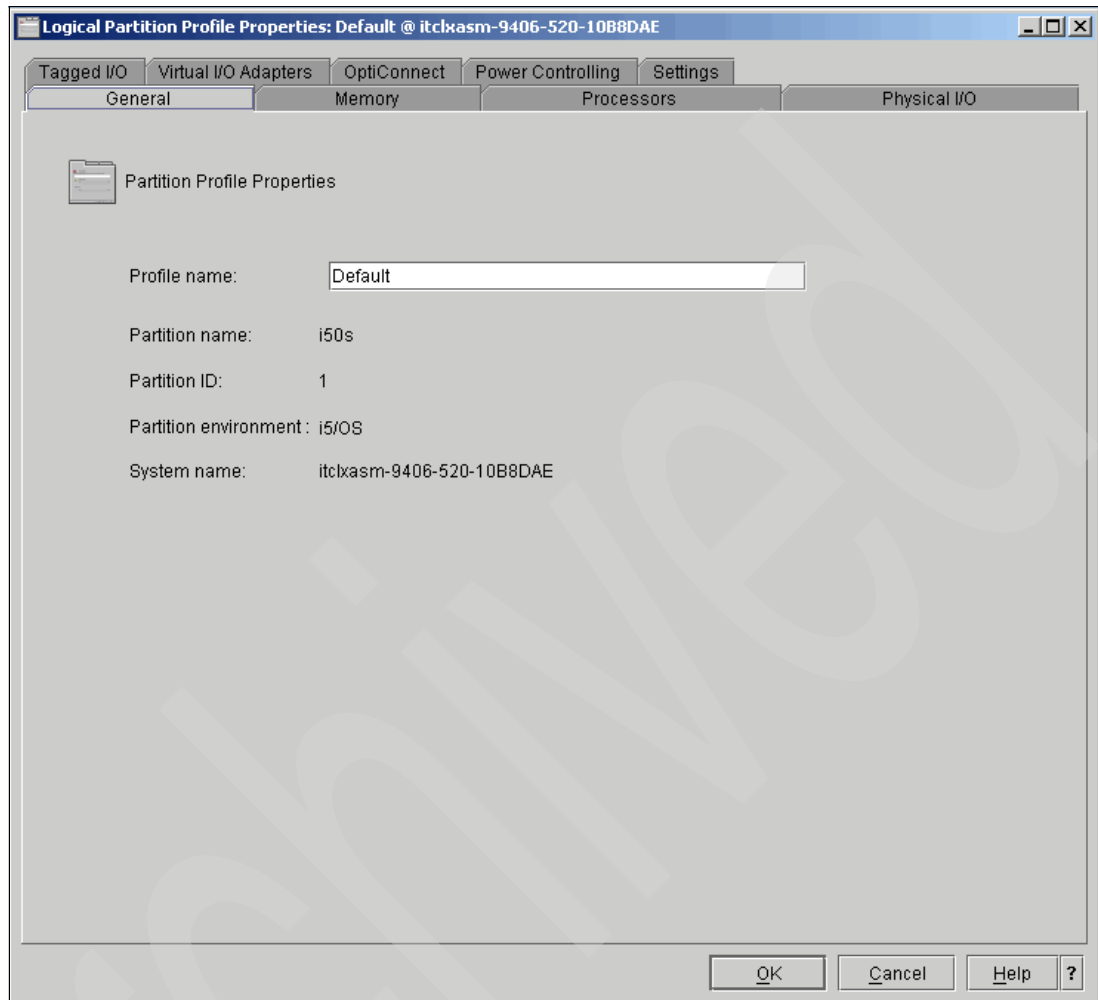


Figure 2-6 Logical Partition Properties

3. Click the SCSI tab for detailed information for each adapter type as shown in Figure 2-7.

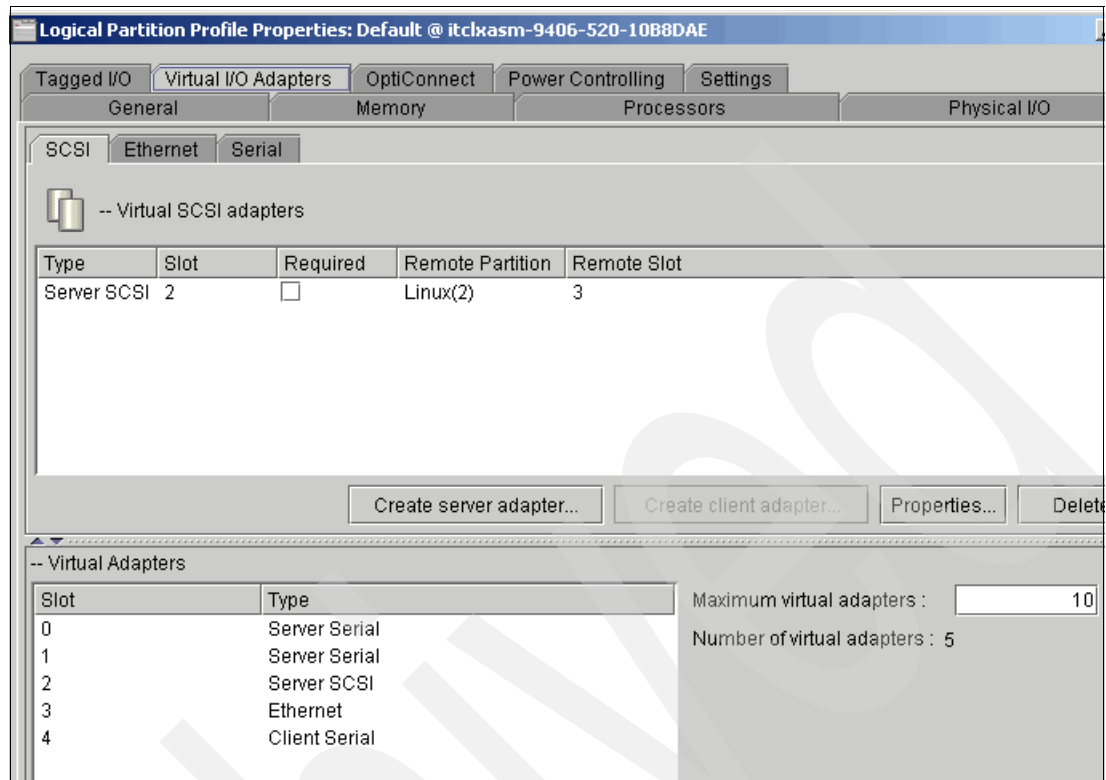


Figure 2-7 Logical Partition Profile Properties - Virtual I/O (SCSI)

Alternatively, click the Ethernet tab as shown in Figure 2-8.

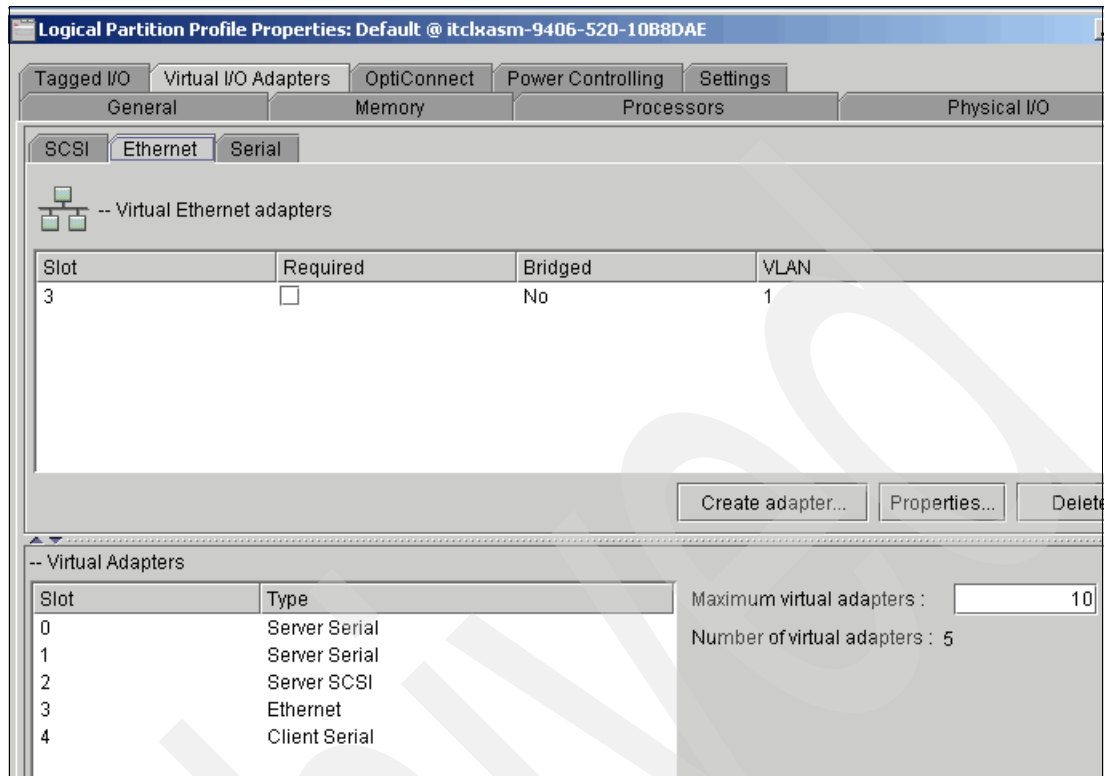


Figure 2-8 Logical Partition Profile Properties - Virtual I/O (Ethernet)

You can also click the Serial tab as shown in Figure 2-9.

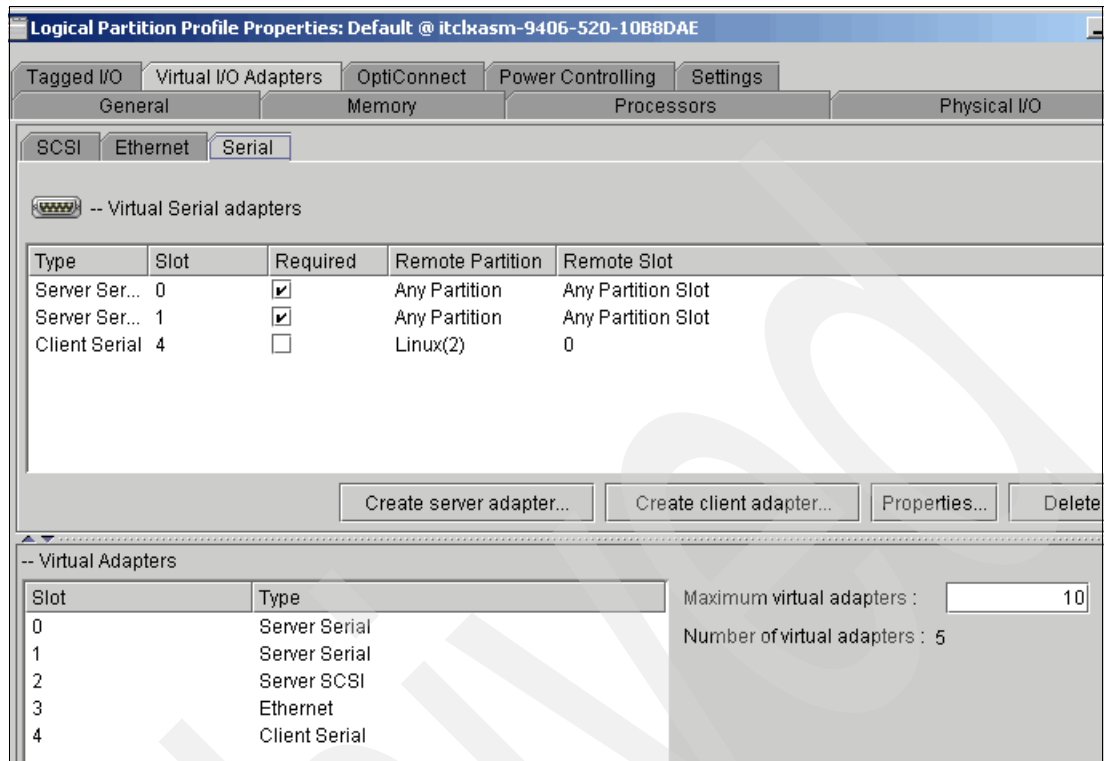


Figure 2-9 Logical Partition Profile Properties - Virtual I/O (Serial)

- The Virtual I/O display shows all of the virtual I/O adapters configured for the partition (Figure 2-10). The configuration of the virtual I/O adapters should be documented to help in the configuration of the guest partition for Linux.

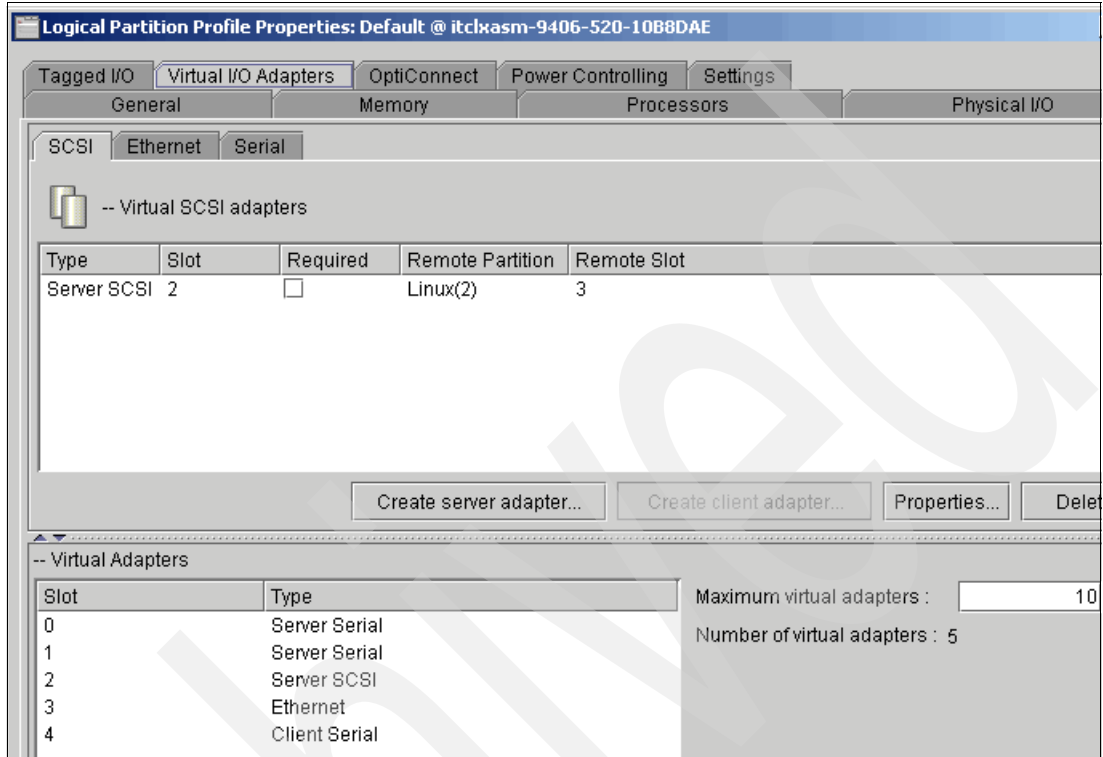


Figure 2-10 Logical Partition Profile Properties - Virtual I/O (SCSI)

One way to document this configuration is to build a diagram of the virtual I/O adapters. As an example, Figure 2-11 represents the Virtual I/O adapter configuration of the i5/OS partition.

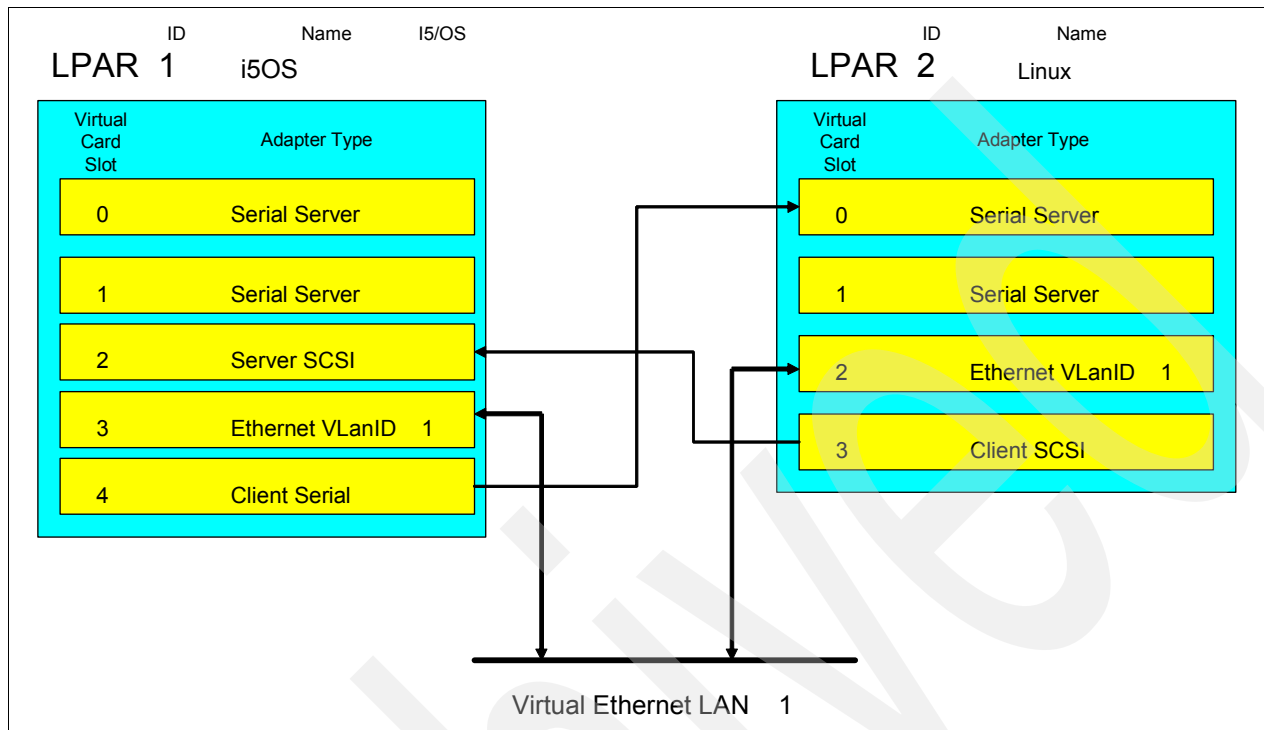


Figure 2-11 Example Virtual I/O configuration

The lines connecting the client and server adapters shown the relationship between the adapters. To complete this diagram, the properties of the two guest partitions (InstallSvr and SLES10) would need to be reviewed to determine what additional virtual I/O adapters are configured. The information presented in this configuration diagram could now be used when creating the Linux guest partition to determine the placement of the Virtual Server SCSI adapter.

Direct I/O adapter

Native I/O is storage devices (DASD, CD/DVD, tape and so on) that are physically *owned* by the logical partition that Linux will be running in. In this case, a physical storage adapter will be allocated to the logical partition that Linux will be running in. Keep in mind that only I/O Adapters (IOA) can be allocated to the Linux partition. The Linux operating system does not support I/O Processors (IOP) and they cannot be allocated to the partition. For planning purposes, the current allocation of physical adapters should be reviewed to determine what IOAs are available for allocation to the Linux partition.

You can review the physical hardware installed in the Managed System as well as current hardware assignment by selecting the I/O tab on the Properties page for the managed system as shown in Figure 2-12.

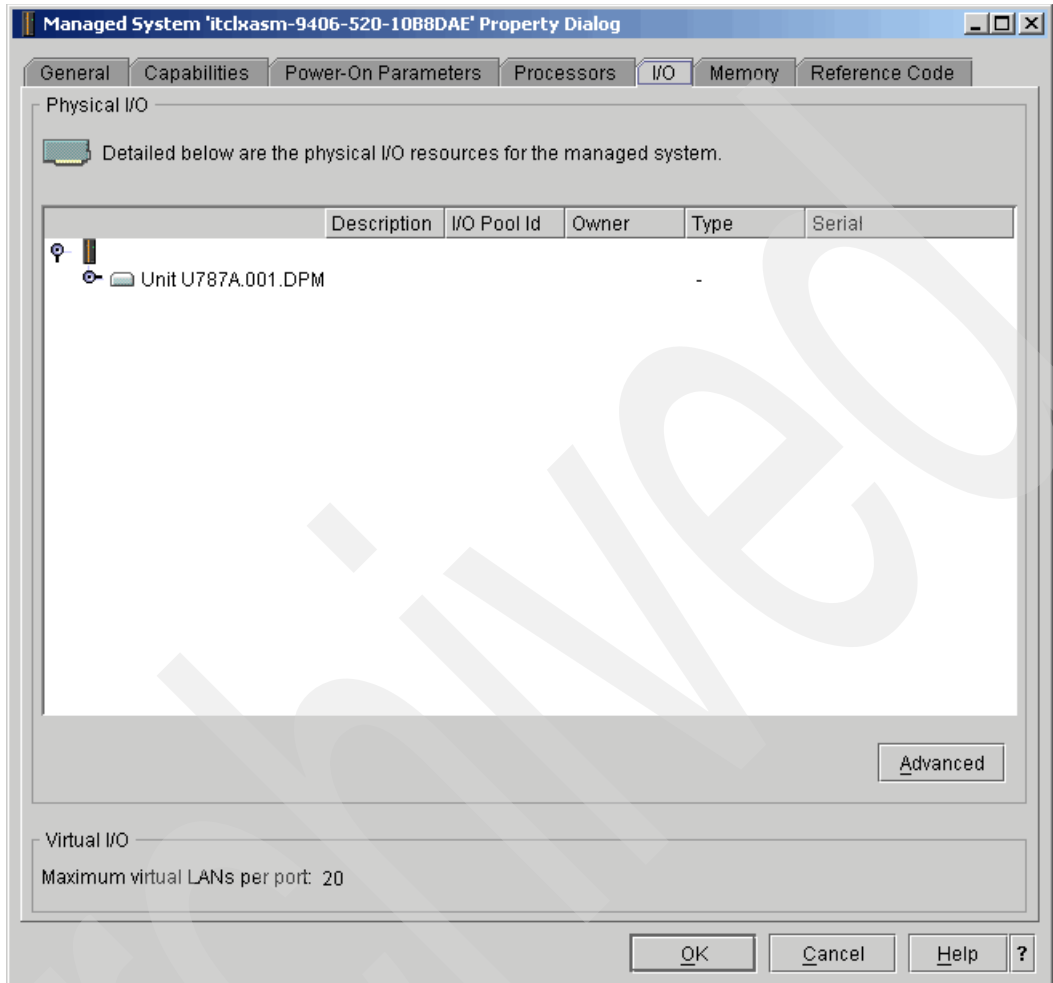


Figure 2-12 Managed System Properties - I/O

The hardware is displayed in a hierarchical form with units containing buses which contain IOPs and IOAs. Clicking on the symbol in front of each level will display the next level. Figure 2-13 shows the fully expanded list of hardware.

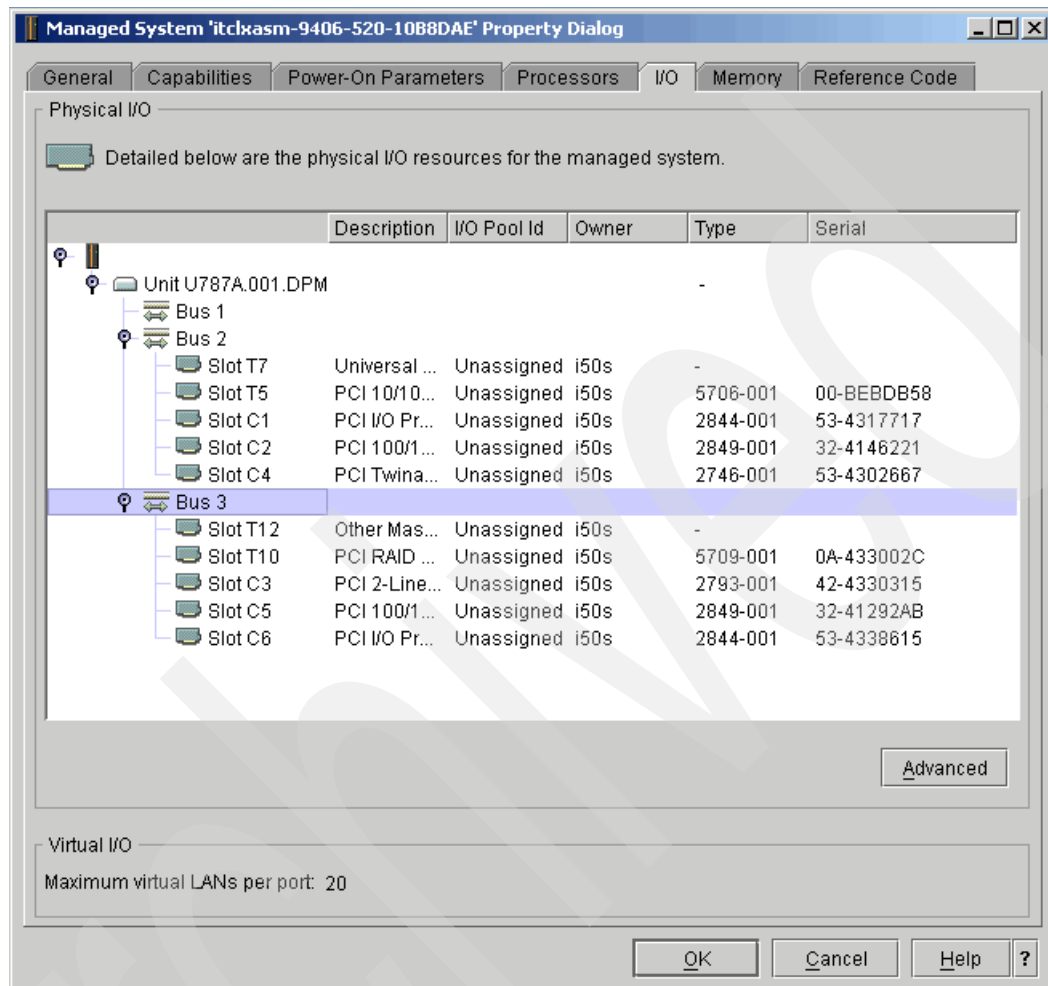


Figure 2-13 Managed System Properties - Expanded I/O Listing

The owner column indicates the active partition that currently has the resource assigned. When allocating physical resources to a Linux partition, remember that I/O adapters can be allocated, however, IOPs cannot be allocated. The ability to allocate an IOA that is downstream (or associated with) an IOP is dependent on the level of firmware the Managed System is running as well as the assignment of the IOP to another partition. Early versions of the system firmware did not allow the assignment of IOAs without the corresponding IOP, later versions of the firmware do allow such an assignment. Remember that sometimes when an i5/OS partition is defined, the hardware assignment will take place at the IOP level which will automatically take the related IOAs as well. In this case, you have to modify the partition or profile definition of the i5/OS partition and remove the IOA so that it can be assigned to a different partition.

2.3.4 Network configuration

The System i platform supports the allocation of both virtual and native network adapters to a logical partition (in fact both types of adapters can be allocated to the same logical partition). With virtual network adapters, network packets are “copied” between the memory allocation

of the partitions by the hypervisor and provides for very fast, reliable, and secure communications between logical partitions. Information on how network traffic can be forwarded between a virtual and physical network will be discussed in the networking chapter of this book.

In addition to virtual network adapters, physical network adapters can also be allocated to a logical partition. For Linux partitions, physical network adapters are typically used when Linux is being implemented as a firewall or when there is a heavy network bandwidth requirement on the workload being supported on Linux.

There are a number of considerations for determining the allocation of network adapters to a Linux partition. These considerations include the following:

- ▶ Requirements for fast intra-partition access (such as Open Database Connectivity (ODBC) to DB2/400)
- ▶ Requirement for external (that is, outside the System i platform) access to the Linux operating system
- ▶ Requirement for direct access (that is, firewall implementation) by the Linux partition to network traffic

While the assignment of virtual and physical network adapters will be very dependent on the environment as well as the availability of system resources, some general recommendations can be made:

- Linux partitions that are going to perform firewall functions should have a physical network adapter allocated to it for the external traffic, and one or more virtual network adapters for routing authenticated traffic to other partitions.
- A partition that is going to impose significant network traffic between itself and another partition (such as a Linux partition running the Apache Web server accessing DB2/400 data via ODBC) should implement virtual network adapters in both partitions to ensure that the traffic is passed at the gigabit speed provided by the virtual LAN.

2.4 Planning for virtual I/O

This section provides information in terms of planning the Linux implementation using virtual I/O.

2.4.1 Considerations for virtual I/O scenarios

With virtual I/O, the resources are owned and managed by the i5/OS logical partition that is hosting the Linux logical partition. The i5/OS shares these resources (disk, tape, CD-ROM, and so on) with the hosted Linux server. The i5/OS provides the DASD protection and some backup/restore facilities for the Linux environment. Virtual ethernet provides 1 Gb communications paths between logical partitions without requiring additional hardware resources.

Linux is started from the hosting i5/OS logical partition by varying in the network server description (NWSD). This Linux server can only be active when the i5/OS hosting logical partition is active. If the i5/OS hosting logical partition is in restricted state, then the NWSDs are in a varied off state.

The advantage of virtual I/O is that you do not need any dedicated hardware for the Linux partition, virtual ethernet is both a secure and fast way of communicating, and there is no

need for additional backup procedure beside i5/OS backup because Linux virtual disks are part of the i5/OS file system.

Both Red Hat Enterprise Linux 4 and SUSE Linux Enterprise Server 10 are able to work in a virtual I/O scenario.

2.4.2 Hardware requirements for a Linux logical partition

This section provides information about the minimum hardware requirements as well as supported devices for a Linux partition on the System i platform in which the Linux partition will be using all virtual resources.

Minimum configuration requirements

Each Linux partition requires the following minimum hardware resources:

Note: This configuration can vary depending of the Linux distribution and version that you want to install. Refer to a specific Linux distribution's guideline for the actual values.

- ▶ Processor unit
 - A whole processor, if you configure the partition with dedicated option
 - 0.1 processing unit, if you configure the partition with shared option
- ▶ Memory
 - 256 MB of memory
- ▶ Disk storage
 - One virtual storage adapter (vscsi) for virtual I/O
 - Approximately 1 GB of disk storage
- ▶ Network interface
 - One virtual ethernet adapter for virtual I/O

Note: In addition to these minimum hardware requirements, you also need to apply all critical fixes and updates for the HMC, firmware, and i5/OS.

Communication options

Linux on a System i platform can establish a TCP/IP connection through either a directly attached network adapter or through a virtual ethernet adapter. In a virtual I/O scenario, we are using a virtual ethernet adapter.

Virtual network adapters

Virtual ethernet provides the same function as using a 1 Gb ethernet adapter. i5/OS and Linux partitions on a System i platform can communicate using TCP/IP over the virtual ethernet communication ports. Virtual ethernet provides a very high speed, secure mechanism for communication among partitions on a single physical system.

Up to 4,094 separate virtual ethernets can be defined. Each partition can have up to 65,534 virtual ethernet adapters connected to the virtual switch.

The enabling and setup of a virtual ethernet does not require any special hardware or software. After a specific virtual ethernet is enabled for a partition, a network device named ethXX is created in the partition, where XX can be from 0 (that is, eth0) to 99 (that is, eth99). The user can then set up a TCP/IP configuration appropriately to communicate with other

partitions using the `ifconfig` command or any wizard that comes with the Linux operating system.

Virtual devices

Linux on System i platform supports virtual devices, as illustrated in Table 2-2.

Table 2-2 Virtual devices

Device	Linux driver	In Linux work as	In i5/OS work as
Virtual console	hvc_console	Server	Client
Virtual tape	st & ibmvscsic	Client	Server
Virtual CD	sr & ibmvscsic	Client	Server
Virtual disk unit	sd & ibmvscsic	Client	Server
Virtual SCSI	ibmvscsic	Client	Server
Virtual serial	is applicable using hvcs driver	Client	Client
Virtual ethernet	ibmveth	N/A*	N/A*

*Virtual ethernet works as peer-to-peer communication.

2.5 Planning for direct I/O

This section provides information in terms of planning a Linux implementation using direct I/O.

2.5.1 Considerations for direct I/O scenarios

Linux is able to communicate directly with System i platform hardware, such as disk units, ethernet adapters, tapes, and so on. Such a scenario is called *direct I/O scenario*. Direct I/O means that I/O resources dedicated to Linux partition are owned by and are under the control of Linux. It does not depend on i5/OS for any of its resources, consequently, i5/OS cannot use these resources, because they are allocated to the Linux server. For directly attached hardware, all failure and diagnostic messages are displayed within the Linux server.

With direct I/O, a Linux server can be active even when the i5/OS logical partition is not active, which is not the case in a virtual I/O scenario.

The disadvantage of direct I/O is that it demands more hardware administration. Also, with direct I/O, you need to implement additional backup procedures for Linux.

Both Red Hat Enterprise Linux 4 and SUSE Linux Enterprise Server 10 are able to work in a direct I/O scenario.

2.5.2 Hardware requirements for a Linux logical partition

The minimum hardware requirements for a Linux partition is impacted by a number of factors including the Linux distribution and the I/O structure (that is, virtual or native). This section provides general information about hardware requirements as well as supported devices for Linux partition on System i platform.

Note: The configuration can vary depending on the Linux distribution and version that you want to install. Refer to a specific Linux distribution's guideline for the actual values.

The minimum hardware requirements are:

- ▶ Processor unit
 - A whole processor if you configure the partition with dedicated processor option
 - 0.1 processing unit if you configure the partition with shared processor option
- ▶ Memory
 - 256 MB of memory
- ▶ Disk storage
 - One physical IOA for native I/O
 - One hard disk (at least 1 GB) (disk can be a native disk or virtual storage space)
- ▶ Network interface
 - One physical NIC (Network Interface Card) for native I/O or one virtual NIC (Network Interface Card) for virtual networking

Note: In addition to these minimum hardware requirements, you also need to apply all critical fixes and updates for the HMC and firmware.

Communication options

Linux on a System i platform can establish a TCP/IP connection through either a directly attached network adapter or through a virtual ethernet adapter. In a direct I/O scenario, we are using directly attached network adapters.

Directly attached network adapters

A Linux partition can own its physical ethernet adapters. If you have multiple ethernet adapters, you might consider dedicating one or more to the Linux partition. A dedicated adapter eliminates the extra step involved in using the virtual ethernet to communicate with the network as routing methods. You can even limit the traffic between one server with others.

Supported hardware resources

This section lists the devices and adapters supported at the time of writing of this book. Network adapters

Table 2-3 lists the network adapters supported on Linux on System i platform for native I/O, or directly attached option.

Table 2-3 Network adapters supported for native I/O Linux partition

Description	Linux Device Driver	Linux Feature Number	i5/OS Feature Number
PCI-X 1Gb ethernet-SX Fiber	e1000	0620	5700
PCI-X 10/100/1G Base-TX ethernet	e1000	0621	5701

Description	Linux Device Driver	Linux Feature Number	i5/OS Feature Number
PCI-X 10/100/1G 2-port Base-TX ethernet	e1000		5706
PCI-X 1Gb 2-port ethernet-SX Fiber	e1000		5707
PCI 100Mb TokenRing	olympic	0603	2744
PCI 1Gb ethernet	acenic	0601	2743
PCI 1Gb ethernet UTP	acenic	0602	2760
PCI 10/100Mb ethernet	pcnet32	0623	2849

Storage adapters

Table 2-4 lists the storage adapters supported on Linux on System i platform for native I/O, or directly attached.

Table 2-4 Storage adapters supported for native I/O Linux partition

Description	Linux Device Driver	Linux Feature Number	i5/OS Feature Number
PCI-X Dual Channel Ultra320 SCSI RAID and PCI-X Dual Channel Ultra320 SCSI RAID Blind Swap	ipr	0628	5703,5711
32MB Planar Dual Channel SCSI RAID Enablement	ipr		5709
PCI-X Dual Channel Ultra320 SCSI Blind Swap and PCI-X Dual Channel Ultra320 SCSI	ipr	0645	5710,5712
PCI Bus 2Gb Fibre Channel for 64-bit	lpfc		6228
PCI-X 2Gb Fibre Channel	lpfc		6239
PCI Fibre Channel Disk Controller	lpfc	0612	2766
PCI-X U320 RAID w/Read Cache	ipr	0627	2780
PCI-X Fibre Channel Disk Controller	lpfc	0626	2787

Wide area network (WAN) adapters

Table 2-5 lists the WAN adapters supported on Linux on System i platform for native I/O, or directly attached option.

Table 2-5 WAN adapters supported for native I/O Linux partition

Description	Linux Device Driver	Linux Feature Number	i5/OS Feature Number
PCI Two-Line WAN	icom	0613	2742
PCI Two-Line WAN	icom	0608	4745
PCI Dual WAN/Modem	icom	0609	2772
PCI Dual WAN/Modem (ANSI)	icom	0610	2773
PCI Two-Line WAN w/Modem	icom	0614	2793

Description	Linux Device Driver	Linux Feature Number	i5/OS Feature Number
PCI Two-Line WAN w/Modem (CIM)	icom	0615	2794
PCI Quad Modem	icom	0616	2805
PCI Quad Modem (CIM)	icom	0617	2806

Depending on the type of scenario (RAID or not) and type of DASDs, there are different procedures for preparing disk units, as illustrated in Table 2-6.

Table 2-6 Procedures depending on type of scenario and DASDs

	iSeries DASD - HW RAID	iSeries DASD - no HW RAID	pSeries® DASD - HW RAID	pSeries DASD - no HW RAID
Linux	Reformat to 512 and create RAID set	Reformat to 512	Reformat to 522 and create RAID set	Nothing

2.6 Planning for hybrid scenarios

This section provides information in terms of planning of Linux implementation using direct I/O and virtual I/O in a mixed mode.

2.6.1 Considerations for hybrid scenarios

As Linux can take advantage of virtual and direct I/O resources at the same time, some common configurations are hybrid, for example, combining direct and virtual I/O at the same time. Note that it is also possible to switch from virtual I/O to direct I/O or vice versa once we realize that the other solution will fit better.

Both Red Hat Enterprise Linux 4 and SUSE Linux Enterprise Server 10 are able to work in a hybrid I/O scenario.

2.6.2 Hardware requirements for a Linux logical partition

This section provides information about the minimum hardware requirements as well as supported devices for Linux partition on System i platform.

Minimum configuration requirements

Each Linux partition requires the following minimum hardware resources:

Note: This configuration can vary depending on the Linux distribution and version that you want to install. Refer to a specific Linux distribution's guideline for the actual values.

- ▶ Processor unit
 - A whole processor if you configure the partition with dedicated option
 - A 0.1 processing unit if you configure the partition with shared option

- ▶ Memory
 - 256 MB of memory
- ▶ Disk storage
 - One virtual storage adapter (vscsi) for virtual I/O or one physical IOA for native I/O
 - Approximately 1 GB of disk storage
- ▶ Network interface
 - One virtual ethernet adapter for virtual I/O or one physical NIC (Network Interface Card) for native I/O

Note: In addition to these minimum hardware requirements, you also need to apply all critical fixes and updates for the HMC, firmware, and i5/OS.

2.7 Planning for VPM

Virtual Partition Manager (VPM) is for simple Linux configurations. You can create up to five LPAR partitions where you can have a maximum of one i5/OS partition and up to four Linux partitions without HMC.

VPM leverages i5/OS's DST-like interfaces, that is, 5250 command-line interfaces, which does not require the investment in an HMC.

Table 2-7 summarizes the functions and limitation of VPM compared to HMC.

Table 2-7 VPM and HMC comparison

	VPM	HMC
OS supported	i5/OS and Linux	i5/OS, Linux, and AIX 5L
Max# of partitions	Five (One i5/OS and four Linux)	254
Uncapped partition support	Yes	Yes
Dynamic Resource Movement	No	Yes
I/O support for Linux	Virtual	Virtual and direct
Max# of virtual ethernet	Four	4096
Max virtual disk per partition	64 TB	64 TB

For more info about VPM, refer to *Virtual Partition Manager A Guide to Planning and Implementation*, REDP-4013.

2.8 Power control and console options

This section provides two items of consideration before you actually configure the system for the Linux partition: How to set up the power control option and alternative console option.

2.8.1 Power control

There are four types of power control of the Linux partition:

- ▶ Automatic start of the Linux partition: You choose this option when you add to the system profile the partition profile that you want to be activated by the managed system every time when the system starts up.

Note: Automatic start of the Linux partition requires that the partition have dedicated disk resources as the partition may be started prior to any hosting i5/OS partition being started.

- ▶ Activate the Linux partition from HMC: This option is used to choose a special boot mode other than normal boot mode or to override the default boot mode defined in the partition profile. Special boot modes could include System Managed Services (SMS) boot mode or other diagnostic boot mode which are generally used for diagnosis purpose.

Note: Activation of the partition from the HMC is typically performed for a Linux partition with dedicated disk resources. If activating a Linux partition from the HMC with virtual disk resources it will be necessary to vary on the Network Server in the hosting i5/OS partition prior to activation of the Linux partition from the HMC.

- ▶ Activate the Linux partition from i5/OS: You choose this option when you want to start the Linux partition by activating the NWSD. This option is valid when two conditions are met:
 - NWSD's PWRCTL parameter value is *Yes.
 - Power Controlling option on the Linux partition profile has the i5/OS partition as a power controlling partition.

Then you can start the Linux server from i5/OS by a vary function on the NWSD.

- ▶ Activate the Linux partition with i5/OS IPL: This option starts up the Linux partition when the i5/OS partition starts up if all three conditions are met:
 - NWSD's PWRCTL parameter's value is *Yes.
 - Power Controlling option on the Linux partition profile has the i5/OS partition as a power controlling partition.
 - NWSD's ONLINE parameter's value is *Yes.

2.8.2 Linux console options

In addition to regular console options, there are two special console options for your Linux partitions:

- ▶ Virtual console
- ▶ HMC's VTerm

They can be used for:

- ▶ The initial installation where TCP/IP configuration is not done yet on the Linux partition.
- ▶ Solving problems that result in the Linux becoming inaccessible to the LAN.

Virtual console

Virtual console is served through the server i5/OS partition's Telnet server. Perform the following steps for it to work:

1. From any Telnet client, such as Windows PC or PuTTY, log in to the i5/OS partition over port 2301. For example, on your DOS prompt, type:


```
C:>telnet <i5/OS_server_host_name> 2301
```
2. Select the Linux partition ID.
3. Log in using i5/OS System Services Tool (SST) user ID and its password.
4. Log in as any Linux user ID such as root.

To be able to access the virtual console through the i5/OS partition's Telnet server you will need to define a virtual Serial Client adapter in the i5/OS partition that connects to the virtual Serial Server adapter (slot 0) in the Linux partition. Such a configuration is shown in Figure 2-11 on page 39.

HMC's VTerm

You can open the terminal sessions from the HMC to the i5/OS partition, Linux partition, and AIX partition. You do not need to configure anything for this feature. It just comes with the HMC code itself. To start VTerm for the Linux partition from HMC, there are two ways:

- ▶ If the Linux partition has not been started yet, activate the partition first to open a terminal session.
- ▶ If the Linux partition is already started, on HMC main menu, select **Server and Partition** → **Server Management**. Then choose the Linux partition and right-click it. Choose **Open Terminal Session** option to open a VTerm session.

Archived

Creating logical partitions

This section will present methods for creating logical partitions (LPARs) on the System i platform for supporting Linux instances. Information in this section will include the use of the Hardware Management Console (HMC) for creating logical partitions as well as the use of Virtual Partition Manager for creating of logical partitions. Creating logical partitions with both virtual I/O as well as I/O dedicated to the Linux partitions will be presented.

Our discussion of HMC here is focused on configuration of Logical partitions for Linux. For general discussions of HMC and LPAR, refer to *Logical Partitions on System i5: A Guide to Planning and Configuring LPAR with HMC on System i*, SG24-8000.

Information presented in this chapter is based on the different scenarios that you are likely to be establishing Linux partitions for. Essentially four different scenarios will be presented:

- ▶ Linux partition with all virtual I/O resources via HMC
- ▶ Linux partition with all virtual I/O resources via Virtual Partition Manager
- ▶ Linux partition with virtual disk and physical network via HMC
- ▶ Linux partition with all physical resources via HMC

3.1 Concepts and terminology

Here are concepts and terms you should know when creating the logical partition on the System i platform for the Linux instance.

3.1.1 Logical partitions

By using LPAR, a single physical system can be divided into multiple logical partitions each running their own operating system image.

LPAR provides a framework for systems consolidation on System i platform.

Partition ID

The partition ID is a whole number used to identify logical partitions.

Partition profile

A *partition profile* specifies the resources and settings for a logical partition. This includes the memory, processor and I/O allocations. To activate a logical partition, one of the partition profiles for that logical partition must be activated.

You may have multiple partition profiles with different resource specifications for a single logical partition based on the needs or requirements of that logical partition. Note, however, that only one partition profile can be active at a time. Activating a different partition profile requires the logical partition to be shutdown.

If you have multiple partition profiles, any one of them can be designated as the default partition profile. The HMC will activate the default profile unless you specify a different partition profile to be activated.

It is possible that a partition profile will not activate due to an overcommitment of resources on the system. The HMC shows all of the resources available on the system but does not verify if these resources are already in use by an activated partition profile. As a partition profile is activated, the system attempts to allocate the resources specified. If the resources are already in use, the partition profile will not activate.

Minimum, desired, maximum values

As you set up the partition profile, you will be asked to input the minimum, desired, and maximum values for memory and processor units. If your system resources are not overcommitted, the logical partition will get the desired values. However, if resources are overcommitted, the logical partition will be given a value between the minimum and desired values. If the minimum value cannot be met, the partition profile will not activate, meaning you cannot start the particular server, that is, the operating system, of that partition. The maximum indicates the maximum value that may be dynamically set.

Full system partition profile

A full system partition profile is a partition profile that has been setup to use all the resources of the system including memory, processors, I/O, disks. A full system partition profile can be used for i5/OS only.

When this partition profile is activated, all system resources are committed to the associated logical partition. No other partition profiles and logical partitions will be allowed to be activated as long as this full partition profile and associated logical partition are active. Conversely, a

full system partition profile cannot be activated when other partition profiles and logical partitions are already active.

If additional hardware resources are added to the system and then a full system partition profile is activated, the associated logical partition automatically recognizes and uses the new hardware.

System profile

A system profile is an ordered list of partition profiles on the managed system. When a system profile is activated, the managed systems will attempt to activate the partition profiles in the order they are listed in the system profile.

A system profile is helpful when changing the managed system from one set of logical partitions to another.

It is possible to create a system profile that contains a partition profile with overcommitted resources. The HMC can be used to validate the system profile against currently available resources or against the total system resources. This validation will ensure your I/O devices and processing resources are not overcommitted. However, memory requirements are only estimated. It is possible for a system profile to pass the validation test, but not have enough memory to be activated.

Service partition

The HMC reports hardware errors to IBM. If the HMC is unavailable, then a service partition can report errors. This partition has the authority needed to update the system and other policy parameters without having to power off the System i platform. On System i platform, only an i5/OS logical partition can act as the service partition. This service partition, typically, has a physical connection to a network attached to the HMC and virtual connections to the other logical partitions on the System i platform. This allows the service partition to receive server errors from the other logical partitions and report the server errors to IBM.

3.2 Linux partition with all virtual resources via HMC

This is probably the most common configuration for a Linux partition on the System i platform. In this configuration, all of the resources (that is, disk, network) that will be used by the Linux instance are actually owned by another operating system (i.e., an i5/OS partition). With this configuration only virtual I/O adapters will be allocated to the Linux partition.

This section will walk through the creation of the Logical partition for Linux through the use of the Hardware Management Console.

The Hardware Management console is used to create partitions and partition profiles. Partition profiles are a unique aspect of partitioning on the System i platform (i.e., partition profiles were not used on the previous versions of LPAR on iSeries). If you are new to LPAR or from a previous version of LPAR on the POWER4 version of the iSeries (as with OS/400 V5R2 and previous versions), you have to be familiar with this new concept. Refer to *Logical Partitions on System i5: A Guide to Planning and Configuring LPAR with HMC on System i*, SG24-8000.

HMC has a wizard for both partition and partition profile creation. The following steps describe how to use this wizard.

Note: To create a partition or system profile you must have administrator or operator privileges in the HMC.

Creating the logical partition environment to support Linux can be broken down into three broad areas:

- ▶ Creating the logical partition
- ▶ Adding virtual slots (for support of Virtual I/O) to the i5/OS partition
- ▶ Registering the partition with Open Firmware

3.2.1 Creating the logical partition

This section describes the task of creating the logical partition where you install Linux system.

Creating a new Linux partition

1. In the navigation area of the HMC, select **Server and Partition** → **Server Management**. In the contents pane, right-click **Partitions** and select **Create** → **Logical Partition** as illustrated in Figure 3-1.

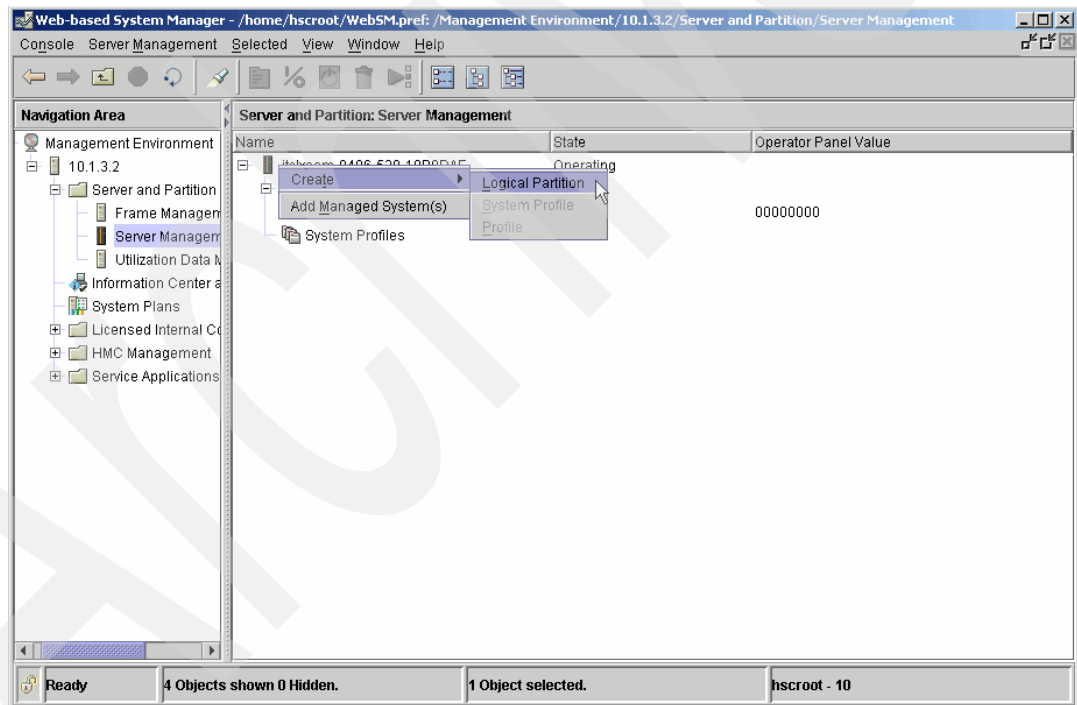


Figure 3-1 Starting the logical partition wizard

2. The Partition wizard screen appears. The Partition ID is an integer number that is used by the system to uniquely identify the partition and typically is left at the default value displayed. Provide the partition name and specify that the partition environment is AIX or

Linux as illustrated in Figure 3-2. Once the fields have been completed, click the **Next** button to continue.

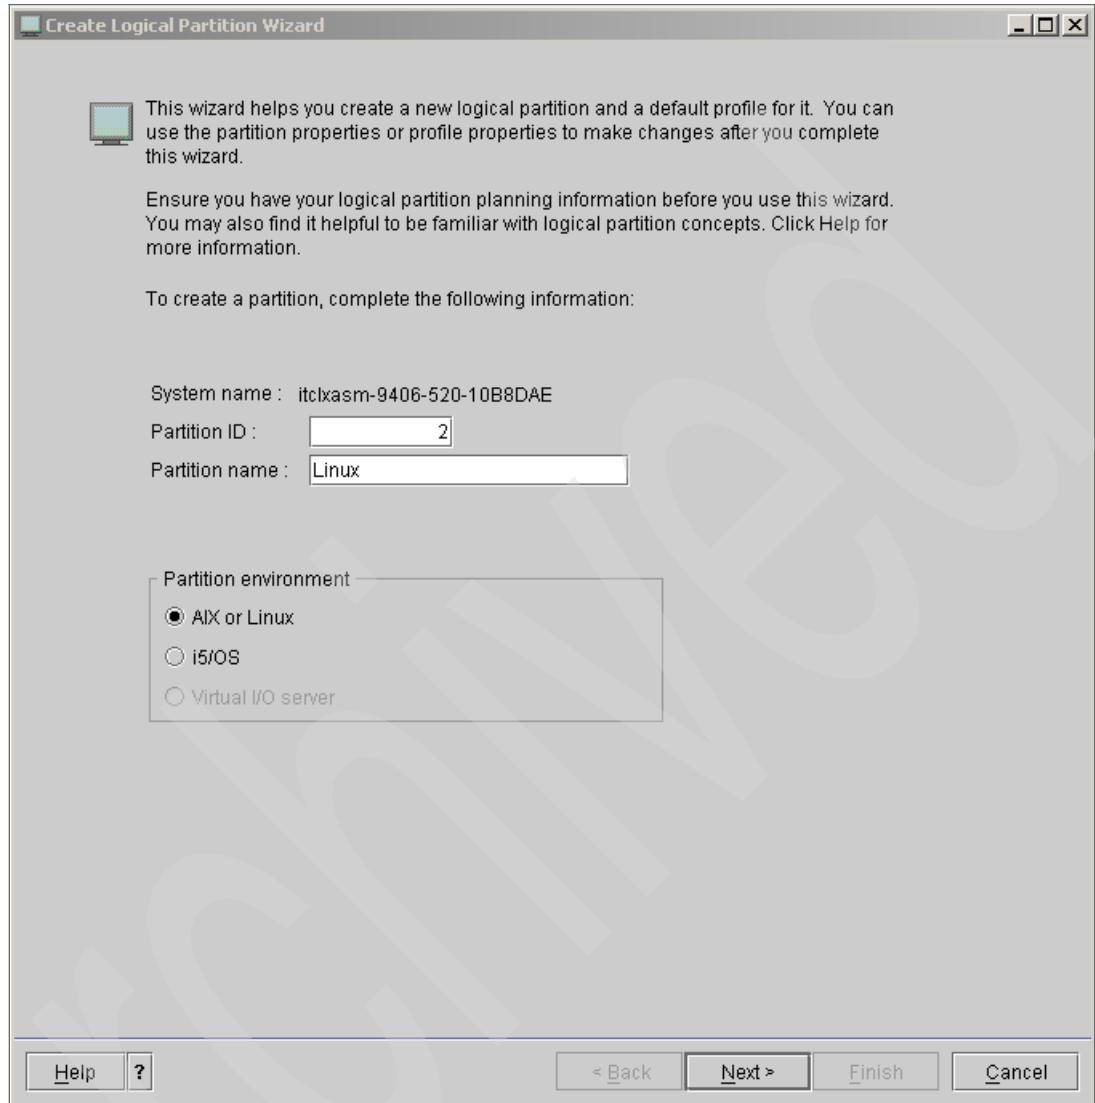


Figure 3-2 Creating a new partition

3. This next window provides you the option for the partition to be a part of a group of partitions that can be managed by a workload application. We choose **No**, as illustrated in Figure 3-3. Click **Next** to continue.

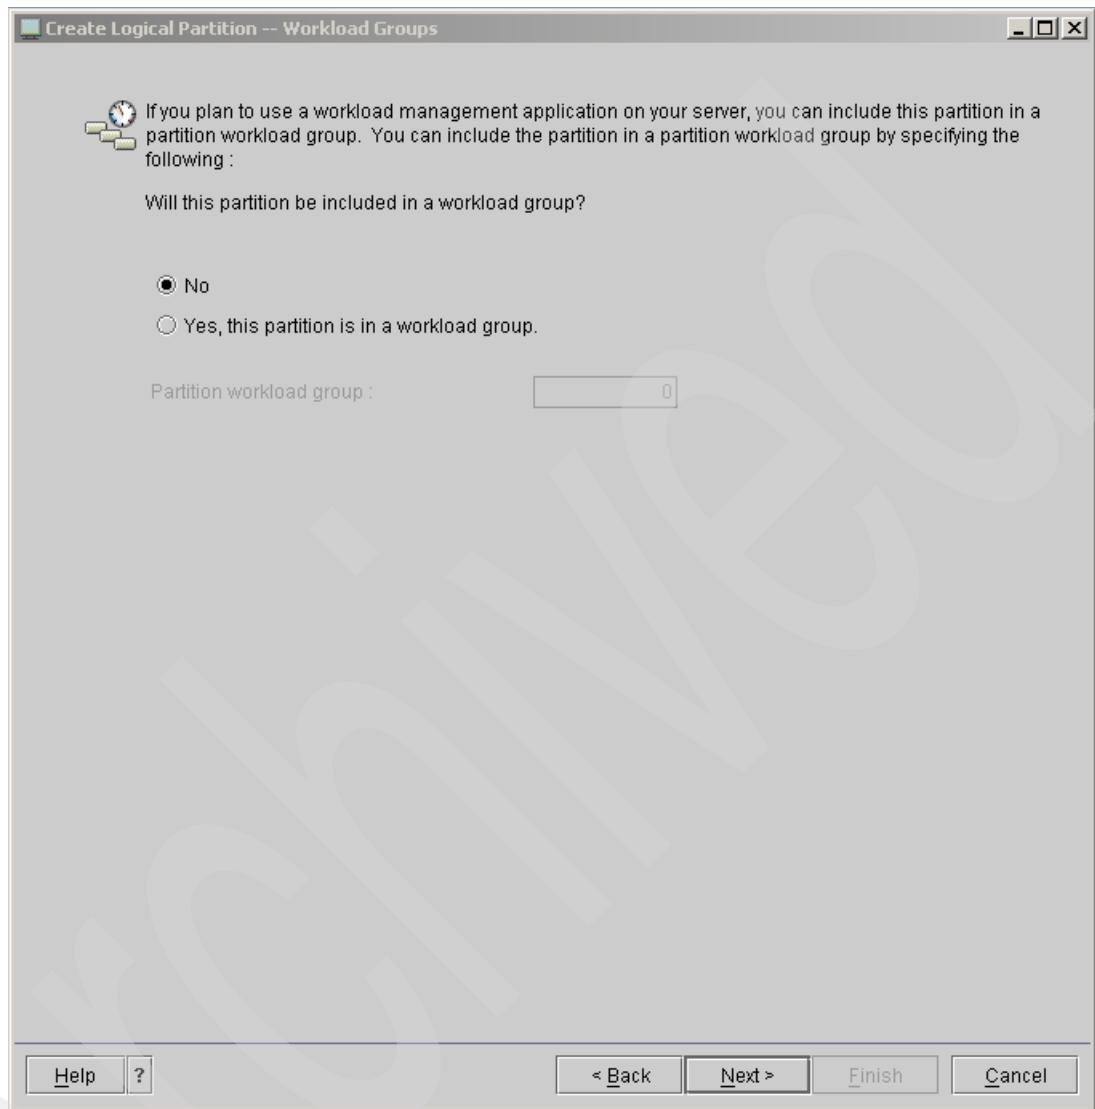


Figure 3-3 Creating a new partition - workload

- The next step is creating a partition profile. A partition profile specifies characteristics of the partition, such as the memory, processors, physical and virtual adapters.

As illustrated in Figure 3-4, give the name of the profile. In our example, we use Default as the profile name. Click **Next** to continue.

Note: Do not check Use all resources in the system option. If you check it, the partition that is associated with this profile will try to obtain all of the physical resources in the system when the partition profile is activated. If a partition owns all of the physical resources, then there can only be one active partition on the system.

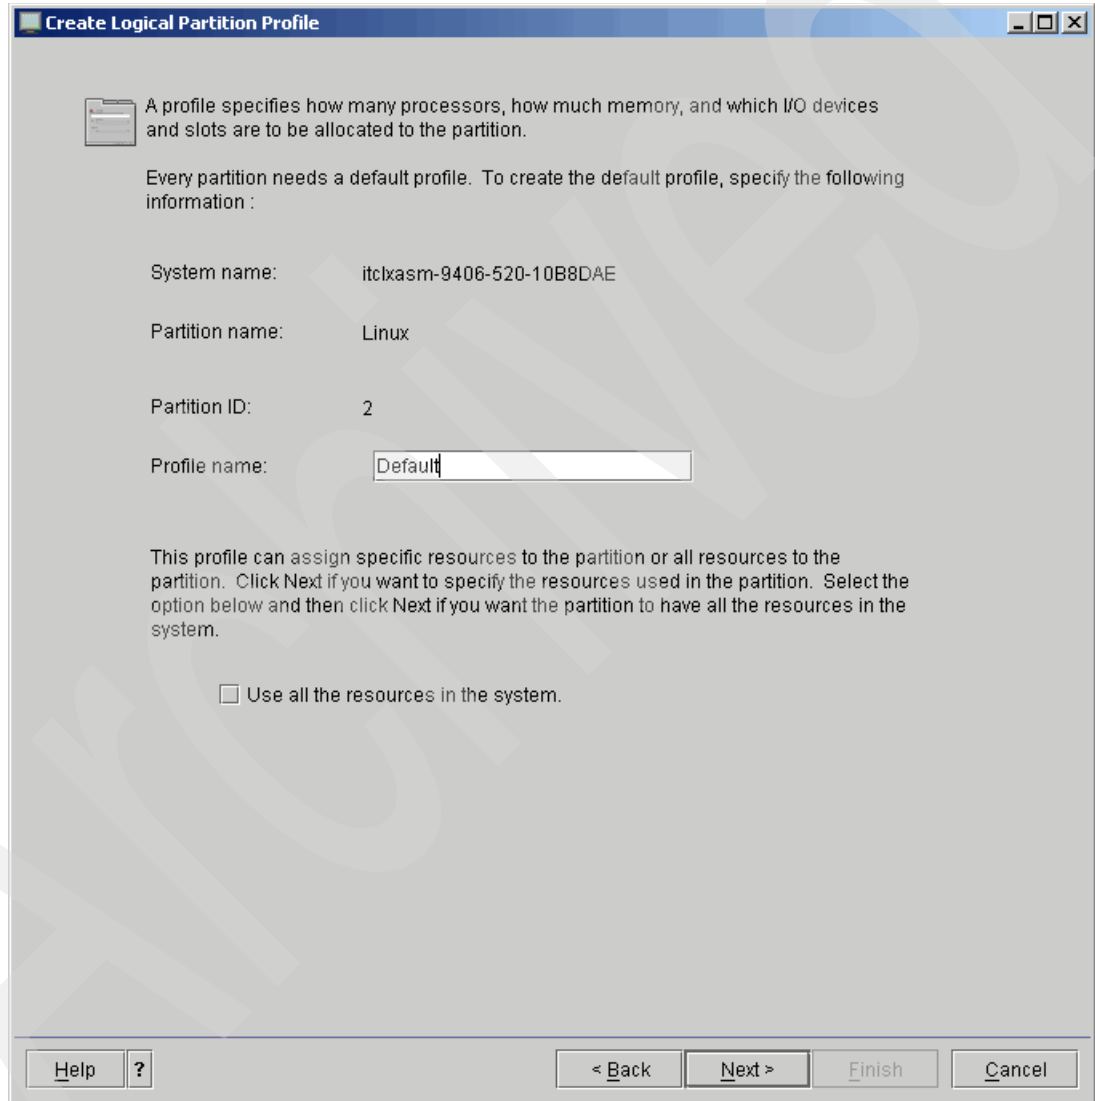


Figure 3-4 Creating a new partition - logical partition profile

5. Specify the memory size that the partition profile will manage, as illustrated in Figure 3-5. Use this page to specify the memory management information for this partition profile. You must specify three memory levels:
- The minimum memory is the minimum amount of memory that the logical partition must have to start up. If the system cannot allocate this much memory, the activation of this partition will fail.
 - The desired memory is the amount of memory that you want the logical partition to have when you activate the logical partition.
 - The maximum memory is the maximum amount of memory that the logical partition is allowed to have when you dynamically move memory.

When the partition is activated, an attempt will be made to allocate the *desired* amount of memory defined for the partition-profile. If the amount of desired memory is not available, then an additional check will be made to see if the amount of unallocated memory left on the Managed System is more then or equal to the *minimum* amount of memory defined for the partition profile. If the amount of memory is equal to or greater then the minimum defined memory, then it will be allocated to the partition, otherwise, the partition will fail to activate.

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

Installed memory (MB): 16384

Current memory available for partition usage (MB) : 15872

Minimum memory	Desired memory	Maximum memory
0 GB	0 GB	0 GB
256 MB	512 MB	512 MB

Help ? < Back Next > Finish Cancel

Figure 3-5 Creating a new partition - memory

In our example, the minimum quantity of memory that the partition will be activated with is 256 MB, the desired amount of memory is 512 MB and the maximum amount of memory is also 512 MB. Click **Next** to continue.

Note: Dynamic memory movement (that is, adding or removing memory from the partition while the Linux operating system is active) is not supported in the 2.6 and 2.4 versions of the Linux kernel (the kernel in current supported distributions of Linux on the System i platform). Therefore, the ability to dynamically change the memory allocation for Linux is not supported on the System i platform.

In addition to the memory allocated to the partition, the hypervisor sets aside memory to manage the partitions memory. This additional memory is called the Hardware Paging Table and is based on the size of the maximum memory setting.

- The next step is to choose Shared or Dedicated processors. Dedicated means you are not sharing the processor with other partitions. With the Shared option, you can utilize the subprocessor feature where each partition can run with 0.1 processing unit at minimum. While the minimum processor (for shared) is .10, the definition can be in the hundredths of processor units (for example, 0.15). This is referred to as *micro-partitioning*. In our example, we choose the Shared option, as illustrated in Figure 3-6.



Figure 3-6 Creating a new Partition - Processor

- For the Shared processor option, we choose the minimum, maximum, and desired processing units of the partition.

When the partition is activated, an attempt will be made to allocate the “desired” amount of processors defined for the partition-profile. If the amount of desired processors is not available, then an additional check will be made to see if the amount of unallocated processors left on the Managed System is more then or equal to the “minimum” amount of processors defined for the partition profile. If the amount of processors is equal to or greater then the minimum defined processors, then it will be allocated to the partition, otherwise, the partition will fail to activate.

As illustrated in Figure 3-7, we define 0.1 unit of a processor to be our minimum processing units, 1 unit of processor power as our desired amount, and 1.5 as the maximum processing units. Click the **Advanced** button.

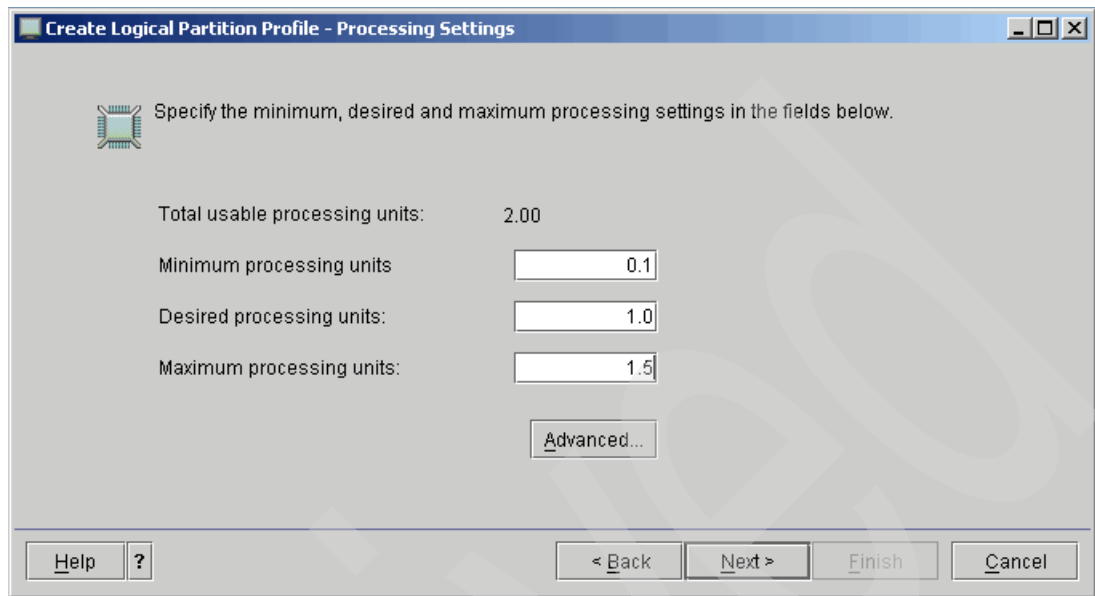


Figure 3-7 Creating a new Partition - Processing Units

8. When shared processing units are defined in a partition profile, additional settings can be set for the number of virtual processors as well as the capped/uncapped setting. As shown in Figure 3-8, the sharing mode can be set to either Capped or Uncapped. An uncapped partition will have the ability to have additional processing units allocated to it by the Hypervisor on an as-needed basis, based on the performance of the overall system. Capped partitions will only have access to the processing units allocated to it when first started (or changed through Dynamic LPAR).

For Uncapped partitions, a weight will be specified that indicates the ratio of left over (or surplus) processing units that will be allocated to the partition. As an example, if two partitions are configured as uncapped with a weight of 10 and 20 respectively and both are busy, then the partition with a weight of 20 will get two-thirds of the surplus processor resource while the partition that has a weight of 10 will get one-third of the surplus processor resource.

In addition to the sharing mode, the number of virtual processors is also defined through the Advanced Processing Settings page. The number of virtual processors indicates the number of processor threads that the workload allocated processing units will be spread across.

In this example, we set the sharing mode of the partition to uncapped with a weight of 128 and set the number of virtual processors to a minimum of 1, desired of 1 and a maximum of 2. Once the values have been set click **Ok**, then click **Next** on the Create Logical Partition Profile - Processing Settings window.

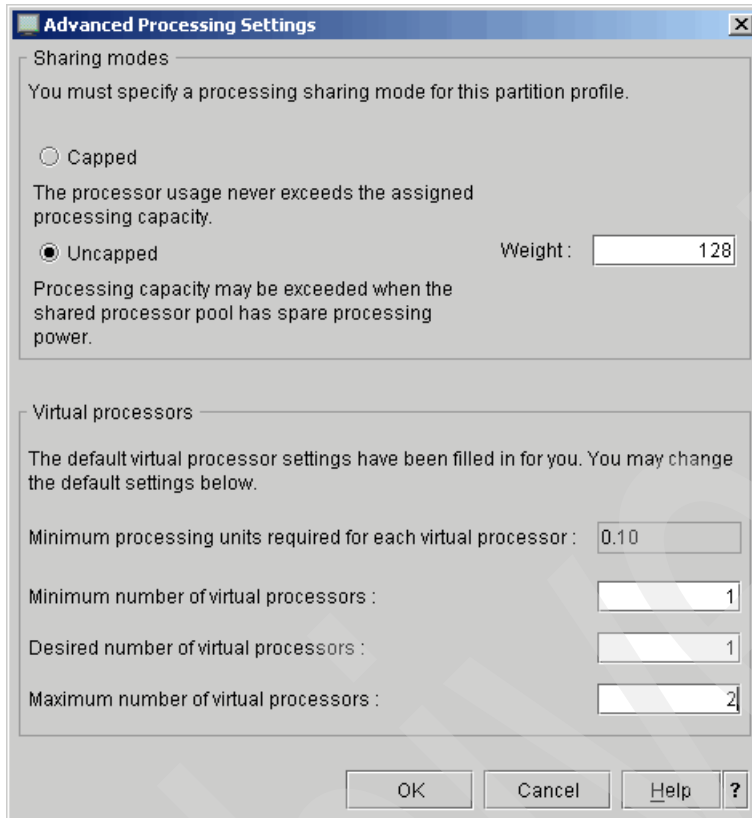


Figure 3-8 Creating a new Partition - Advanced Processing Settings

9. The next dialog displayed is for the allocation of physical hardware resources to the partition. As illustrated in Figure 3-9, the window displays the hardware installed in the

managed system and allows for selection of hardware for the partition profile. Since this partition profile will be using all virtual resources click **Next**.

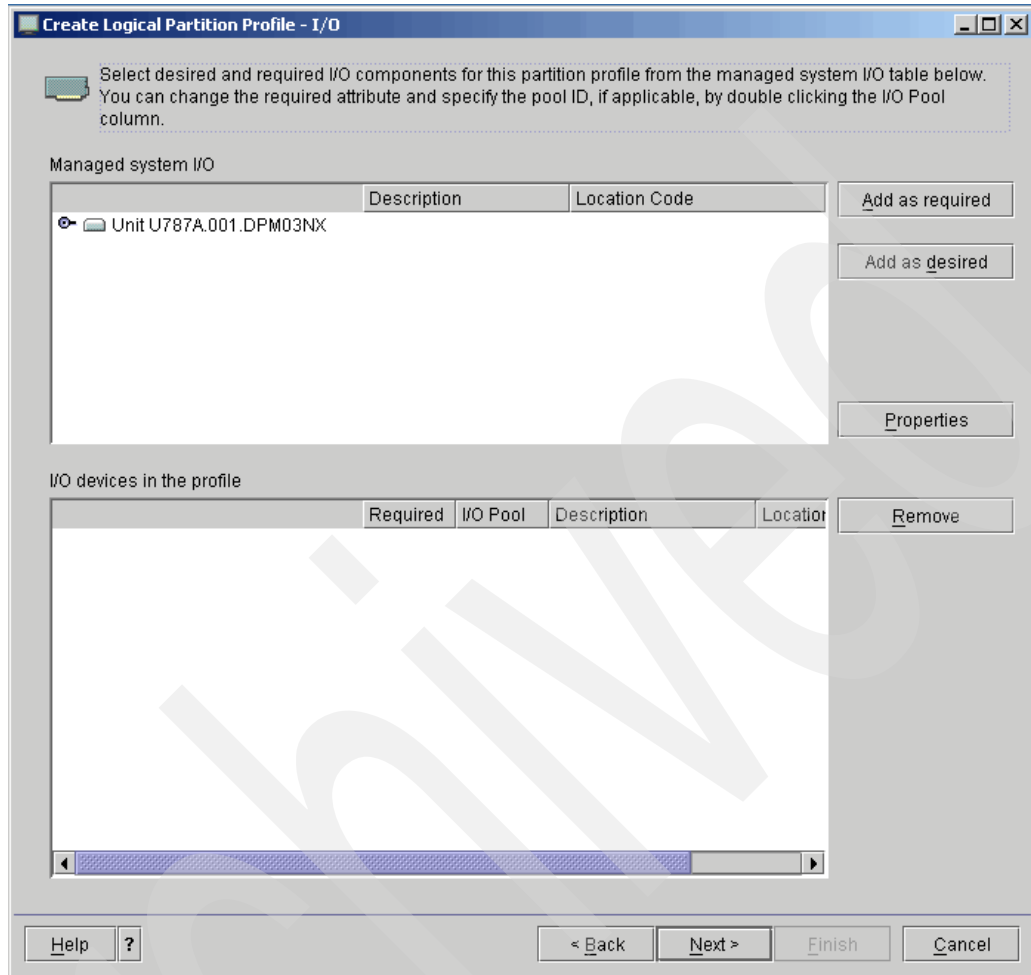


Figure 3-9 Creating a new partition - I/O

10. The next window, illustrated in Figure 3-10, is for adding the I/O pools. Linux partitions do not use I/O pools, so click **Next** to continue.

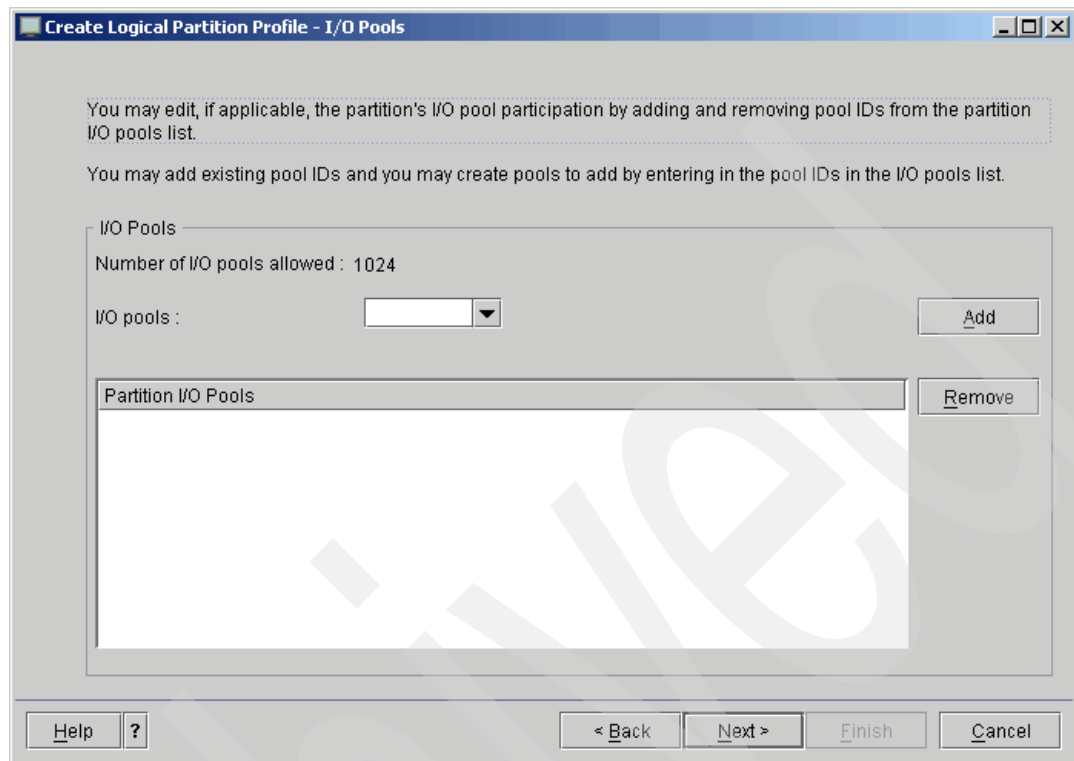


Figure 3-10 Creating a new partition - I/O pools

11. The Virtual I/O Adapters window appears as shown in Figure 3-11. Since this partition will be using hosted I/O (that is, I/O that is owned by another partition), select **Yes** as illustrated in Figure 3-11. Click **Next** to continue.

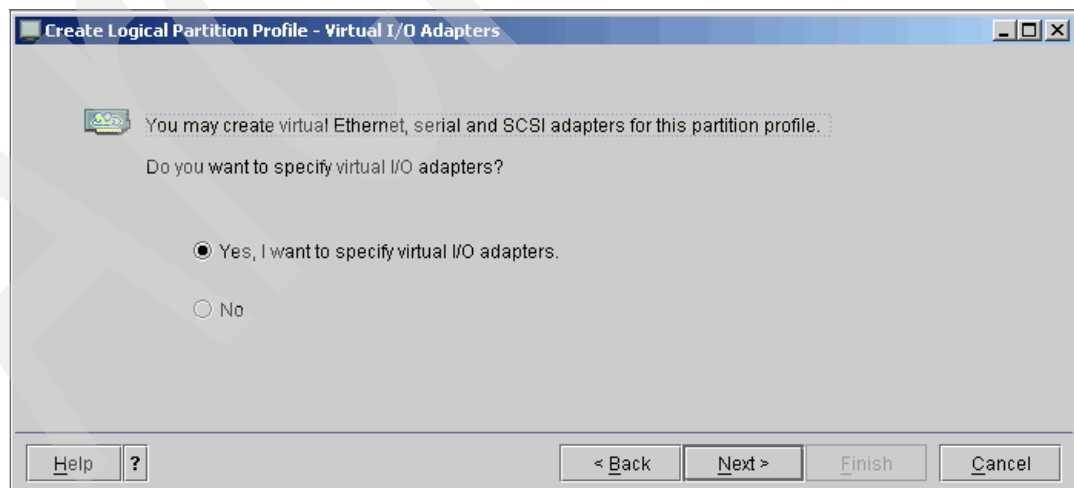


Figure 3-11 Creating a new partition - Virtual I/O Adapters

12. Every logical partition created on the System i platform has two virtual adapters automatically created for it as shown in Figure 3-12. As we will see later, the two Server Serial adapters are typically used to establish terminal access to the operating system running in the logical partition.

The Virtual Adapters window shows the currently defined adapters in the lower portion of the window and provides for the creation of additional adapters.

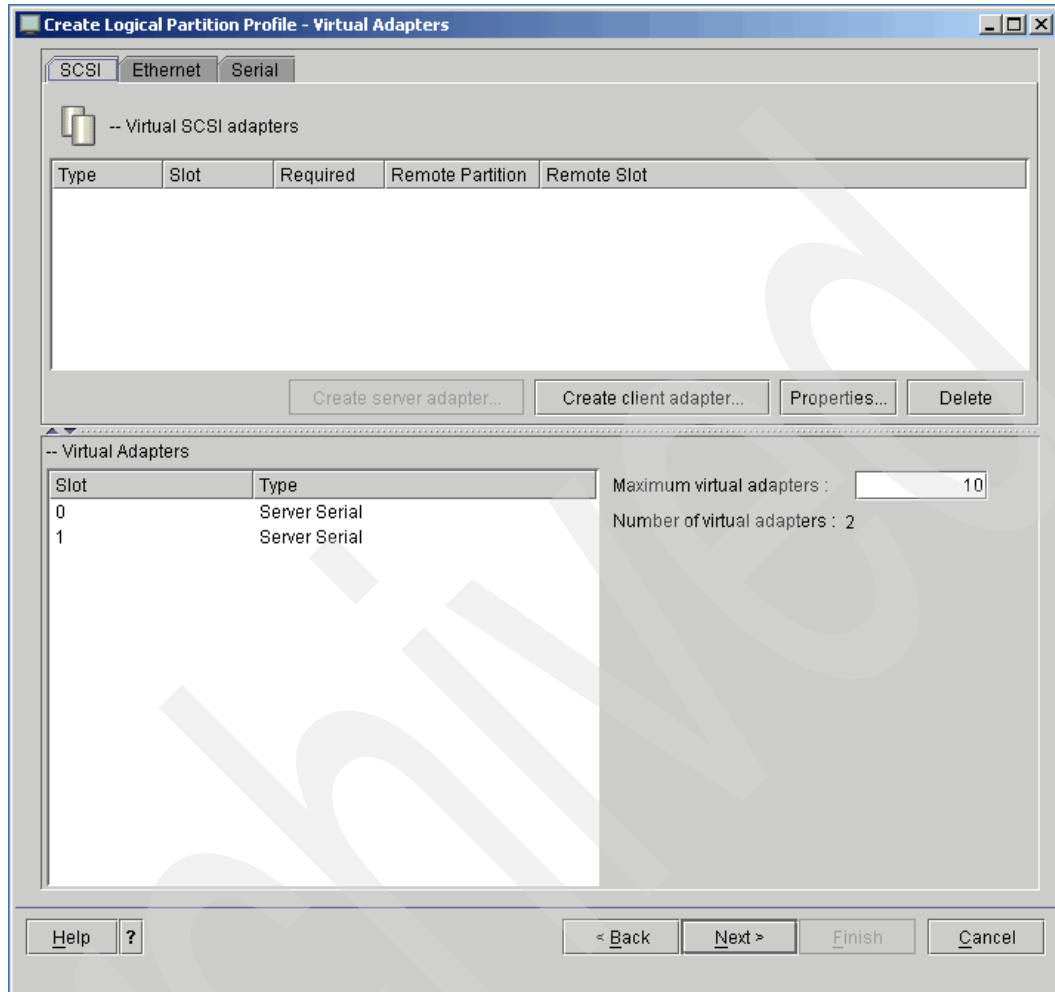


Figure 3-12 Creating a new partition - Virtual Adapters

Documenting the virtual slot definitions as the logical partitions are built can be beneficial both in avoiding partition configuration problems as well as providing the basis for an overall system configuration document. For this example, assume that the Managed

System has a single i5/OS partition that only has the default virtual I/O adapters as shown in Figure 3-13.

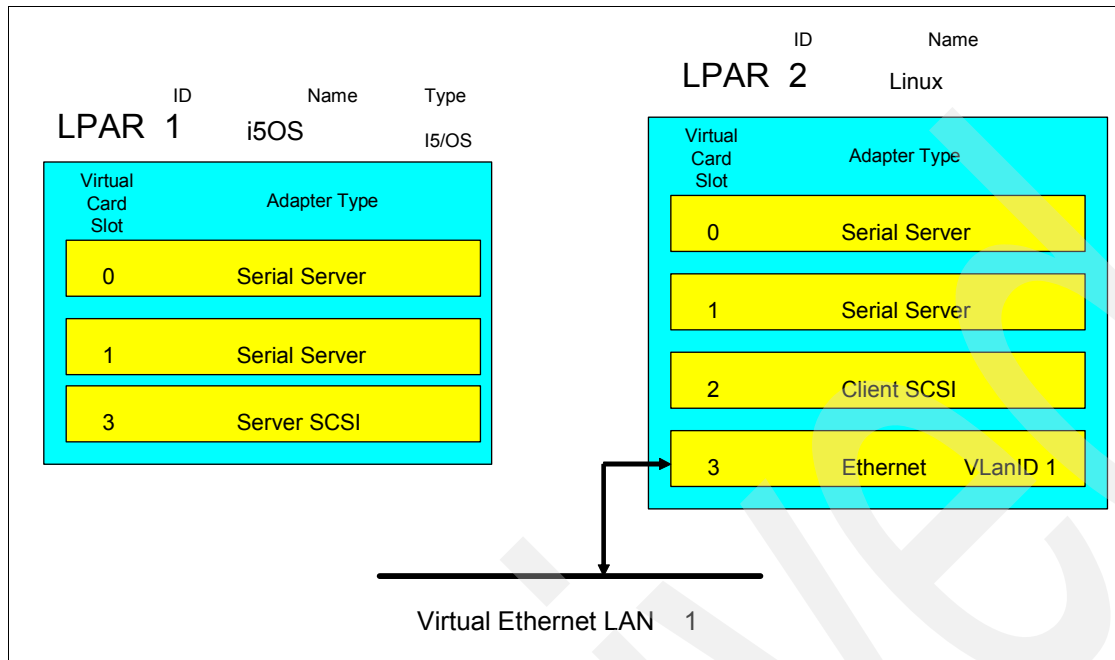


Figure 3-13 Virtual I/O Configuration Diagram

The Virtual I/O Configuration diagram will be updated as we continue to build the LPAR environment.

13. Access to I/O devices hosted by another partition will require the creation of a Small Computer System Interface (SCSI) adapter pairing between this partition (referred to as the *guest partition*) and the partition that owns the I/O resources (referred to as the *host partition*). We will create the client side of that pairing here. Click **Create client adapter**.
14. As shown in Figure 3-14, the definition of the Virtual SCSI client adapter consists of three items. The client slot is the next available virtual slot number for this partition and is typically left at the system determined value. The Server partition is the name of the i5/OS partition that will host the I/O resources for this partition. In this example, we select V5R4 from the drop-down list. The Server partition slot is the virtual slot number in the hosting i5/OS partition where the SCSI Server adapter will be created. Typically, the virtual I/O adapters definition of the i5/OS hosting partition will be reviewed prior to creation of the

Linux partition. By reviewing the i5/OS virtual adapter definition you would have an indication of the adapter slots that have been used and the available slot numbers.

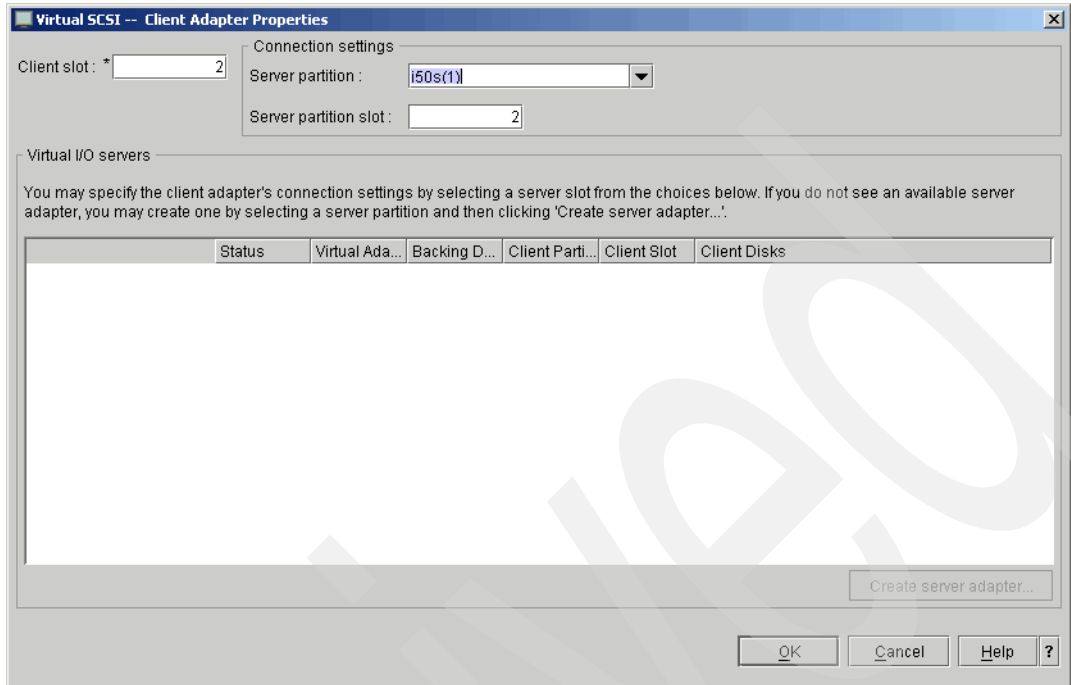


Figure 3-14 Creating a new Partition - Virtual SCSI adapter

15. In addition to the Virtual SCSI client adapter, you have to create a Virtual Ethernet adapter to support network communications for the partition. Back on the Create Logical Partition

Profile - Virtual Adapters, click the **Ethernet** tab to display the current defined Ethernet adapters for the partition as shown in Figure 3-15.

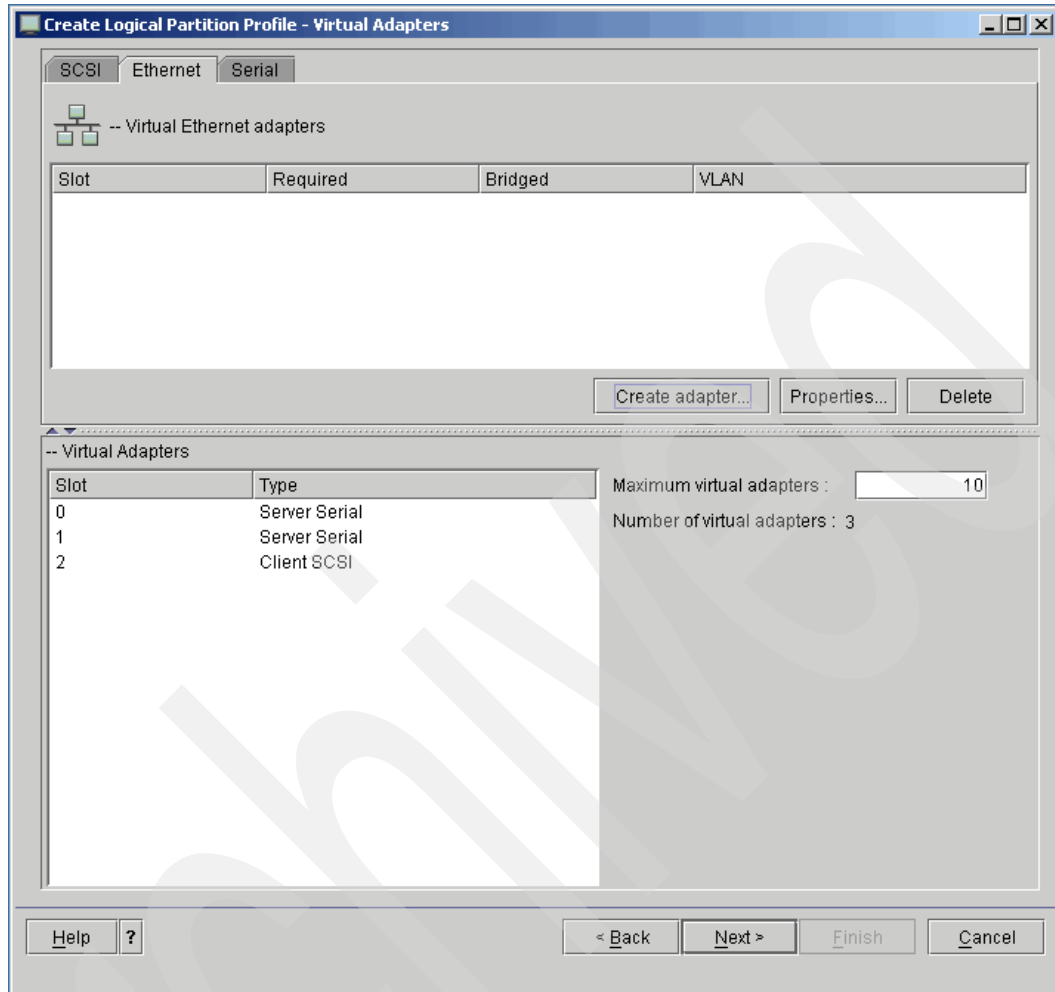


Figure 3-15 Create a new Partition - Virtual Adapters (Ethernet Tab)

Once the Ethernet tab has been selected, click the **Create adapter...** button. Depending on your system configuration, an error dialog box similar to that shown in Figure 3-16 may be displayed. This error occurs when the System i platform does not have a partition

created of type VIO Server. This is a typical condition on the System i platform and the error dialog box can be safely ignored.

Note: VIO Server is an AIX-based appliance that can be used to provide virtual I/O resources to guest partitions on System i or System p systems. There is an additional feature code that would need to be ordered to enable VIO Server on the system. A description of the functions and configuration of VIO Server is beyond the scope of this book.

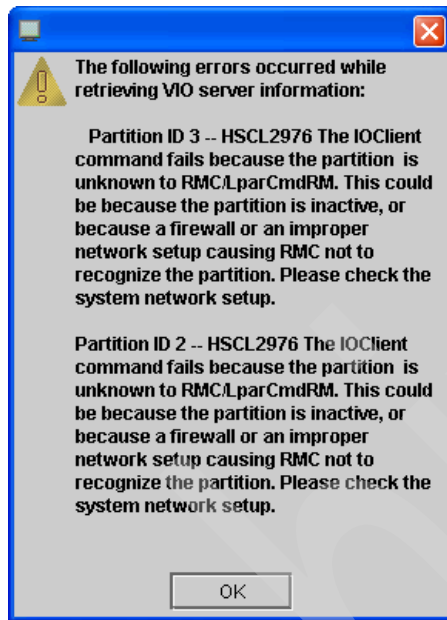


Figure 3-16 Creating a new Partition - RMC Error Dialog

The Virtual Ethernet adapter is defined completely in the fields shown in Figure 3-17. The slot number is the next available slot number as determined by the Create Partition wizard and is usually left as-is. The Virtual LAN ID is the Virtual LAN that you want the partition to have an ethernet adapter connected to. Typically, the Virtual LAN ID is the same LAN ID as the i5/OS partition that is going to be used to route network traffic to the partition. In our

example, we are specifying that the partition should have a connection on Virtual LAN 1. Once you have completed defining the fields, click **OK**.

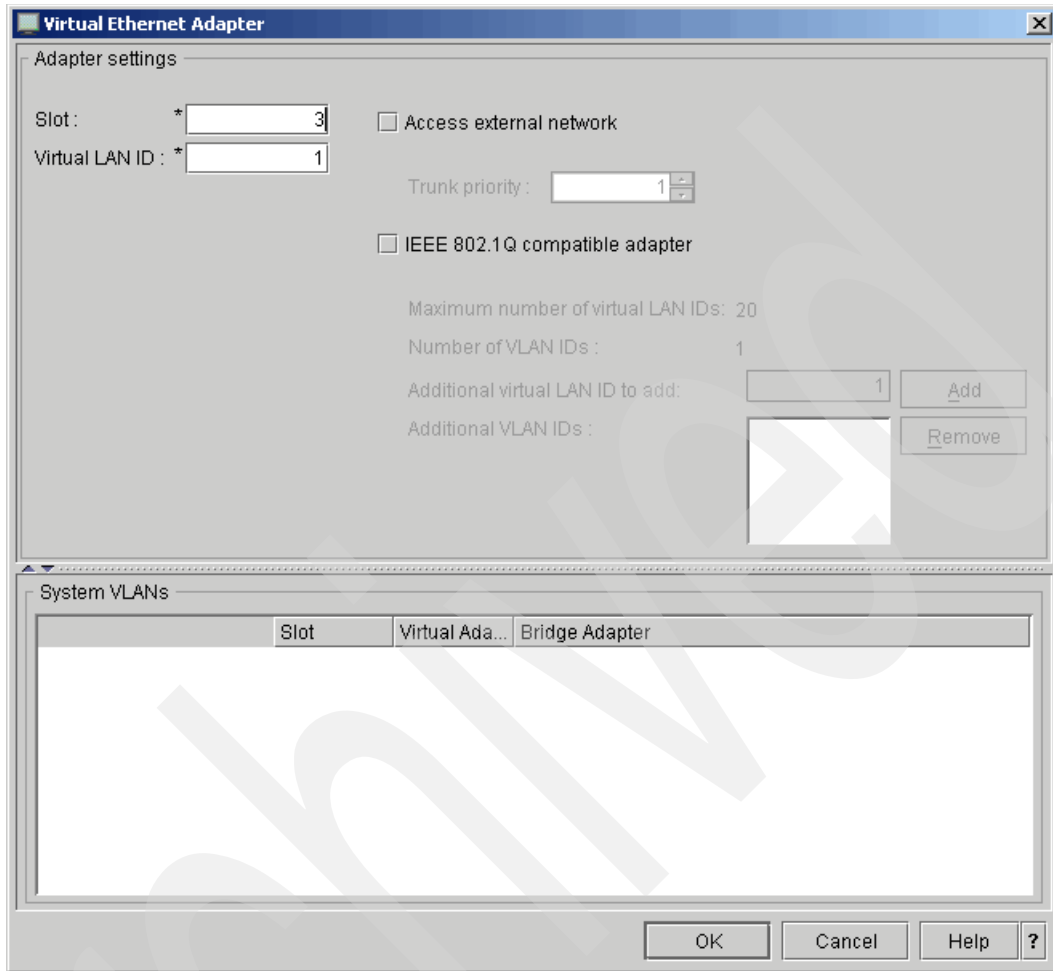


Figure 3-17 Creating a new Partition - Virtual Ethernet Adapter

16. Once the Virtual Ethernet Adapter has been created, click **Next** on the Create Logical Partition Profile - Virtual Adapters.

Now that the virtual I/O adapters have been created in the Linux partition, the virtual I/O diagram can be updated (as shown in Figure 3-18).

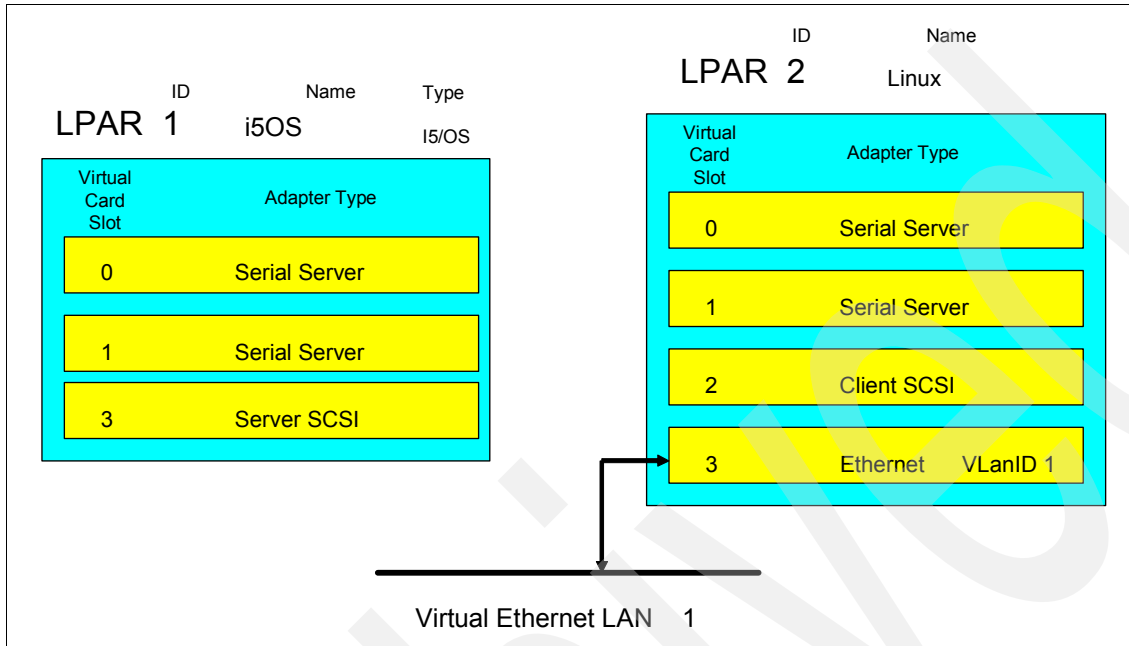


Figure 3-18 Virtual I/O Diagram

While the Server SCSI adapter has not yet been created, it is shown in the diagram as a reminder of the slot definition used in the virtual client SCSI adapter created in the Linux partition. Later, when the Virtual SCSI server adapter is created in the i5/OS partition, the diagram will be updated again to show the connection between the two adapters.

17. A Power Controlling Partition is a partition that is allowed to make power requests (activate, shutdown) for a partition via the Hypervisor. Figure 3-19 shows the definition screen for the Power Control Partition. Since the I/O for this partition is being hosted by another i5/OS partition, it is typical to have the i5/OS partition start the Linux partition

through a vary on of the Network Server Description (NWSD). In our example, we select the name of the i5/OS hosting partition (V5R4), click Add and then click Next.

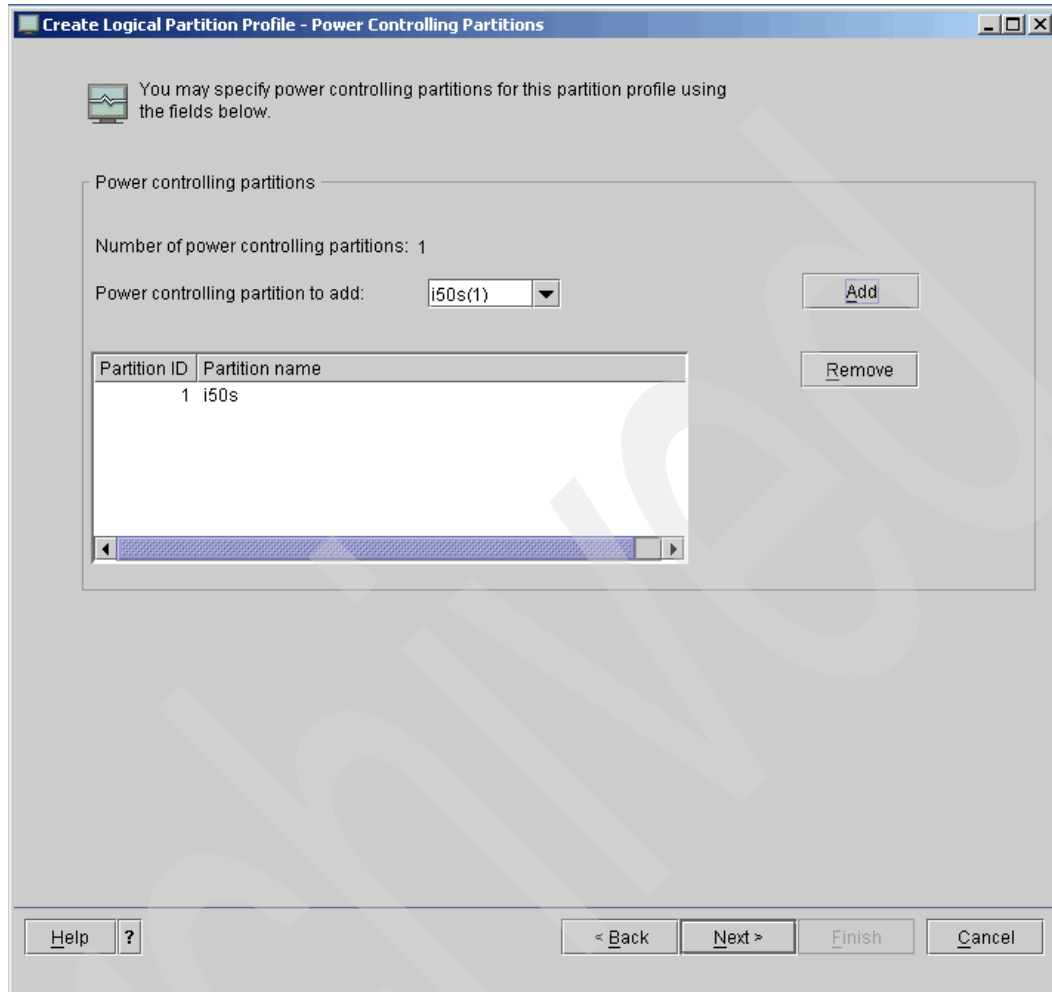


Figure 3-19 Create a new Partition - Power Controlling Partitions

18. Optional settings, such as boot mode, connection monitoring, and automatic activation of the partition when the Managed System starts are shown in Figure 3-20. For Linux partitions that have their I/O hosted by another partition the settings should be left at

Normal for boot mode, and the optional settings for connection monitoring and automatic start should be left as unset. Click **Next**.

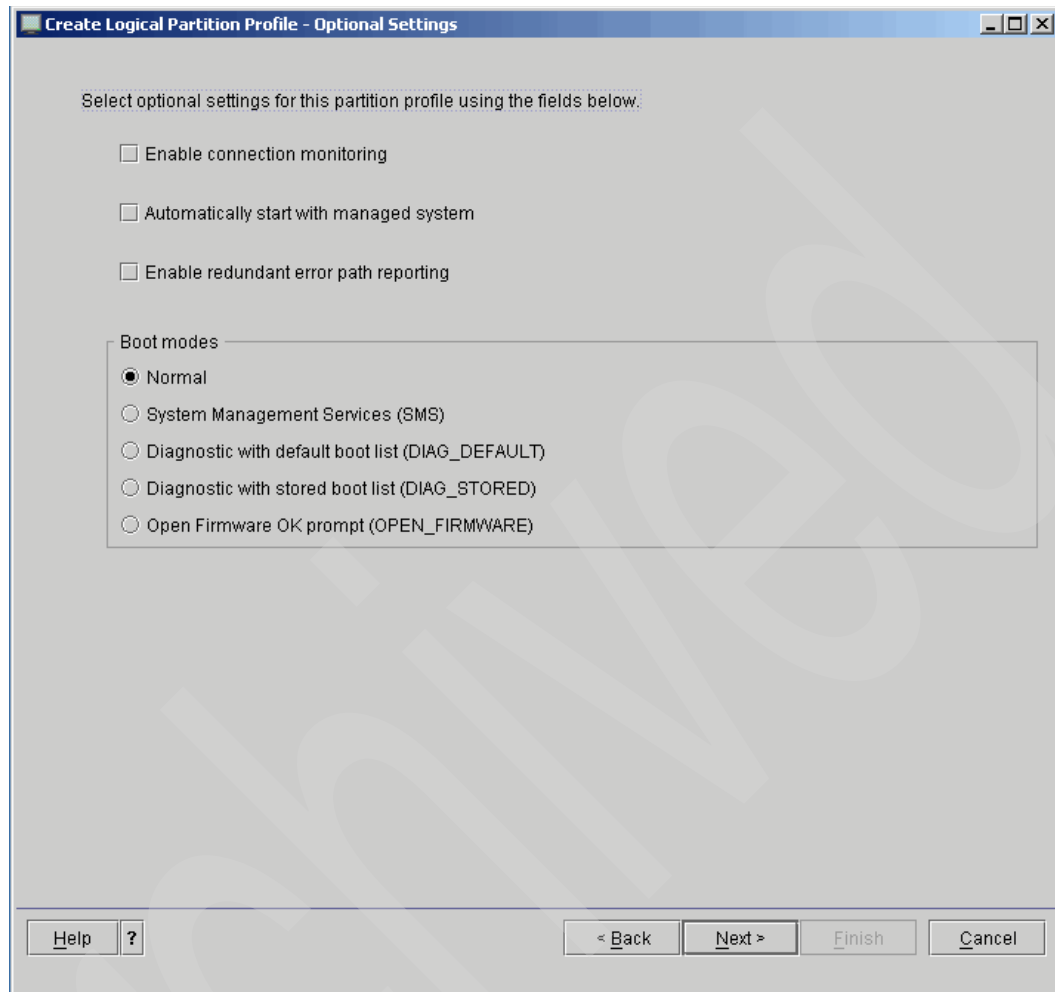


Figure 3-20 Create a new Partition - Optional Settings

19. At this point the partition definition is complete the a summary of the definition is displayed as shown in Figure 3-21.

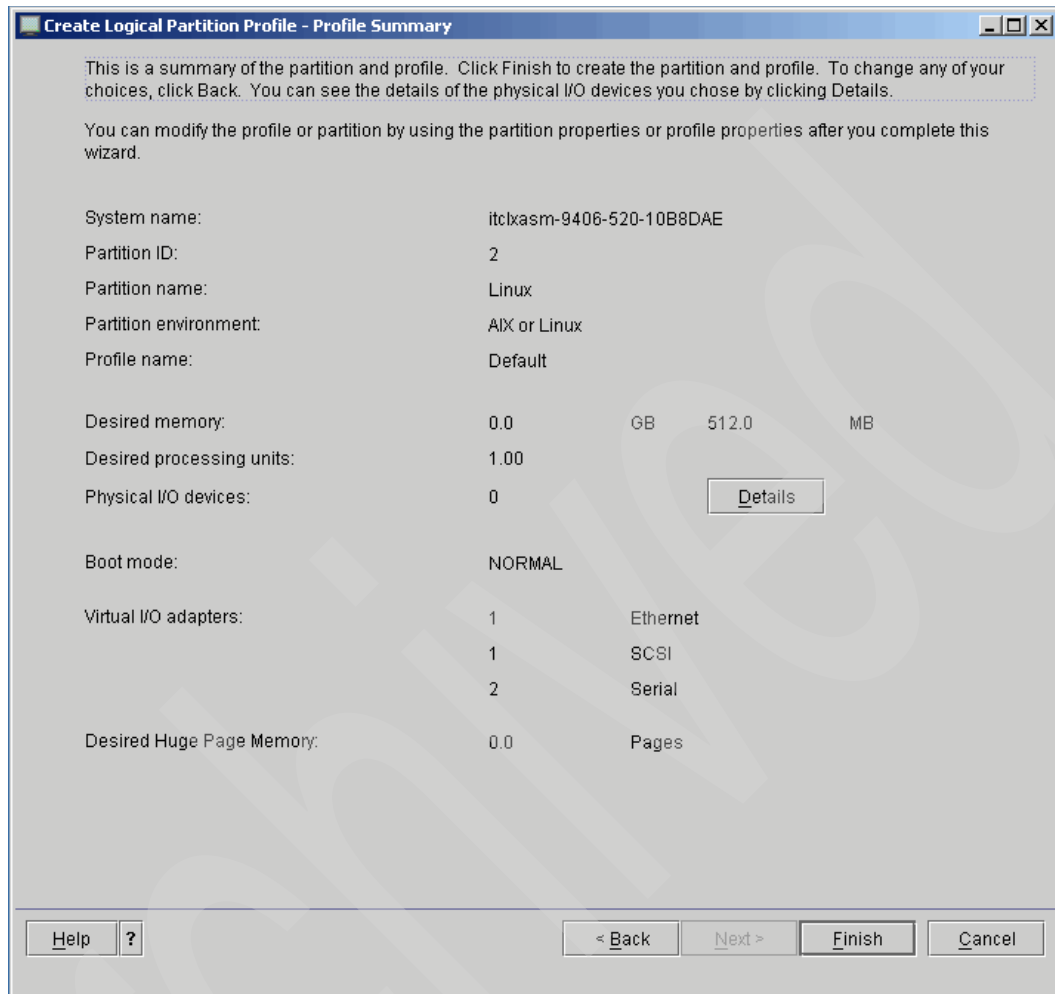


Figure 3-21 Partition Profile Summary

20. Click **Finish** to cause the partition to be created. A working dialog will be displayed as shown in Figure 3-22.




Figure 3-22 Working dialog

Adding new virtual adapters to hosting i5/OS partition

Once the Linux partition has been created, virtual adapters need to be created in the i5/OS partition to support virtual I/O and optionally virtual console. The virtual adapters will need to be created in the partition profile of the i5/OS partition and will also need to be added dynamically (via Dynamic LPAR) to the running state of the i5/OS partition.

Note: As indicated above, the adapters to support virtual I/O will need to be added twice. The adapter is added to the partition profile of the i5/OS partition to ensure that the change will be permanent across IPLs of i5/OS and deactivations or reactivations of the Logical Partition. Dynamic LPAR is used to add the virtual I/O adapter to the running state of the partition to make the adapter available for immediate use.

1. From the HMC, right-click the profile of the i5/OS partition and select **properties**. The properties window for the logical partition will displayed as shown in Figure 3-23.

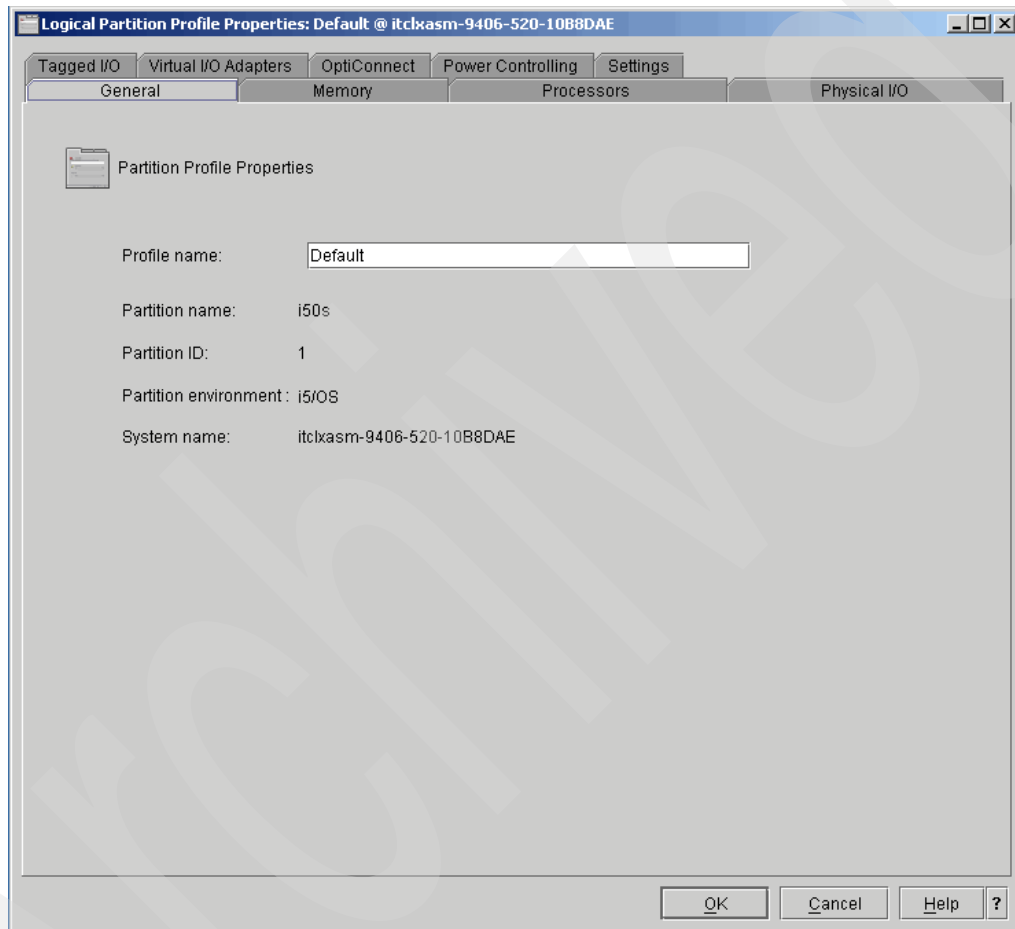


Figure 3-23 Hosting Partition Profile Properties

- Click the Virtual I/O Adapters tab to display the virtual I/O adapters that are currently defined in the partition profile. A list of adapters will be displayed as shown in Figure 3-24.

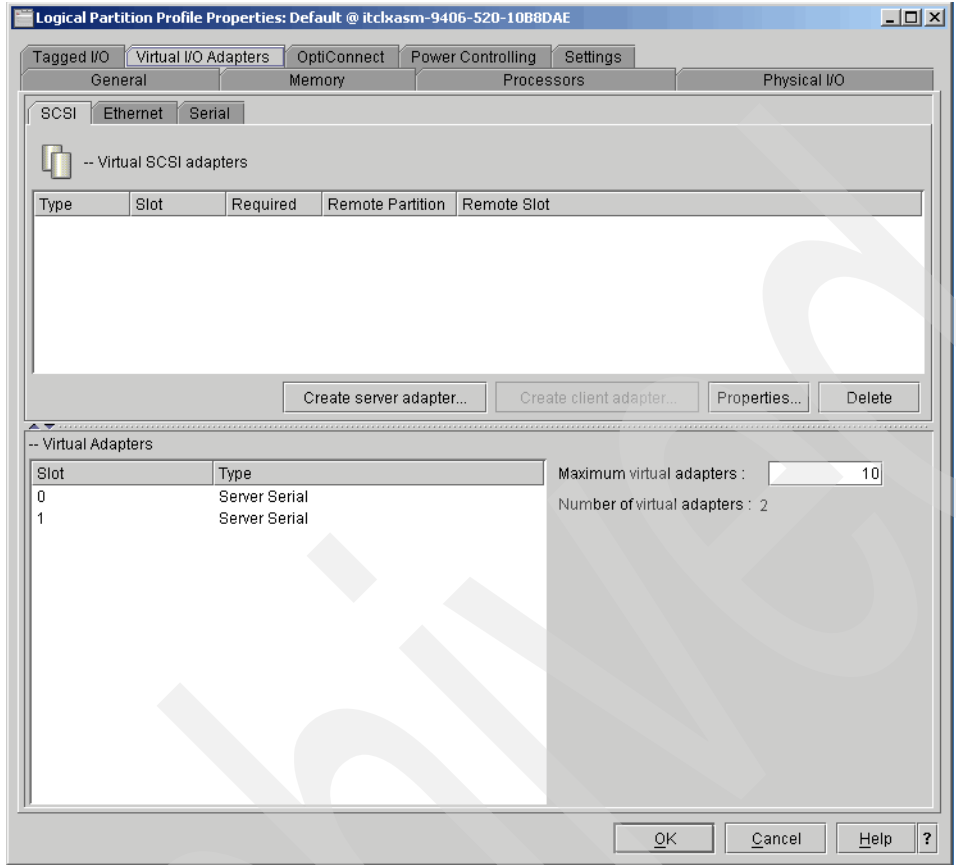


Figure 3-24 Hosting Partition Properties -- Virtual I/O Adapters

- To create the virtual SCSI Server adapter that will provide the I/O resource to the Linux partition, ensure that the SCASI tab is selected, and then select the **Create server adapter** button.
- In the Virtual SCSI Server Adapter Properties window as shown in Figure 3-25, complete the following fields:
 - Server slot:** This is the slot number for the adapter being created. This slot number must match the slot number that the Virtual Client Adapter in the Linux partition is pointing to.
 - Slot connection settings:** Only the selected client partition can connect
 - Client partition:** The name of the Linux partition selected from the pull-down list
 - Client partition slot:** The slot number of the Virtual SCSI Client adapter in the Linux partition.

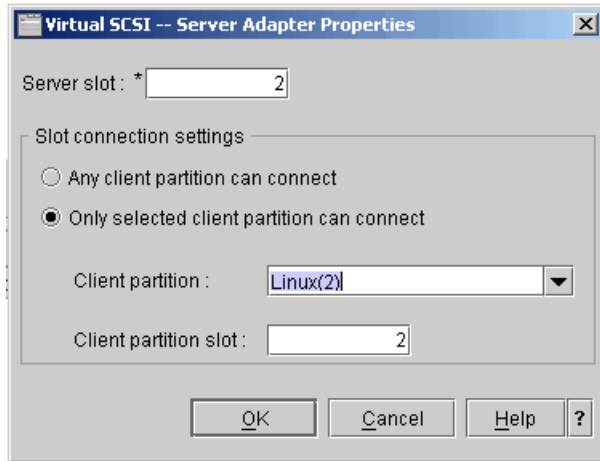


Figure 3-25 Hosting Partition Properties - Create Virtual SCSI Server Adapter

To help clarify, the Virtual SCSI Client adapter in the Linux partition needs to point to the Virtual SCSI Server adapter in the Hosting i5/OS partition and likewise, the Virtual SCSI Server adapter in the Hosting i5/OS partition needs to point to the Virtual SCSI client adapter in the Linux partition.

As an example, look at Figure 3-26 which shows the properties of the Virtual SCSI server adapter from the hosting i5/OS partition and mappings to the properties of the Virtual SCSI client adapter in the Linux partition.

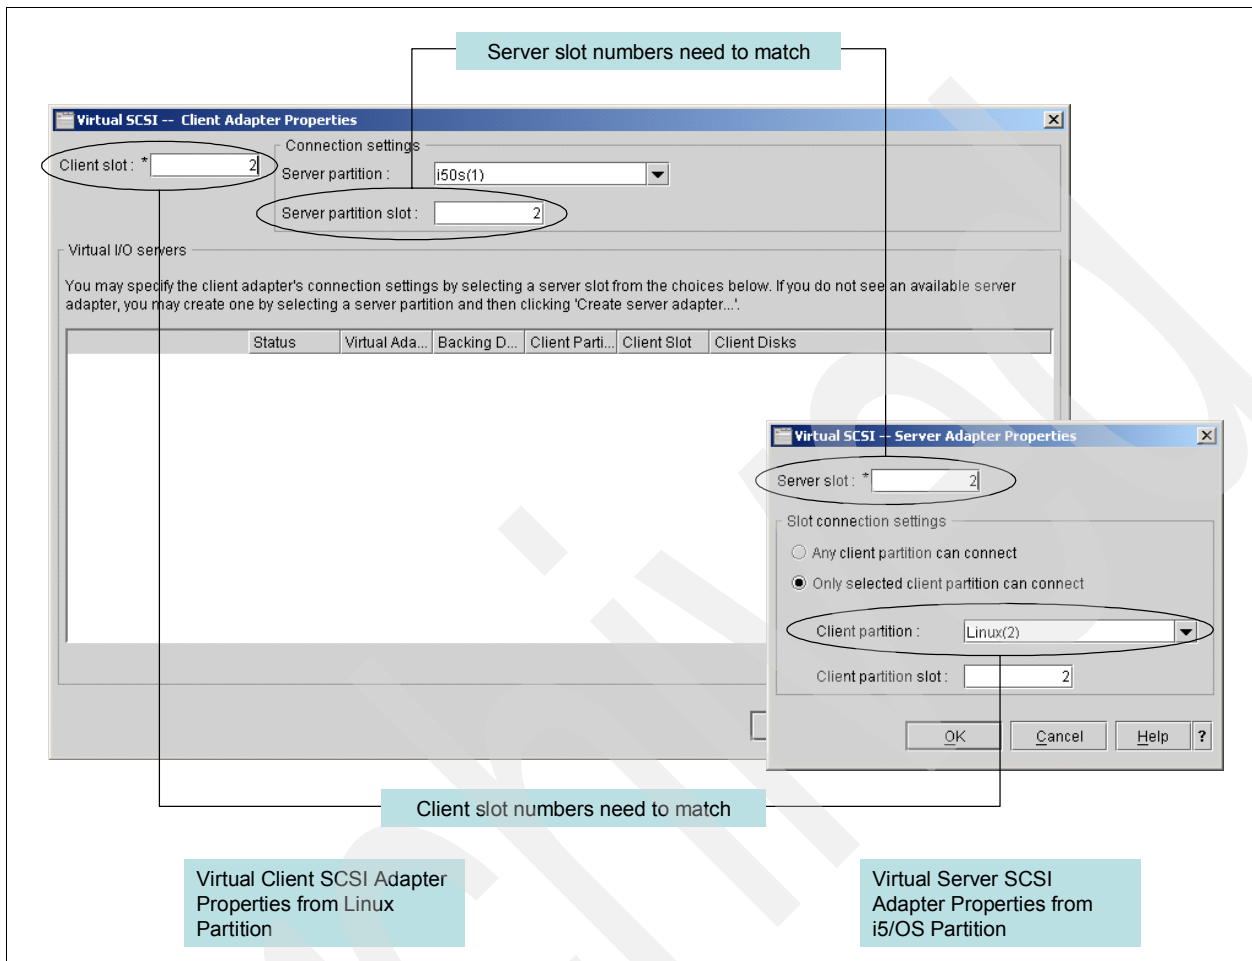


Figure 3-26 Virtual SCSI Adapters -- Mapping Example

- Once the fields of the Virtual SCSI - Server Adapter Properties have been completed, click **OK**. In the Logical Partition Profile Properties window, click **OK**.

At this point the Virtual SCSI Server adapter has been added to the profile of the hosting i5/OS partition which means that the next time the i5/OS partition is started from the HMC the adapter will be available. To add the adapter to the i5/OS partition without a restart of i5/OS will require that the adapter be added to the running state of the partition through the use of Dynamic LPAR.

Note: The Virtual SCSI Server adapter should be added twice to ensure that it is available to support virtual I/O to Linux without a restart of i5/OS as well as to ensure that the adapter is available across restarts of i5/OS.

- To add the Virtual SCSI Server adapter to the hosting i5/OS partition, right-click the i5/OS partition (not the partition profile) and then select **Dynamic Logical Partitioning** → **Virtual Adapter Resources** → **Add / Remove** as shown in Figure 3-27.

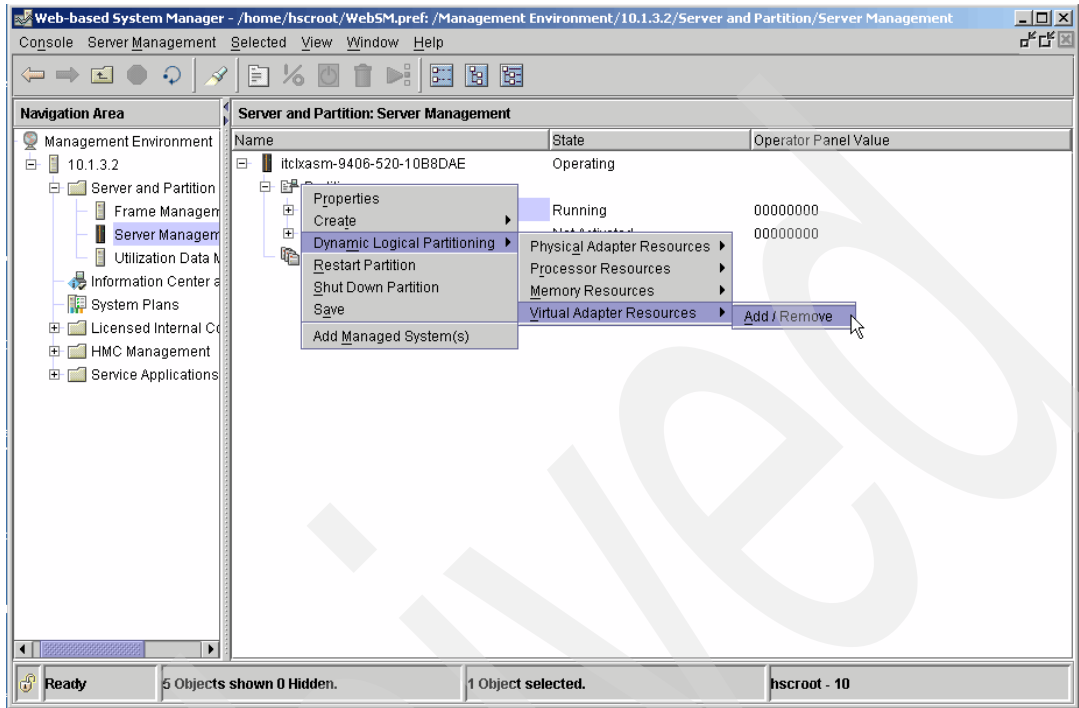


Figure 3-27 Hosting i5/OS Partition - Start Dynamic Logical Partitioning

7. In the Dynamic Logical Partitioning - Virtual Adapters window (shown in Figure 3-28), ensure that the SCSI tab is selected and then click **Create server adapter**.

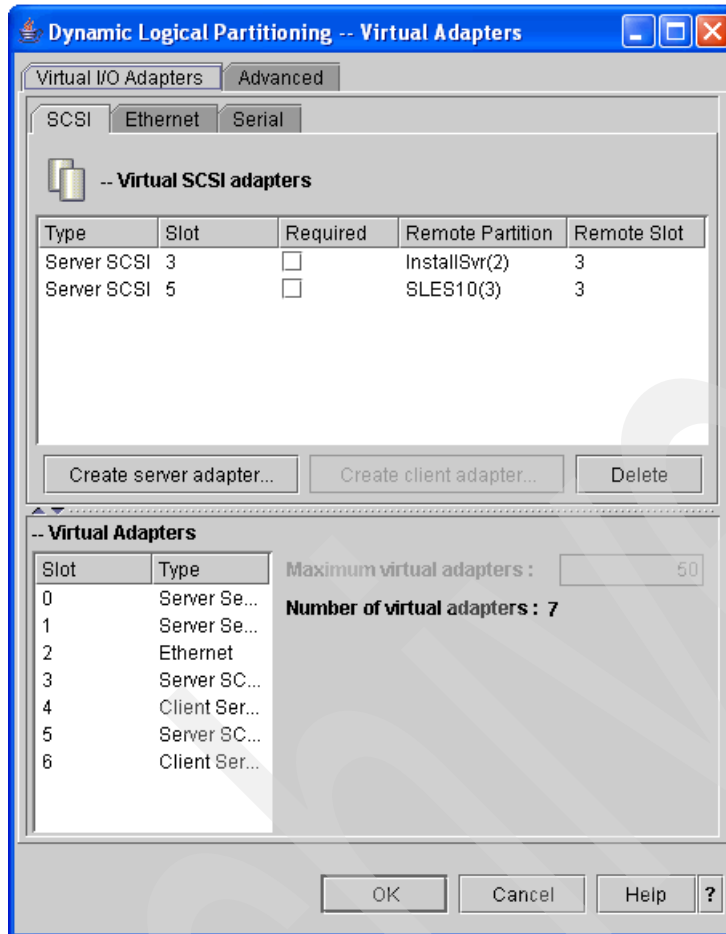


Figure 3-28 Hosting Partition - Dynamic Logical Partitioning - Virtual Adapters

8. In the adapter properties window shown in Figure 3-29, enter the following information:
 - **Server slot:** This is the slot number for the adapter being created. This slot number must match the slot number that the Virtual Client Adapter in the Linux partition is pointing to.
 - **Slot connection settings:** Only selected client partition can connect.
 - **Client partition:** The name of the Linux partition selected from the pull-down list.
 - **Client partition slot:** The slot number of the Virtual SCSI Client adapter in the Linux partition.

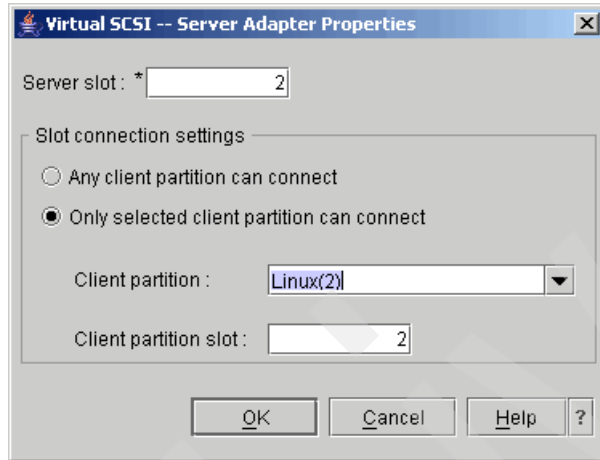


Figure 3-29 Hosting Partition - Dynamic Logical Partition - Server SCSI Partition Properties

Note: The information entered for the adapter should match the information defined for the Virtual Server SCSI adapter in the partition profile of the hosting i5/OS partition.

9. Once the adapter properties have been defined, click **OK**. In the Dynamic Logical Partition - Virtual Adapters window, click **OK**.
 - Now that the virtual SCSI server adapter has been added to the i5/OS partition, the virtual I/O configuration diagram is updated to reflect the attachment of two SCSI adapters as shown in Figure 3-30.

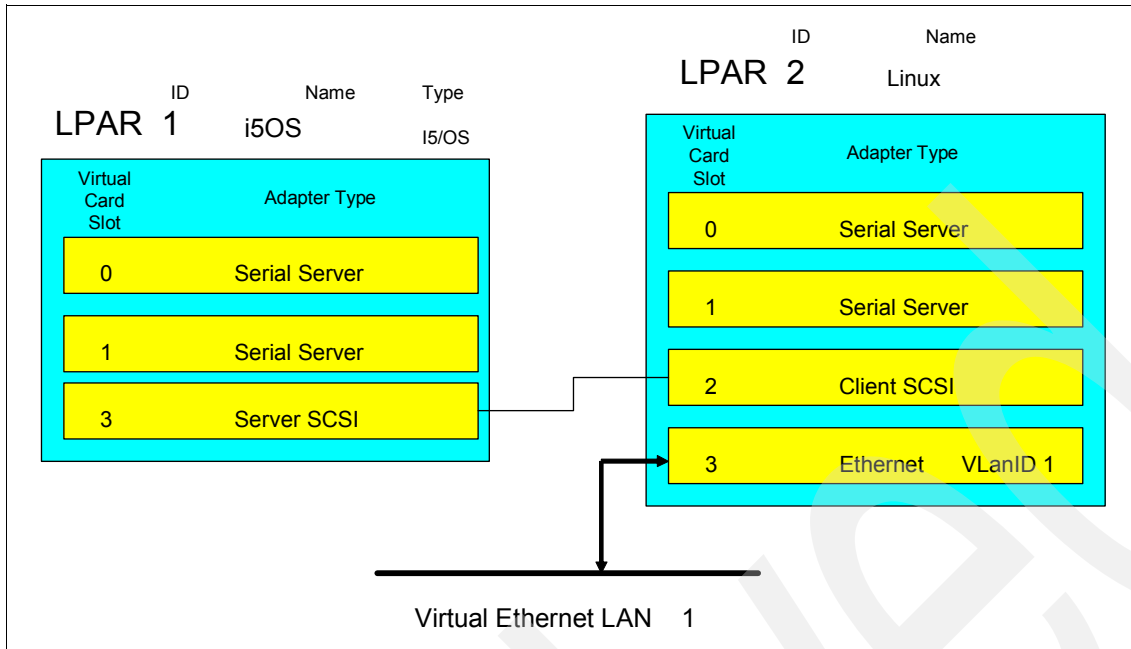


Figure 3-30 Virtual I/O Configuration

Registering the Linux partition definition

At this point, the properties of the logical partition for Linux have been defined and the properties of the hosting i5/OS partition have been updated to support virtual I/O for the Linux partition. The last step in the partition definition is registering the partition definition with the firmware of the System i platform which requires the following steps:

1. From the HMC, right-click the **Linux** partition and then click **Activate** as shown in Figure 3-31.

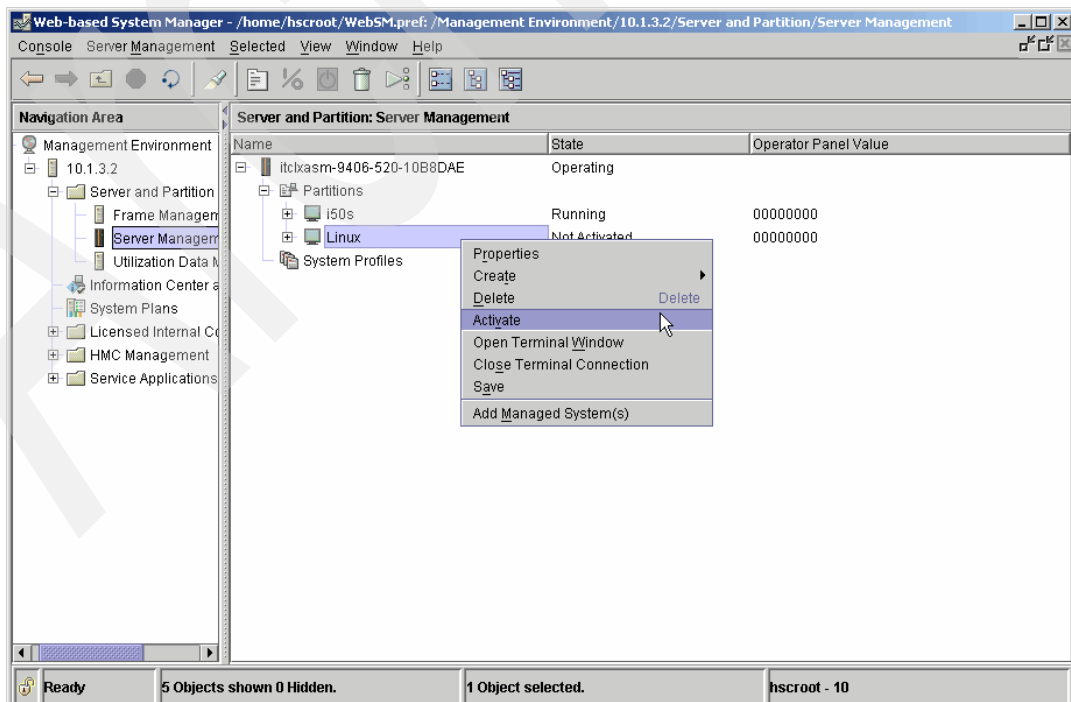


Figure 3-31 Logical partition activation

2. From the Activate Logical Partition window shown in Figure 3-32, select **Open a terminal window or console session** and then click **OK**.

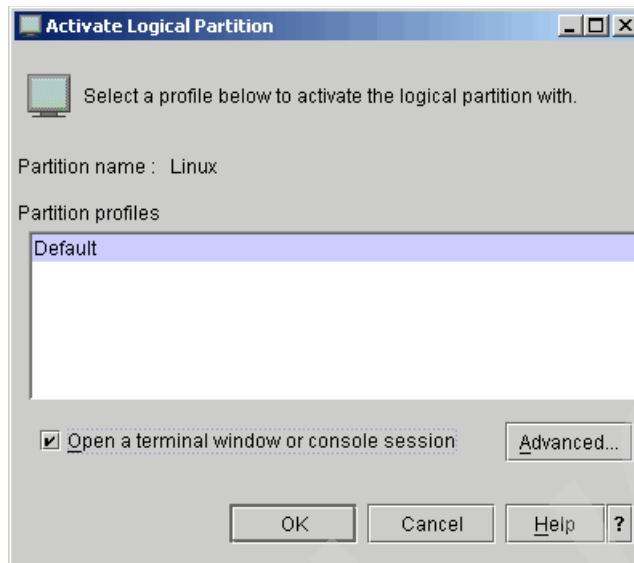


Figure 3-32 Activate Logical Partition

A message may be displayed in the virtual console window prompting for console selection. If this message is displayed, press the **0** key.

At this point, the resources defined in the partition profile will be reserved by the hypervisor and the partition will be started. Since there is no I/O defined yet, the partition will eventually boot into the firmware of the partition.

3. Now that the partition has been registered with the firmware of the System i platform it can be shut back down. From the HMC, right-click the Linux partition and select **Shut Down Partition**.

4. A window will be displayed prompting for confirmation of the shutdown as shown in Figure 3-33. Since there is currently no operating system installed in the partition, click **Immediate** and then click **OK**.

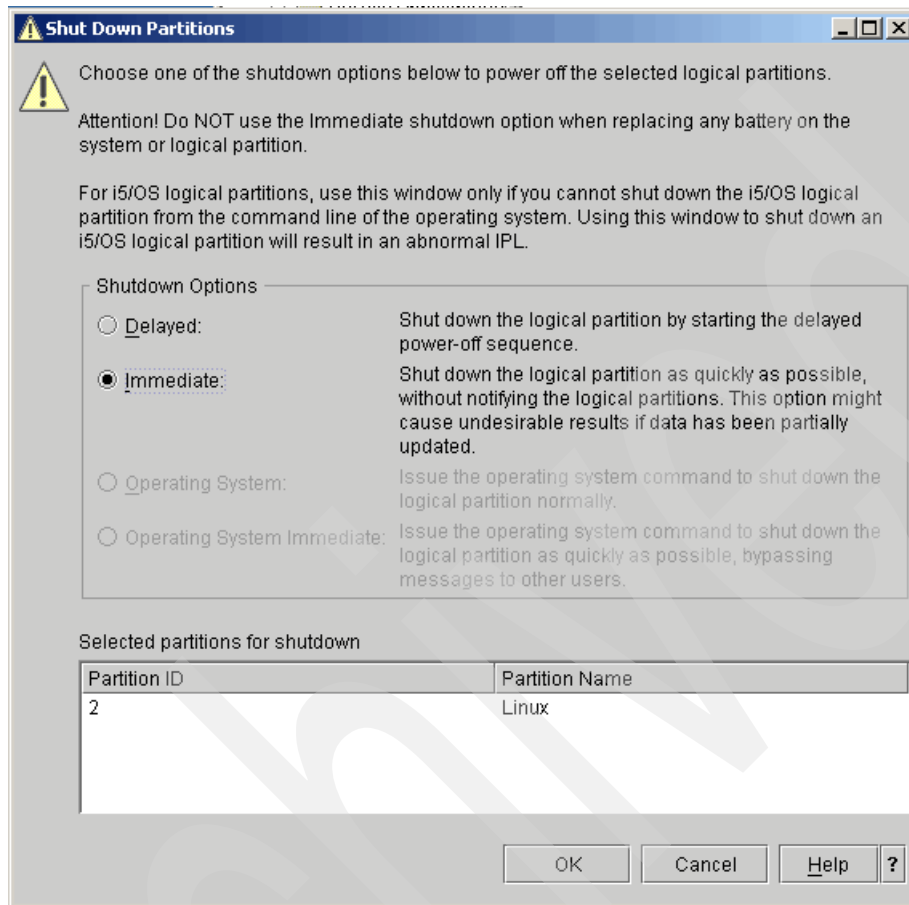


Figure 3-33 Shut down partitions option screen

5. Since an *immediate* shutdown was requested, a confirmation screen will be displayed as shown in Figure 3-34. Click **Yes**.

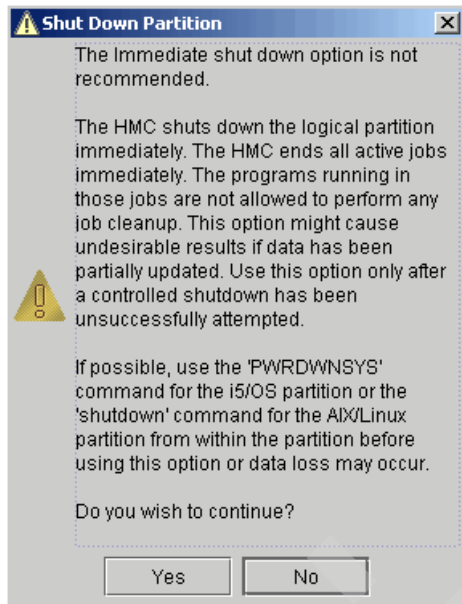


Figure 3-34 Partition shutdown: immediate shutdown confirmation

6. Depending on the version of HMC, an additional warning dialog box may be displayed concerning replacement of a cache battery (Figure 3-35). Click **No** to indicate that a cache battery is not being replaced. At this point the partition will be shut down.




Figure 3-35 Cache Battery Replacement dialog

Note: If the Virtual Terminal for the Linux partition is still displayed, it can be closed by clicking the X in the upper right-hand corner.

Definition of the logical partition for Linux is now complete. Continue to Section 4.1, “Supporting virtual I/O” on page 118 for definition of the i5/OS components to support virtual I/O

3.3 Linux partition with all virtual resources via Virtual Partition Manager

Virtual Partition Manager provides the capability to create and manage Linux partitions without the use of the HMC. Virtual Partition Manager supports the needs of small and medium business customers that want to add simple Linux workloads to their System i server. Virtual Partition Manager is included with i5/OS V5R3.

You can find detailed information about Virtual Partition Manager in *Virtual Partition Manager A Guide to Planning and Implementation*, REDP-4013. Information presented in this section is limited to an example of using Virtual Partition Manager to establish an initial Linux partition. This example assumes that all system resources are currently owned by the i5/OS partition.

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 System i platform, 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 System i 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 enable 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.

3.3.1 Removing resources from i5/OS

Again, we assume all system resources are currently owned by the i5/OS partition. Therefore, you need to remove the resources to be used by your Linux partition from the i5/OS partition first by performing the following steps:

1. From the Main menu or the i5/OS command line, start the System Service Tools by typing STRSST as shown in Figure 3-36.

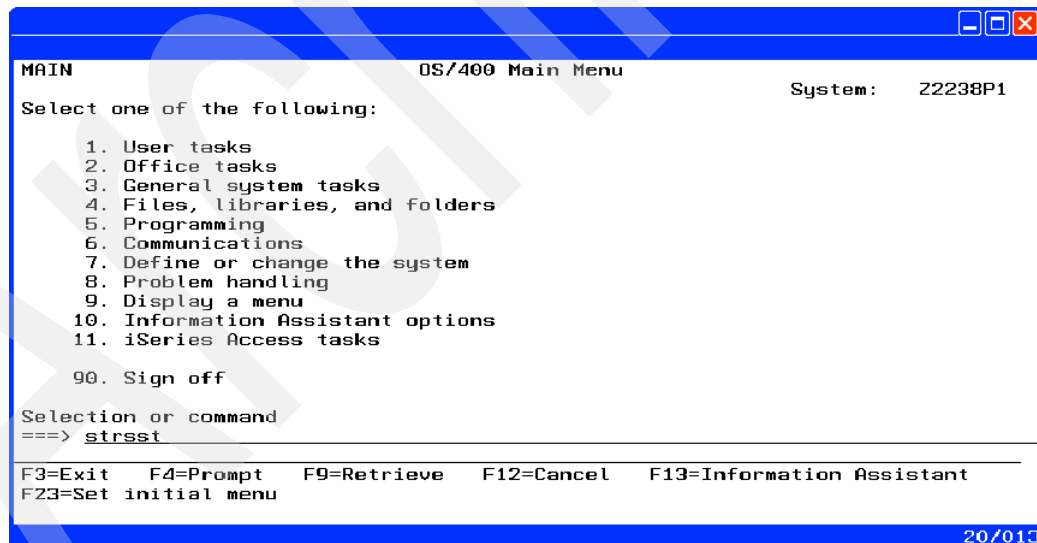


Figure 3-36 Main menu

2. Enter your user ID and a password as seen in Figure 3-37. 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.

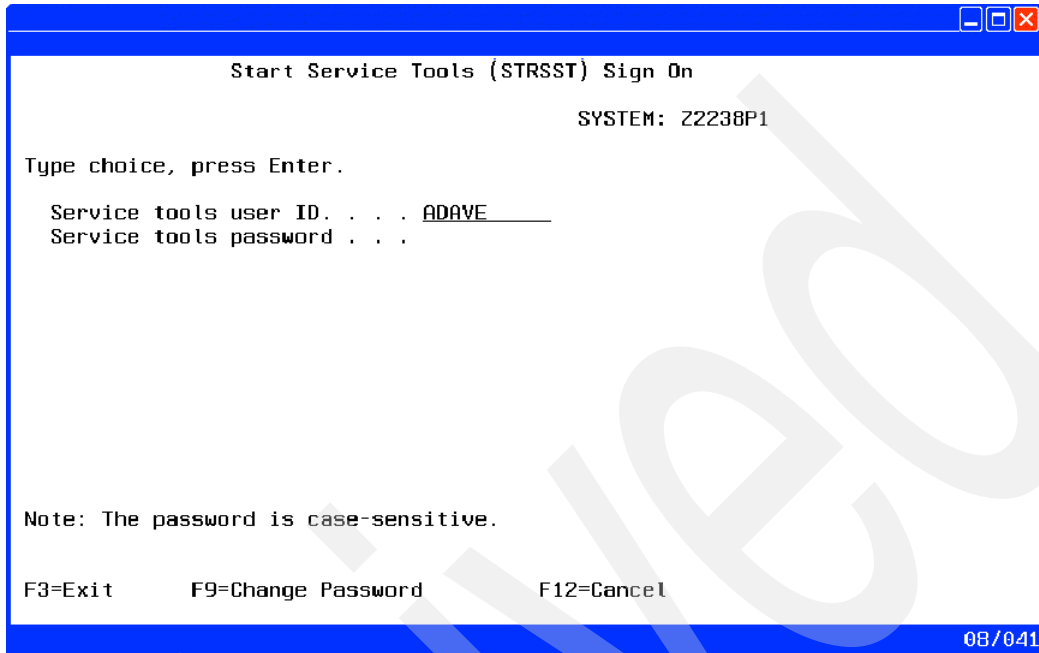


Figure 3-37 Start Service Tools sign-on screen

3. From the System Service Tools (SST) menu as seen in Figure 3-38, select option **5 Work with system partition**.

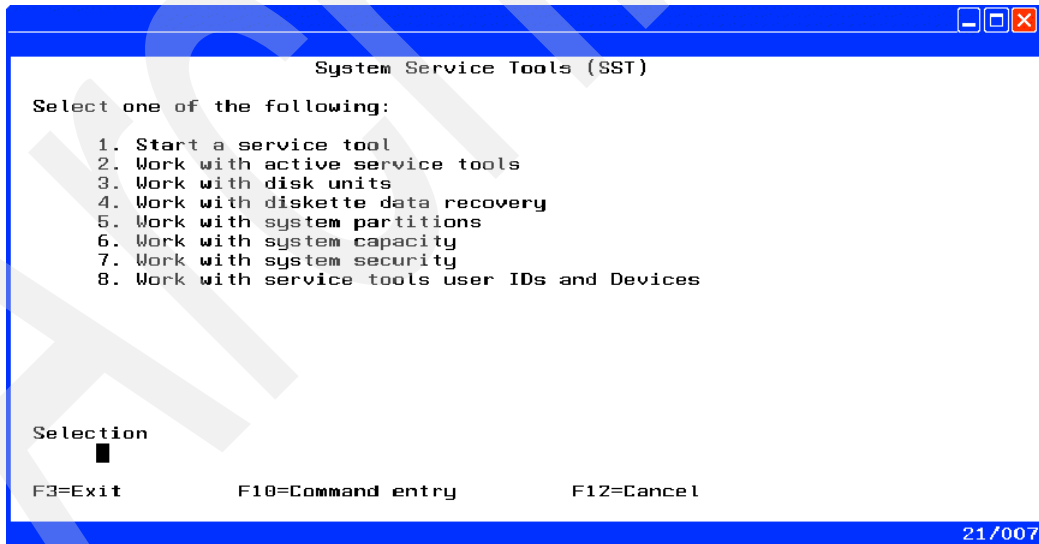


Figure 3-38 System Service Tools

A message appears as shown in Figure 3-39. The 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 Virtual Partition Manager, an error dialog box appears indicating that the tool is already in use.

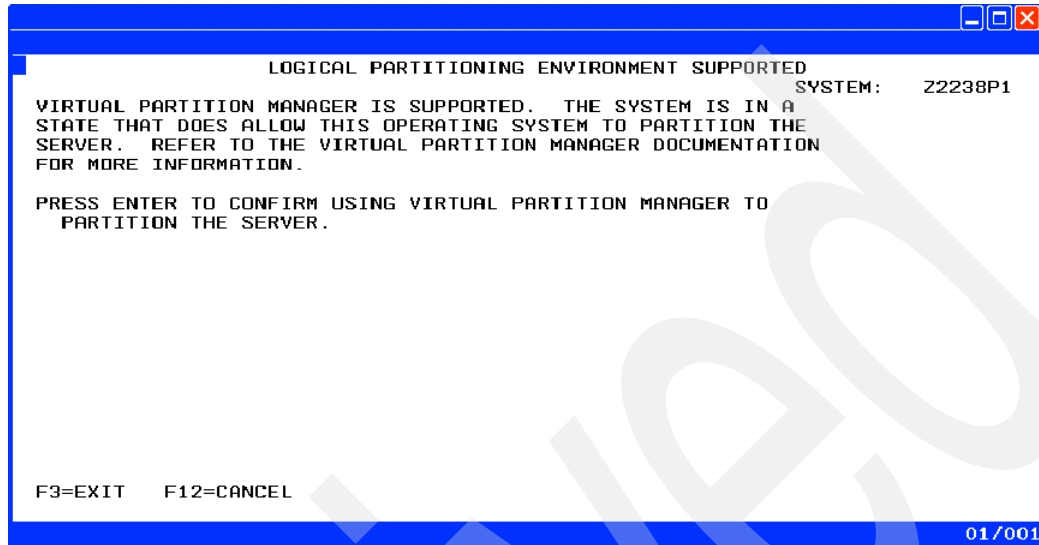


Figure 3-39 Initial information display

4. Press **Enter** at the message. The logical partition management tasks appear, as seen in Figure 3-40.

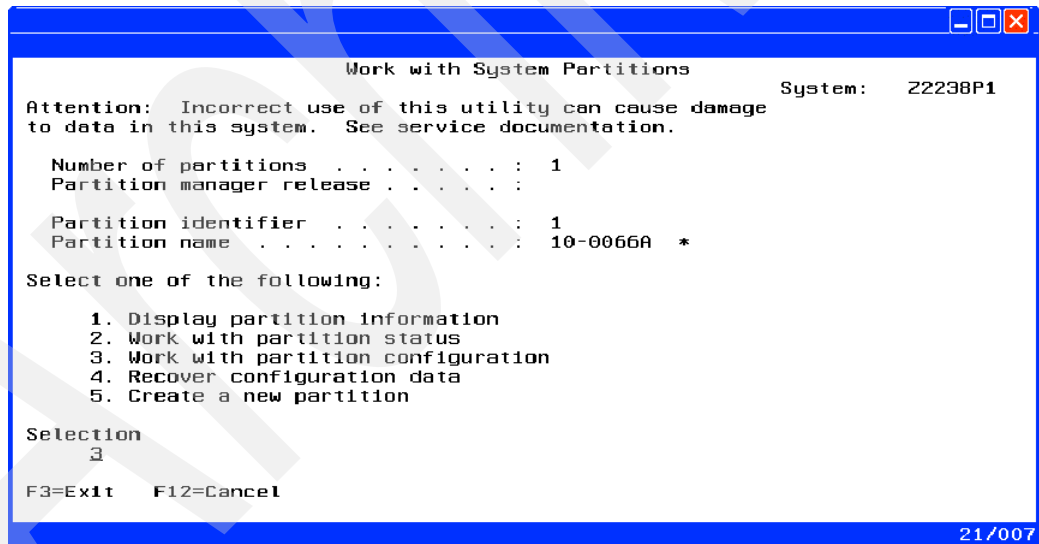


Figure 3-40 Work with System Partitions

5. Select option **3 Work with partition configuration**, as shown in Figure 3-40.

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.

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 to Change partition configuration for your i5/OS instance as shown in Figure 3-41.

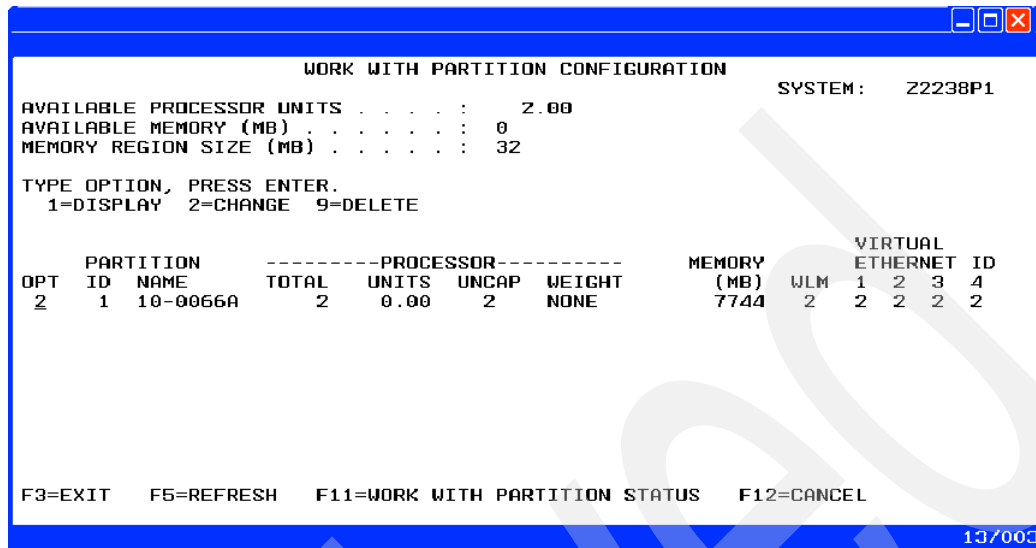


Figure 3-41 Work with Partition Configuration

With a new system, one i5/OS partition is defined and the Change Partition Configuration display shows the defaults as shown in Figure 3-42. This is where we remove resources from your i5/OS instance so that additional logical partitions for Linux.

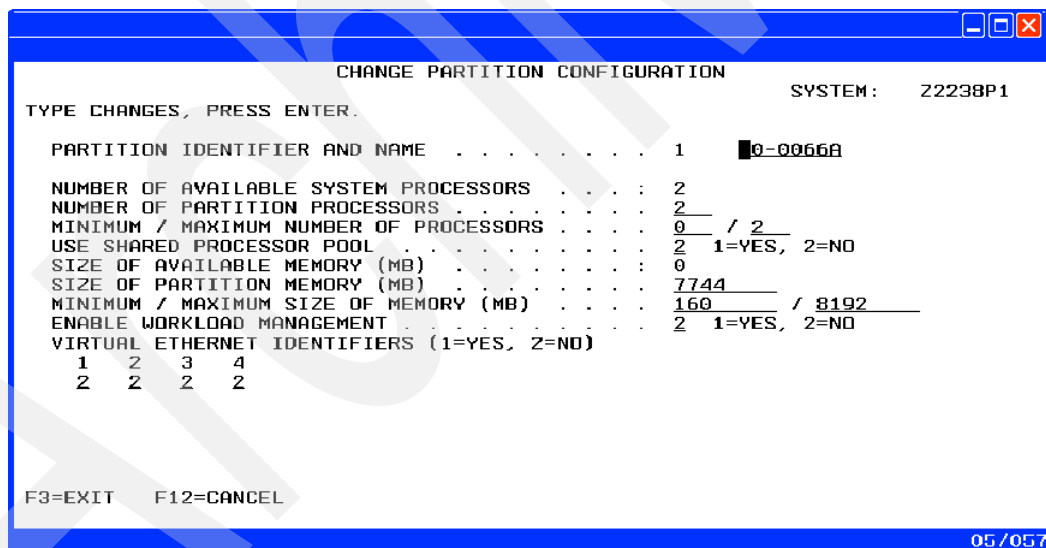


Figure 3-42 Change Partition Configuration

7. Make several changes here, based on the resources you want to set for the i5/OS partition. You have to assign CPU and memory allocations. We examine all of the changes, step-by-step, as highlighted in Figure 3-43.

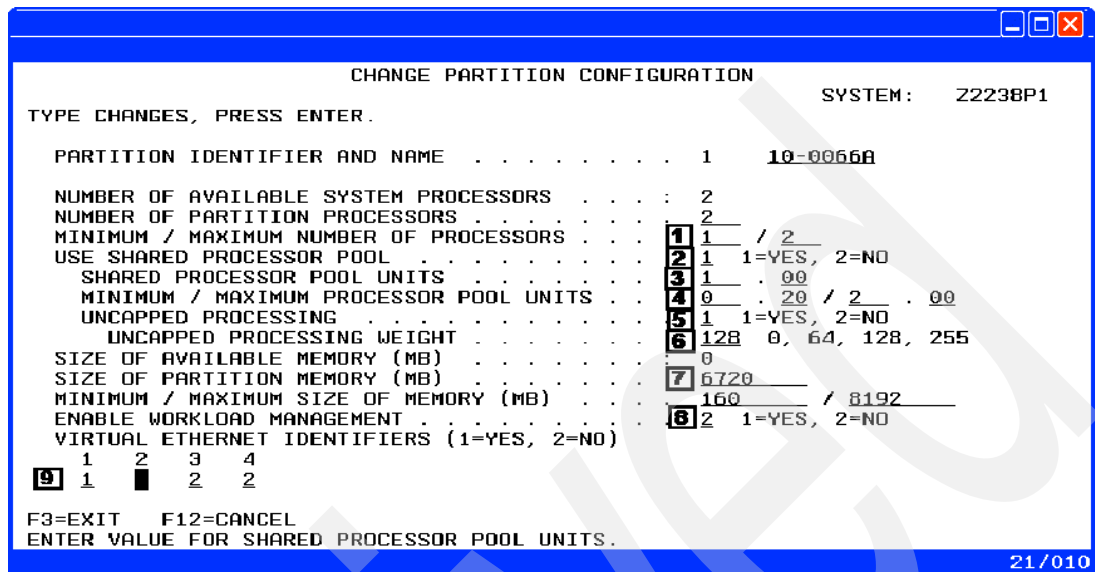


Figure 3-43 Change partition configuration options

- **1 Minimum / Maximum Number of Processors:** Enter the value for the minimum number of processors you anticipate you will require for your i5/OS environment, and the maximum number of processors available for i5/OS. In our example environment, we had two 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.
- **2 Use Shared Processor Pool:** The shared processor pool is a group of physical processors that provide processing capacity that can be shared amount multiple logical partitions. The shared processor pool allows you to assign partition processors to a logical partition. In our example, since we are going to create four additional Linux partitions, we do not have sufficient resources to allocate a full processor for every partition. Another advantage for changing the partition configuration to use shared processor pool is that the partition has the capability to use processor resources unused by other partitions.

Note: When the current value of 2 is changed to 1, indicating shared processor pool is desired, the system prompts for new shared processor pool values *once the Enter key is pressed*.

- **3 Shared Processor Pool Units:** Specifies the total number of processing units that are available to the partition after the resources are removed. In this example, the i5/OS partition is left with 1.00 processing units, or a full processor after we remove the CPU resources.
- **4 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 2.00 processing units, or 2 processors.
- **5 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 is not assigned excess shared processor cycles and is a capped partition.

A capped partition indicates that the logical partition will never exceed its assigned processing capacity. You can use the capped mode if you know that a software application would never require more than a certain amount of processing power. Any unused processing resources are 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 process or pool has any unused processing power.

Note: Considering that Virtual Partition Manager does not support dynamic movement of resources, you can define your partitions as uncapped partitions to automatically assign unused processing units to a partition that needs the resources.

- **6 Uncapped Processing Weight:** With Virtual Partition Manager, since you can only create four additional Linux partitions, the 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 3 received. A value of zero is required to be set for partitions that are capped.
- **7 Size of Partition Memory:** Linux partitions require a minimum of 256 megabytes. In this example, the value indicates the amount of main storage that remains with the i5/OS partition. Make the new value multiples of the value set as the LMB size during your initial setup using ASM. For example, you cannot set a value of 6700 because it gives an error message, line the one shown in Figure 3-44.

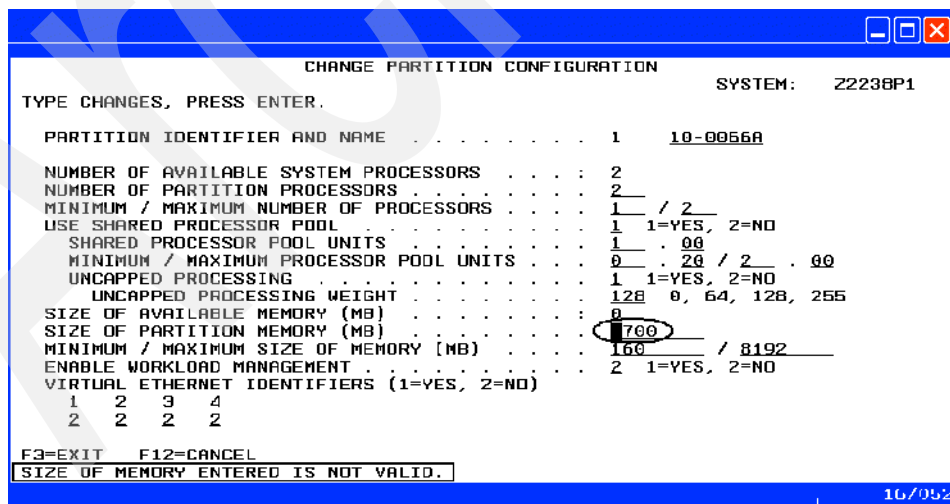


Figure 3-44 Change Partition Configuration Memory Error

- **8 Enable Workload Manager:** The default value for Virtual Partition Manager is set to 2=No, meaning that the partition is not allowed to use the future workload management

tool within the partition to automatically adjust resource assignments, such as the IBM Enterprise Workload Manager.

- **9 Virtual Ethernet Identifiers:** A value of 1 indicates you enabled 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 creates a virtual Ethernet adapter under i5/OS with a resource type of 268C and an associated resource name that is used during the setup of the 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 shown in Figure 3-45 appears.

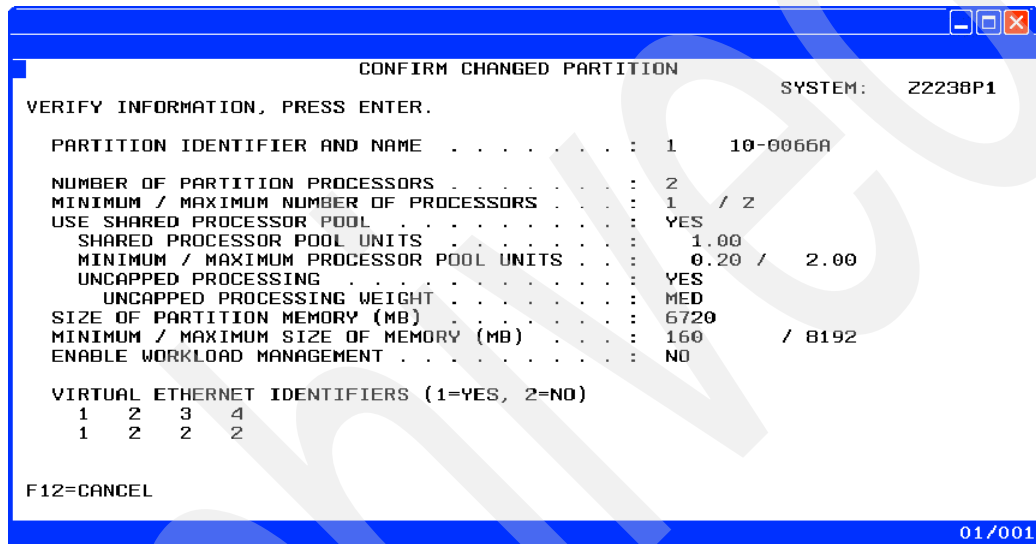


Figure 3-45 Confirm changed partition

9. Press **Enter** again on the Confirm Changed Partition screen as shown in Figure 3-46. This completes the changes required for the i5/OS partition configuration.

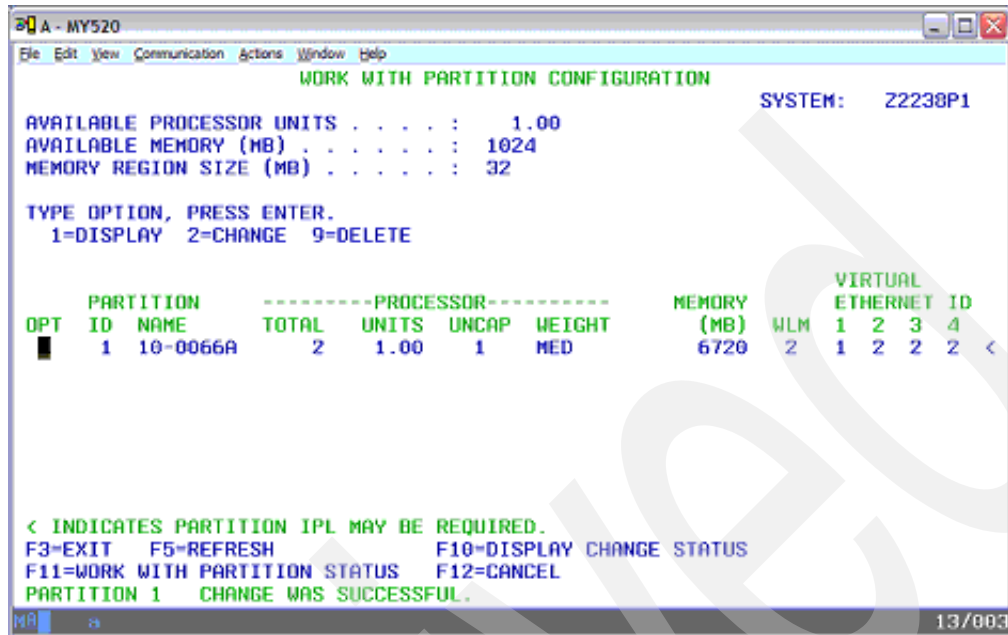


Figure 3-46 Partition change successful

Notice that the changes made to i5/OS resources require an IPL of the system. There is no need to perform this task at present, you can perform the IPL once you have defined 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 press **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 3-47 shows the available CPU and memory resources for creating new partitions. In the next section, we use these resources to define new Linux partitions.

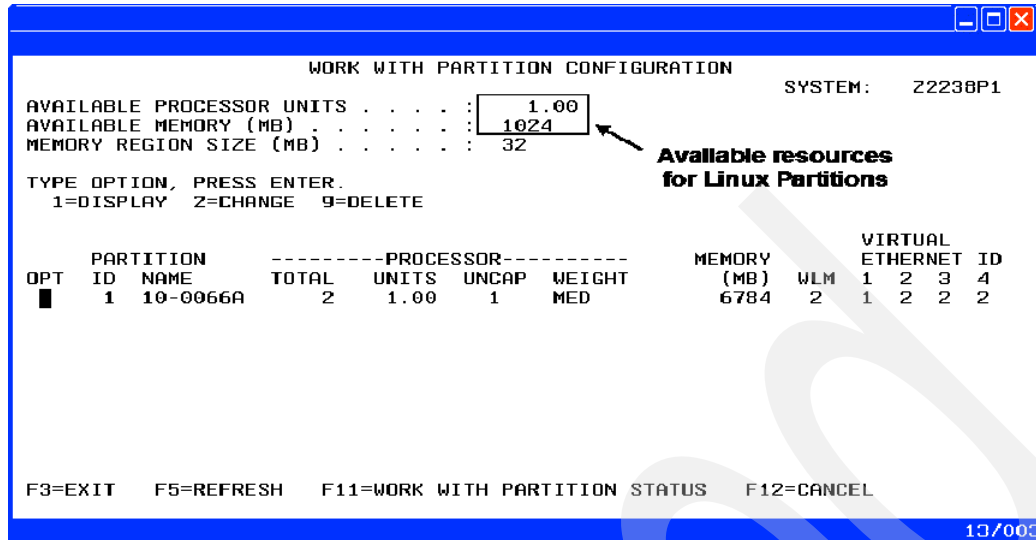


Figure 3-47 Available resources for Creating New Partitions

12. You are now ready to create Linux partitions using Virtual Partition Manager, and the resources that you removed from the 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 section walks through creation of the Linux partition.

3.3.2 Defining new Linux partitions

1. From the Work with System Partitions display, select option 5 to create a new partition as shown below in Figure 3-48.

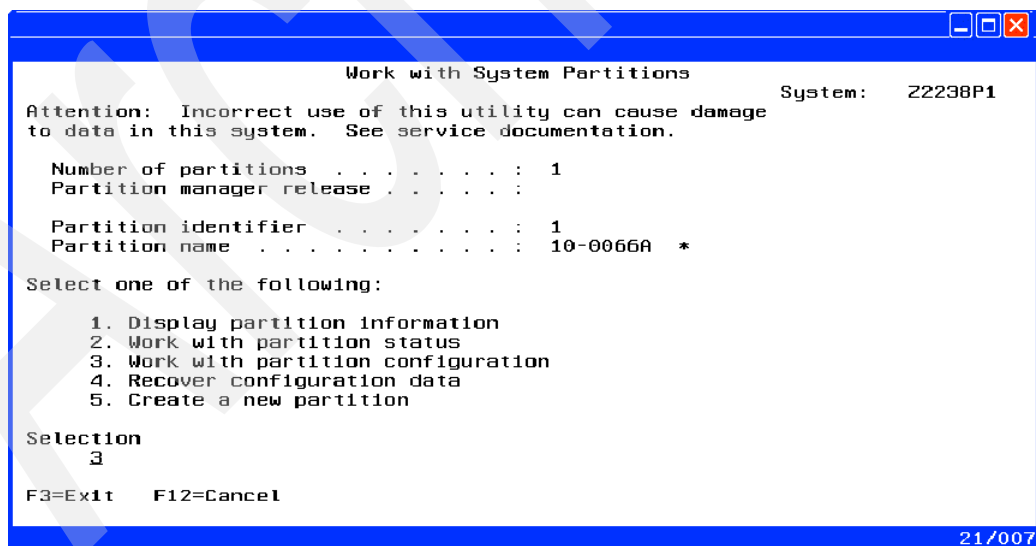


Figure 3-48 Work with system partitions

2. Assign the following values for creating the new Linux partition as shown in Figure 3-49:
 - **1 Partition Identifier and Name:** Enter the partition name for your Linux partition. You can also change the partition identifier if you choose to. In this example, the default given (next partition identifier number) by the system is selected.

- **2 Number of Partition Processors:** If you are defining a Linux partition with dedicated processors, then enter the value of the processors that are 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 also affects how you setup you Min/Max number of processors, and also how you setup your Min/Max shared processor pool units later. For example, if you set a value of 4, then set the maximum number of processors in the next field (as shown in 3 below) to 4. Similarly, you need at least 0.40 processing units specified in step 4.
- **3 Minimum / Maximum number of Processors:** Enter the value for minimum number of processors that you need for your Linux environment, and the maximum number of processors that are available for Linux.
- **4 Shared Processor Pool Units:** Assuming that you are going to use the shared processor pool, you now need to specify the total number of processing units available to the partition during startup. 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 with a dedicated CPU.

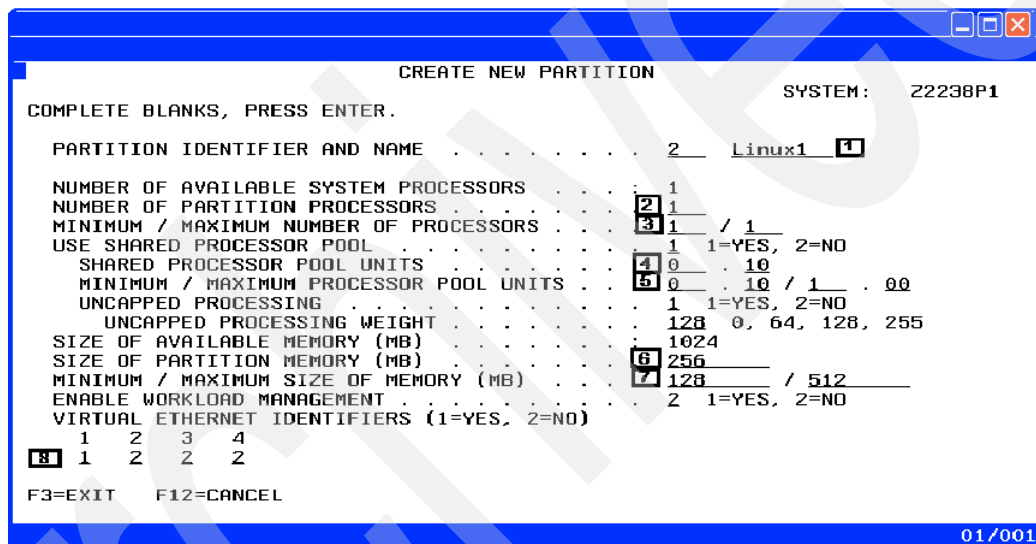


Figure 3-49 Create new Linux partition

- **5 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 that you want your partitions to utilize unused processing cycles
- **6 Uncapped processing:** You have the option to have your Linux partition shared capped, or shared uncapped. The shared processor concept is discussed in some detail in the planning section of this book. You can also reference the shared processor section in the IBM Information Center for more information about capped and uncapped processors on the Web at:

<http://www-1.ibm.com/servers/eserver/iseries/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, define your partitions as uncapped partitions to automatically assign unused processing units to a partition that needs the resources.

- **7 Size of Partition Memory:** The minimum memory requirement for Linux partition is 256 megabytes.
- **8 Virtual Ethernet Identifiers:** A value of 1 indicates that 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 four virtual Ethernet communications links.

Once these values are set, you get a confirmation display, as shown below in Figure 3-50.

3. Press **Enter**. The Work with Security Partitions display appears again. You can repeat the previous steps to define additional Linux partitions, if necessary.

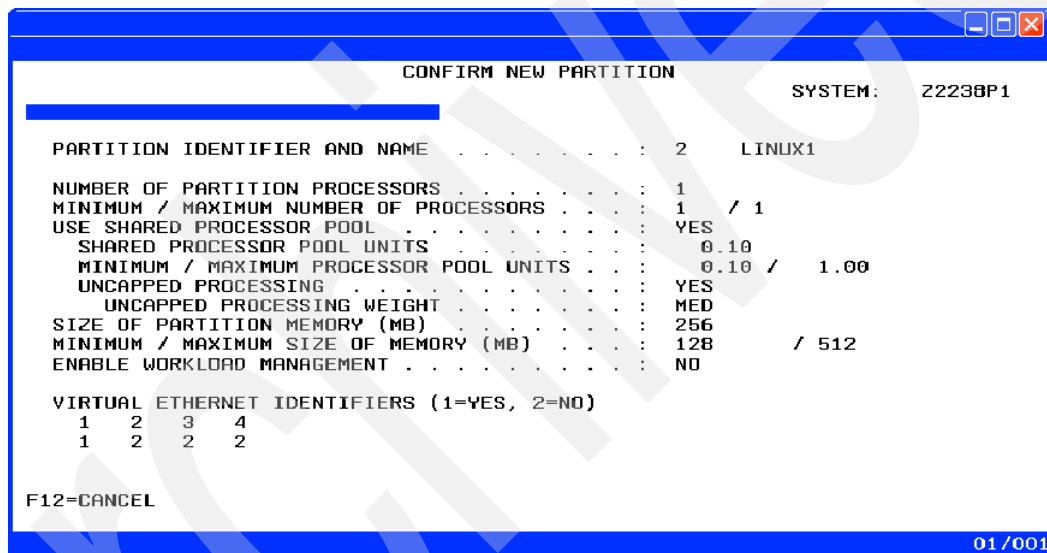


Figure 3-50 Confirm new partition

4. Once you define all of your partitions, you can view them using Option 3 from Work with System Partitions as shown in Figure 3-51. In this example, Linux4 was defined as a capped processor partition.

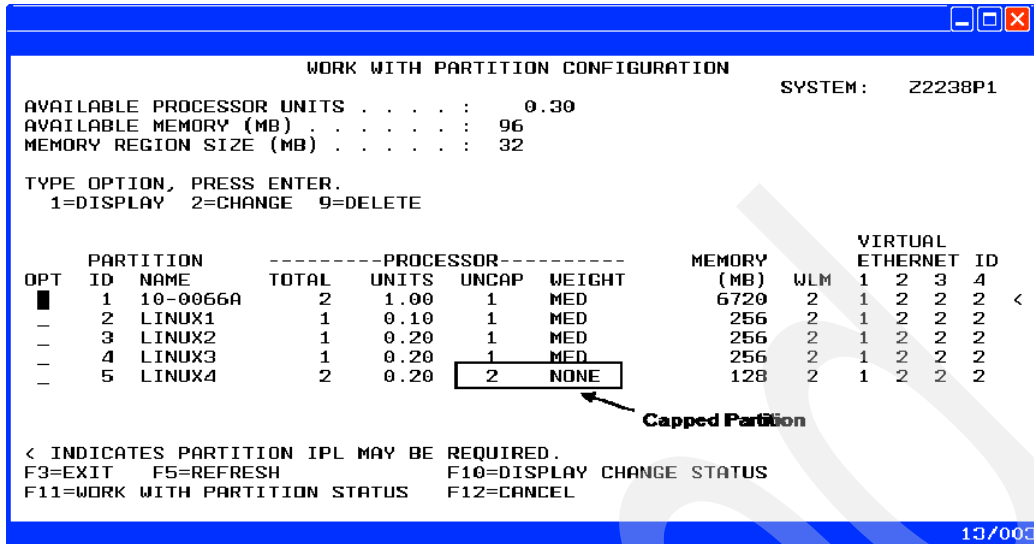


Figure 3-51 View of all newly-defined partitions

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 shown in Figure 3-52 and discussed in 3.3.3, "Recovering configuration data" on page 99.
6. Exit from the Work with Partition Configuration display to return to the Work with System Partitions. You can now start the system IPL to finish the updates required to the hypervisor to enable the new Linux partitions.
7. Select the **F10** function key as shown in Figure 3-52.

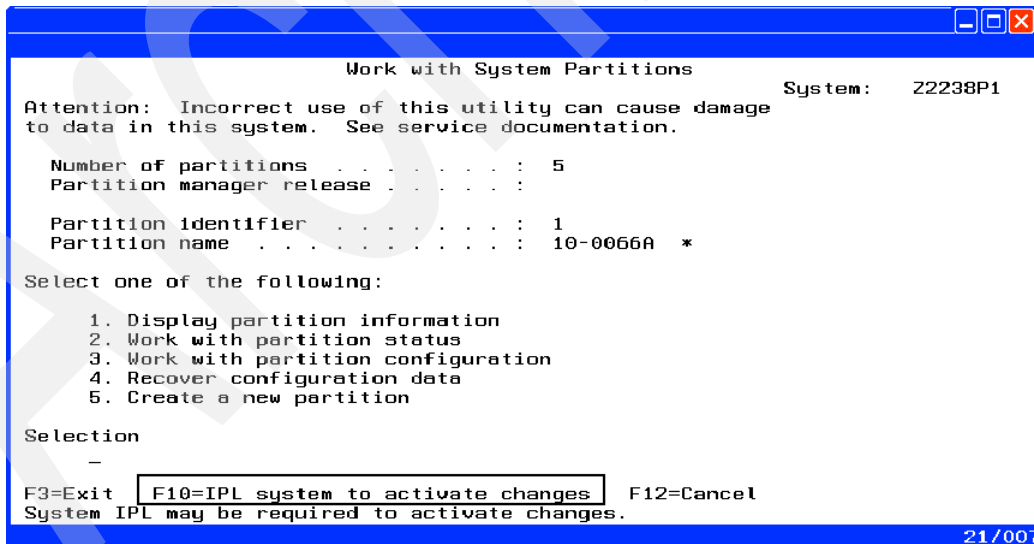


Figure 3-52 Ready to IPL with Work with System Profiles

3.3.3 Recovering configuration data

If for some reason you want to restart from the very beginning by deleting all of the partitions, select option 5 from the Work with System Partitions display, and then select Option 7 to clear configuration data, as shown in Figure 3-53.

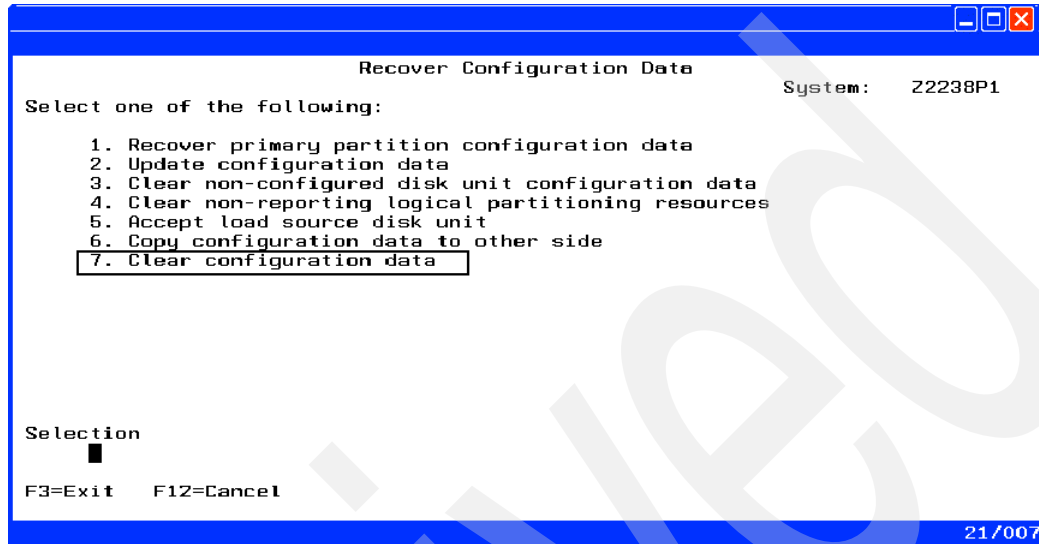


Figure 3-53 Clear configuration data

Be careful when using 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.

3.4 Linux partition with virtual disk and native network via HMC

The configuration of a Logical partition with Linux that incorporates both virtual disk and native network is simply the addition of a native network adapter to the logical partition definition. Section 3.2, “Linux partition with all virtual resources via HMC” on page 55 already presents the steps required to create the virtual I/O components within the logical partition. This section will present the additional steps required to add a native network adapter to the LPAR configuration.

Note: This section assumes that a logical partition has already been defined that uses all virtual I/O resources. It only provides the steps necessary to add the native network adapter.

To add the native network adapter, perform the following steps:

1. From the HMC (or WebSM interface), right-click the partition profile for the Linux partition and select **Properties**.

- When the Logical Partition Profile Properties window is displayed, click the **Physical I/O** tab as shown in Figure 3-54.

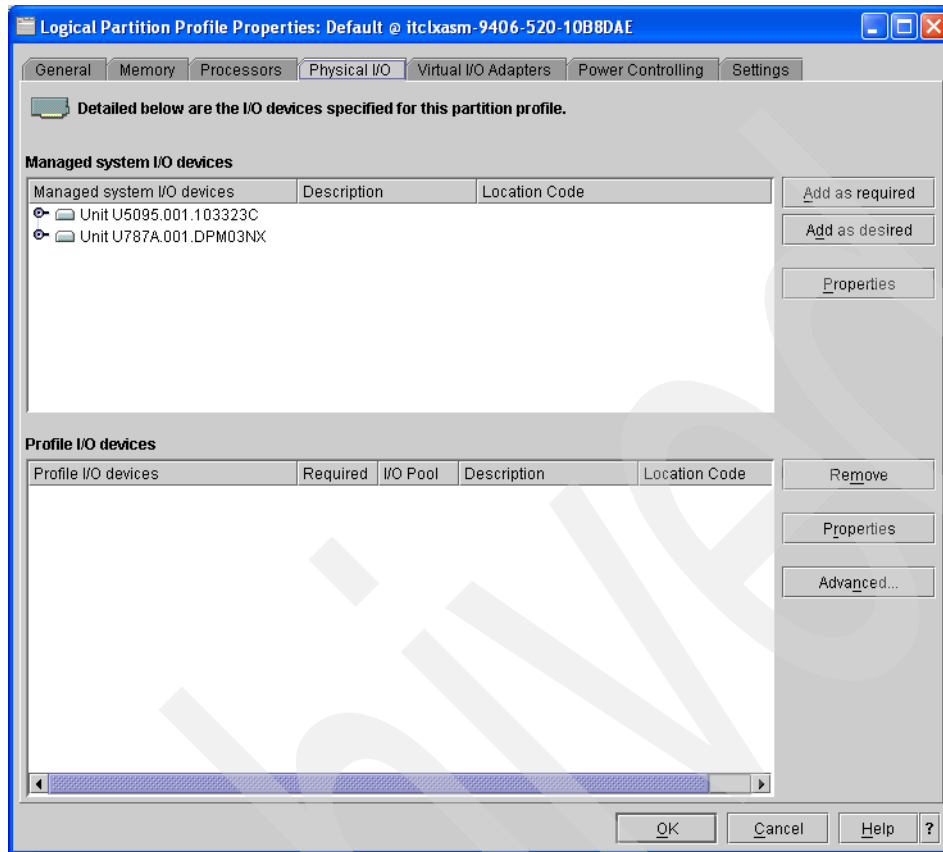


Figure 3-54 Logical Partition Profile Properties - Physical I/O

- The top portion of the window displays all of the physical hardware installed on the System i platform while the bottom portion shows all of the physical hardware that is defined within the logical partition profile. Use the top portion to navigate to the physical adapter to allocate to the Linux partition. Once the network adapter has been located, click it and then click **Add as required**. The adapter will be added to the partition profile as shown in Figure 3-55.

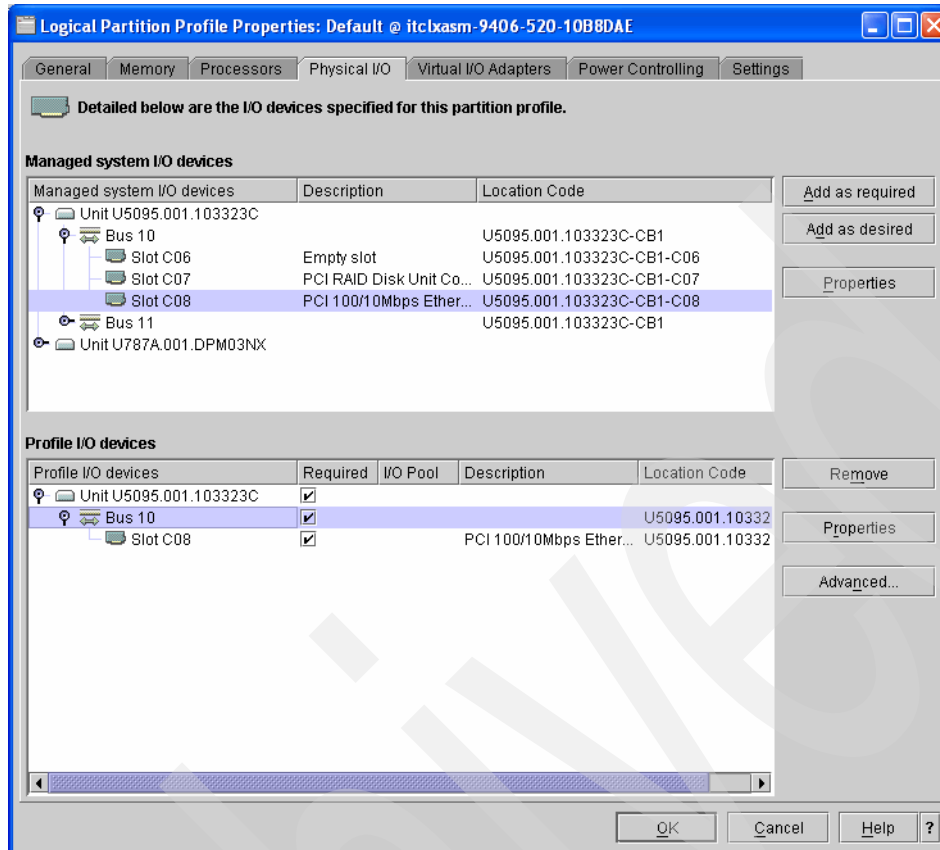


Figure 3-55 Logical Partition Profile Properties - Physical Hardware Assignment

Note: Keep in mind that only I/O Adapters (IOAs) can be assigned to Linux partitions. If the IOA intended for Linux is downstream of an IOP, you should make sure that only the IOA is assigned to Linux. Also, note that if the IOA is downstream of an IOP then it is possible the partition with the IOP will also allocate the IOA if that partition is started first.

- Once the network adapter has been added to the partition profile, click **OK** to finish the partition profile update.

Note: The partition will need to be shutdown and then restarted via the HMC in order to have the partition profile update to take effect. Dynamic LPAR can also be used to add the adapter to the running state of the partition. However, keep in mind that the Dynamic LPAR support packages will need to be installed in the Linux operating system and the adapter would still need to be added to the partition profile.

3.5 Linux partition with all physical resources via HMC

In addition to supporting virtual resources (that is, resources hosted by another partition), Linux partitions can also be established on the System i platform that have actual physical resources allocated to them.

This section will walk through the creation of the Logical partition with physical hardware allocations through the use of the HMC.

Note: Since physical hardware is being allocated to the partition, only HMC can be used. Virtual Partition Manager does not support the allocation of physical hardware to guest partitions.

The Hardware Management console is used to create partitions and partition profiles. Partition profiles are a unique aspect of partitioning on the System i platform (that is, partition profiles were not used on the previous versions of LPAR on iSeries). If you are new to LPAR or from a previous version of LPAR on the POWER4 version of the iSeries (as with OS/400 V5R2 and previous versions), you will need to be familiar with this new concept. Refer to *Logical Partitions on System i5: A Guide to Planning and Configuring LPAR with HMC on System i*, SG24-8000.

HMC has a wizard for both partition and partition profile creation. The steps in the following section describe how to use this wizard.

Note: To create a partition or system profile you must have administrator or operator authority in the HMC.

Creating the logical partition environment to support Linux can be broken down into three broad areas:

- ▶ Creating the logical partition
- ▶ Adding virtual slots (for support of Virtual I/O) to the i5/OS partition (optional)
- ▶ Registering the partition with open firmware

3.5.1 Creating the logical partition

1. In the navigation area of the HMC, select **Server and Partition** → **Server Management**. Then in the right-hand contents pane, right-click **Partitions** and select **Create** → **Logical Partition** as illustrated in Figure 3-56.

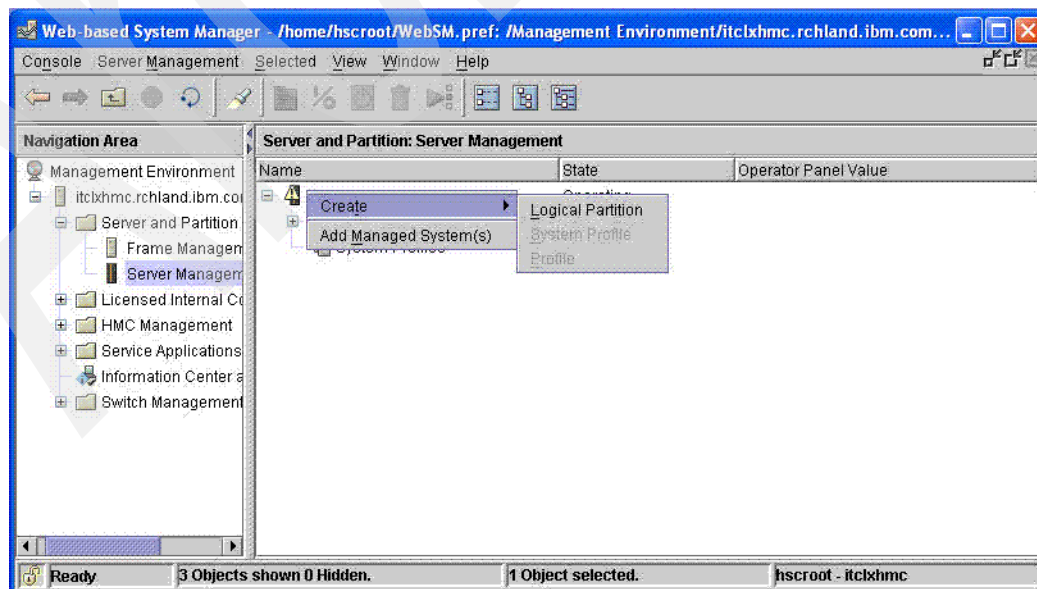


Figure 3-56 Start logical partition wizard

- The partition wizard screen appears. The Partition ID is an integer number that is used by the system to uniquely identify the partition and typically is left at the default value displayed. You need to provide the partition name and specify that the partition environment is AIX or Linux as illustrated in Figure 3-57. Once the fields have been completed click the **Next** button to continue.



Figure 3-57 Creating a new partition

- The next window provides the option for the partition to be part of partitions that can be managed by a workload application. We choose NO, as illustrated in Figure 3-58. Click **Next** to continue.

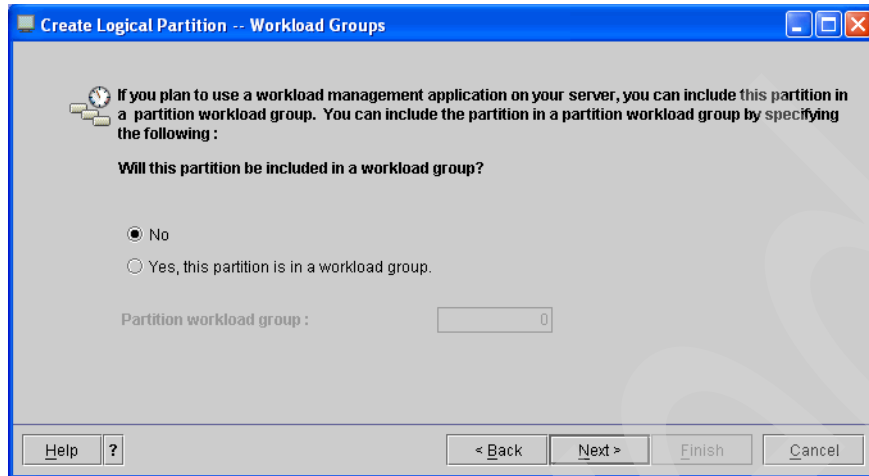


Figure 3-58 Creating a new partition - workload

- The next step is creating a partition profile. A partition profile specifies characteristics of the partition, such as the memory, processors, physical and virtual adapters.

As illustrated in Figure 3-59, give the name of the profile. In our example, we use the profile name "Default". Click **Next** to continue.

Note: Do not check the Use all resources in the system option. If you check it, the partition that is associated with this profile will try to obtain all of the physical resources in the system when the partition profile is activated. If a partition owns all of the physical resources, then there can only be one active partition on the system.

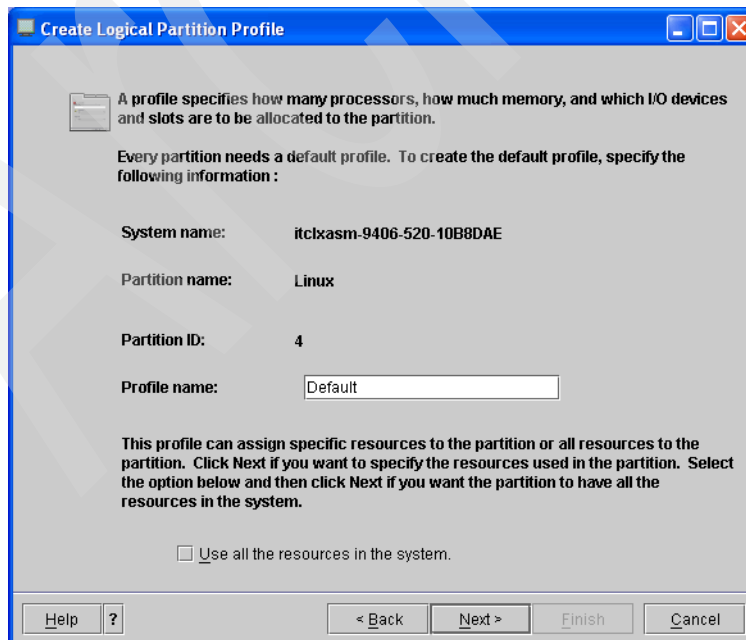


Figure 3-59 Creating a new partition - logical partition profile

5. Specify the memory size that the partition profile will manage, as illustrated in Figure 3-60. Use this page to specify the memory management information for this partition profile. You must specify the following memory levels:

- **Minimum memory:** The minimum amount of memory that the logical partition must have to be activated. If the system cannot allocate this much memory, the activation of this partition will fail.
- **Desired memory:** The amount of memory that you want the logical partition to have when you activate the logical partition.
- **Maximum memory:** The amount of memory that the logical partition is allowed to have when you dynamically move memory.

When the partition is activated, an attempt will be made to allocate the “desired” amount of memory defined for the partition profile. If the amount of memory is not available then an additional check will be made to see if the amount of unallocated memory left on the managed system is more then or equal to the “minimum” amount of memory defined for the partition profile. If the amount of memory is equal to or greater then the minimum defined memory, then it will be allocated to the partition, otherwise the partition will fail to activate.

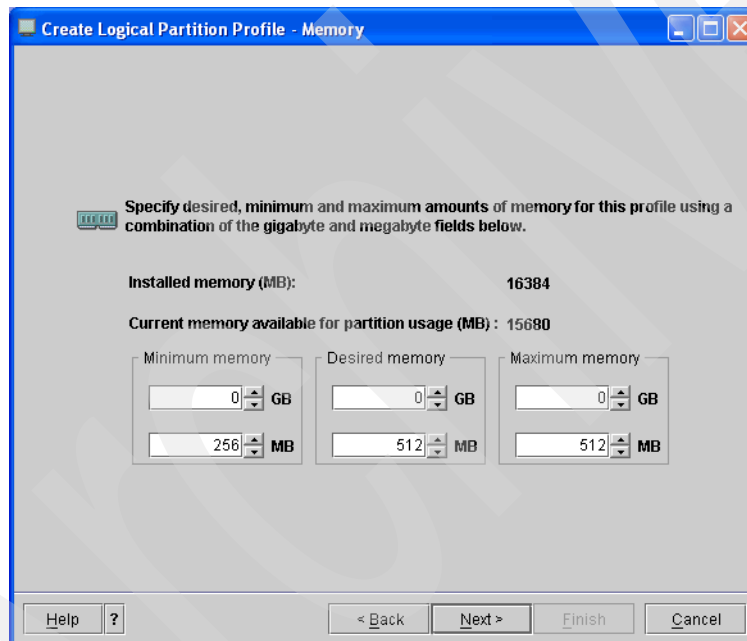


Figure 3-60 Creating a new partition - memory

In our example, the minimum quantity of memory that the partition will be activated with is 256 MB, the desired amount of memory is 512 MB and the maximum amount of memory is also 512 MB. Click **Next** to continue.

Note: Dynamic memory movement (that is, adding or removing memory from the partition while the partition is active) is not supported in the 2.6 and 2.4 versions of the Linux kernel (the kernel in currently supported distributions of Linux on the System i platform). Therefore, the ability to dynamically change the memory allocation for Linux is not supported on the System i platform.

In addition to the memory allocated to the partition, the hypervisor sets aside memory to manage the partition's memory. This additional memory is called the Hardware Paging Table and is based on the size of the maximum memory setting for the partition.

6. The next step is to choose Shared or Dedicated processors. Dedicated means you are not sharing the processor with other partitions. With the Shared option, you can utilize the subprocessor feature where each partition can run with 0.10 processing units at a minimum. While the minimum processor (for shared) is 0.10, the definition can be in the hundredths of processor units (for example, 0.15). This is referred to as micro-partitioning. In our example, we choose the Shared option, as illustrated in Figure 3-61.

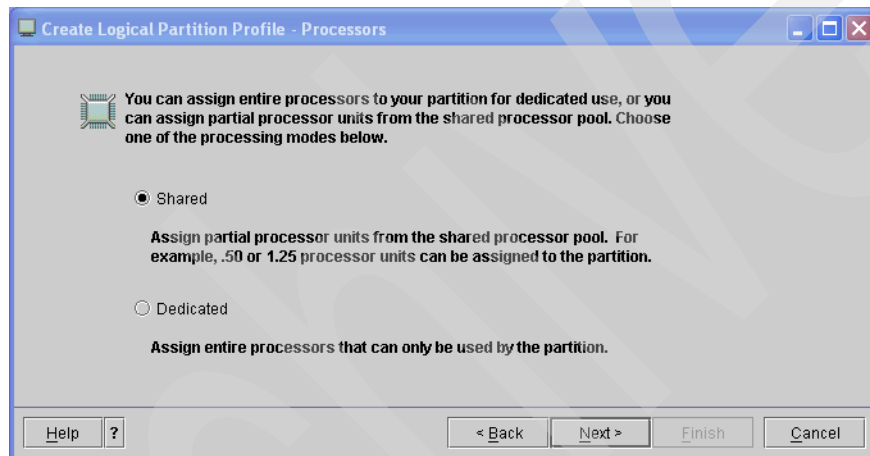


Figure 3-61 Creating a new partition - processor

7. For the Shared processor option, we choose the minimum, maximum, and desired processing units for the partition.

When the partition is activated, an attempt will be made to allocate the “desired” amount of processors defined for the partition-profile. If the amount of desired processors is not available, then an additional check will be made to see if the amount of unallocated processors left on the Managed System is more than or equal to the “minimum” amount of processors defined for the partition profile. If the amount of processors is equal to or greater than the minimum defined processors, then it will be allocated to the partition, otherwise, the partition will fail to activate.

As illustrated in Figure 3-62, we define 0.10 units of processor to be our minimum processing units, 1.0 units of processor power as our desired amount, and 1.50 as the maximum processing units. Check the **Advanced** button.

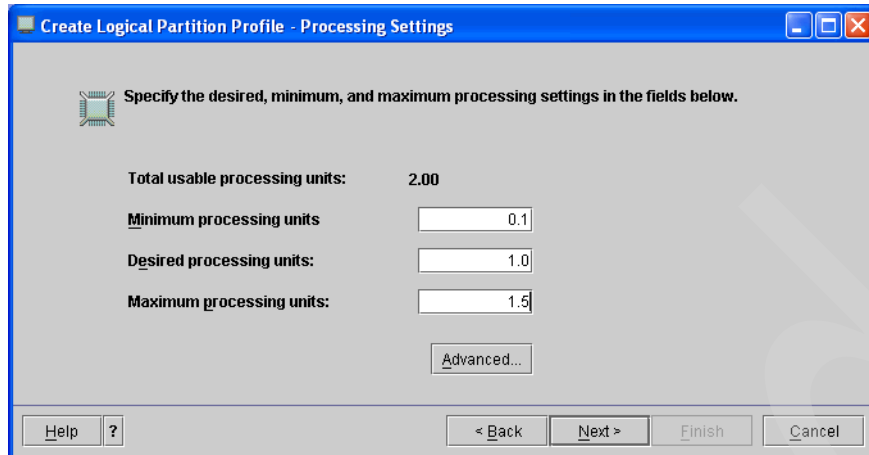


Figure 3-62 Creating a new partition - processing settings

8. When shared processing units are defined in a partition profile, additional settings can be set for the number of virtual processors as well as the capped/uncapped setting. As shown in Figure 3-63, the sharing mode can be set to either Capped or Uncapped. An uncapped partition will have the ability to have additional processing units allocated to it by the Hypervisor on an as-needed basis, based on the performance of the overall system. Capped partitions will only have access to the processing units allocated to it when first started (or changed through Dynamic LPAR).

For Uncapped partitions a weight will be specified that indicates the ratio of left over (or surplus) processing units that will be allocated to the partition. As an example, if two partitions are configured as uncapped with a weight of 10 and 20 respectively and both are busy, then the partition with a weight of 20 will get two thirds of the surplus processor resource, while the partition that has a weight of 10 will get one third of the surplus processor resource.

In addition to the sharing mode, the number of virtual processors is also defined through the Advanced Processing Settings page. The number of virtual processors indicates the number of processor threads that the workload allocated processing units will be spread across.

In this example, we set the sharing mode of the partition to uncapped with a weight of 128 and set the number of virtual processors to a minimum of 1, desired of 1 and a maximum of 2. Once the values have been set, click **Ok**, then click **Next** on the Create Logical Partition Profile - Processing Settings window.

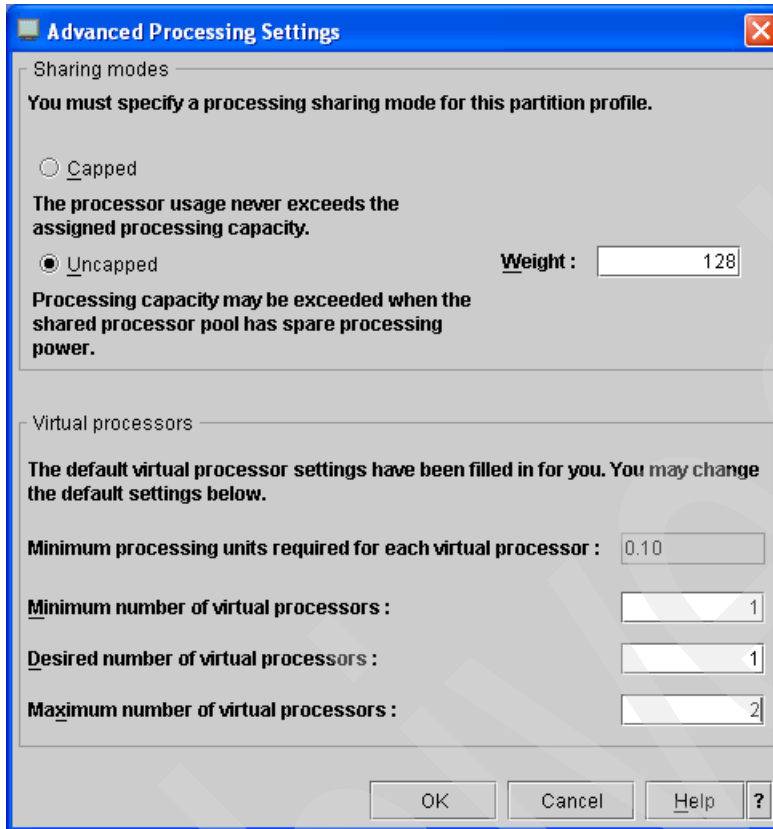


Figure 3-63 Creating a new partition - advanced processing settings

9. The next dialog box displayed is for the allocation of physical hardware to the partition. As illustrated in Figure 3-64, the window displays the hardware installed in the managed system and allows for selection of hardware for the partition profile.

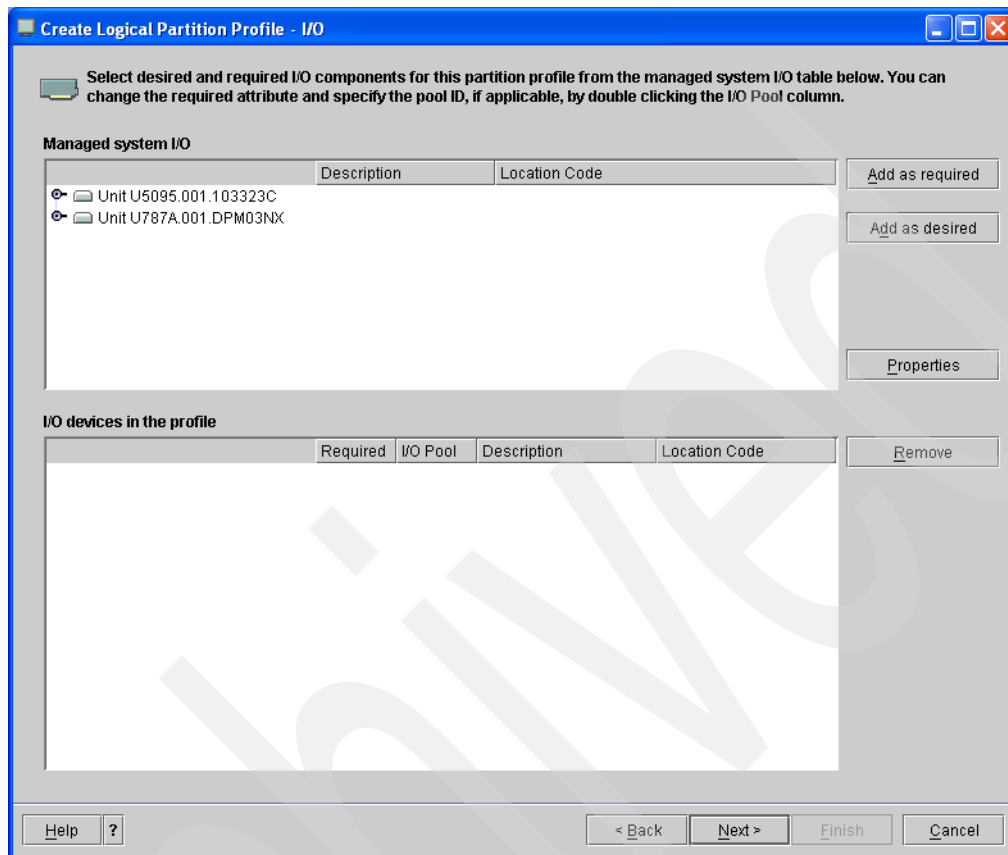


Figure 3-64 Creating a new partition - I/O

For Linux partitions, Input/Output Adapters (IOA) can be allocated for direct attached storage device as well as network adapters. The allocation of the IOA to the partition is accomplished by navigating through the display of system hardware devices in the upper portion of the display and selection of the adapter to allocate to the partition.

The top portion of the window displays all of the physical hardware installed on the System i platform while the bottom portion shows all of the physical hardware that is defined within the logical partition profile. Use the top portion to navigate to the physical adapter to allocate to the Linux partition. Once the storage controller has been located,

select it and then click **Add as required**. The adapter will be added to the partition profile as shown in Figure 3-65.

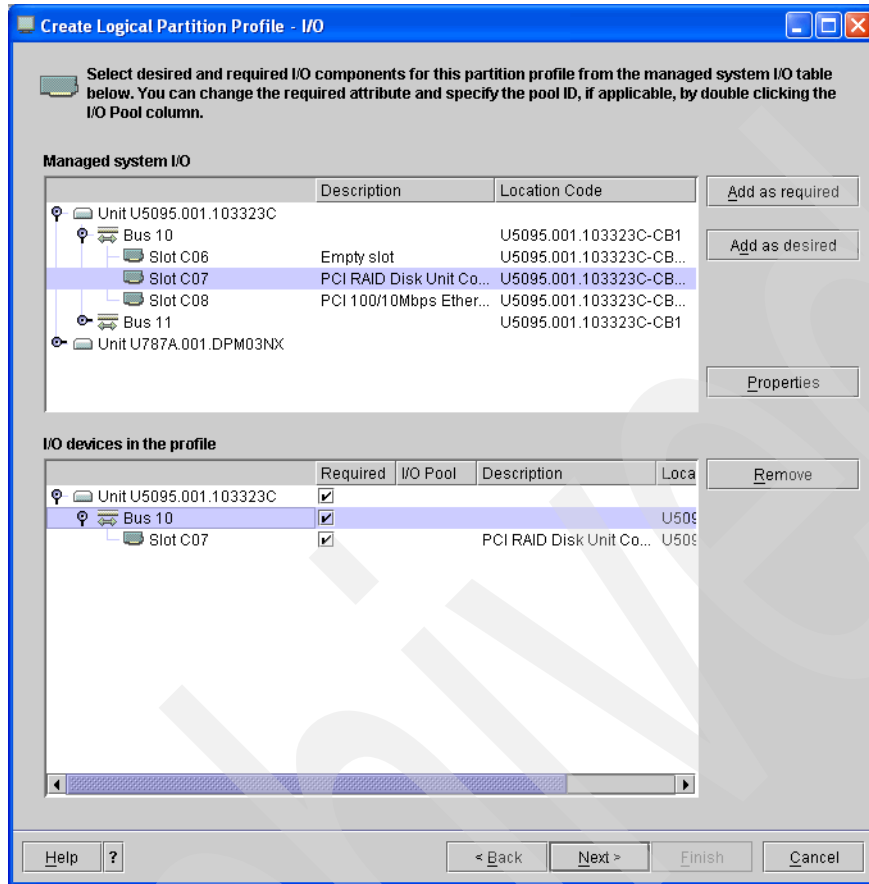


Figure 3-65 Logical Partition Profile Properties - Physical Hardware Assignment

Note: Keep in mind that only I/O Adapters (IOAs) can be assigned to Linux partitions. If the IOA intended for Linux is downstream of an IOP you need to make sure that only the IOA is assigned to Linux. Also make note that if the IOA is downstream of an IOP then it is possible the partition with the IOP will also allocate the IOA if that partition is started first.

10. In addition to assigning physical storage adapters, physical network adapters may also be allocated to a Linux partition. The allocation of the network adapter is the same as that for the storage adapter. Use the top portion to navigate to the physical adapter to allocate to the Linux partition. Once the network adapter has been located, select it and then click

Add as required. The adapter will be added to the partition profile as shown in Figure 3-66.

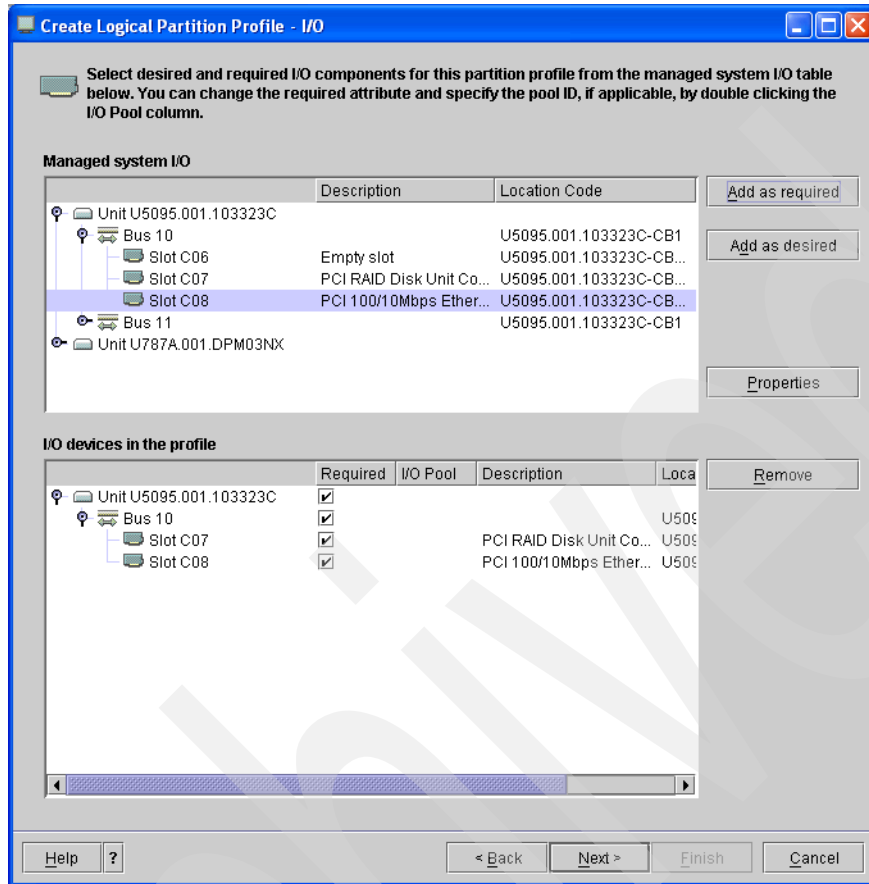


Figure 3-66 Logical Partition Profile Properties - Physical Hardware Assignments

Note: The allocation of a physical storage adapter does not then require a physical network adapter for communications. The definition of both physical and virtual adapters can be mixed in the same partition profile definition to meet any number of requirements for the partition configuration.

11. Once the physical adapters have been added to the partition profile, click **Next**.

12. The next window, illustrated in Figure 3-67, is for adding the I/O pools. In our example, we do not have I/O pools, so click **Next** to continue.

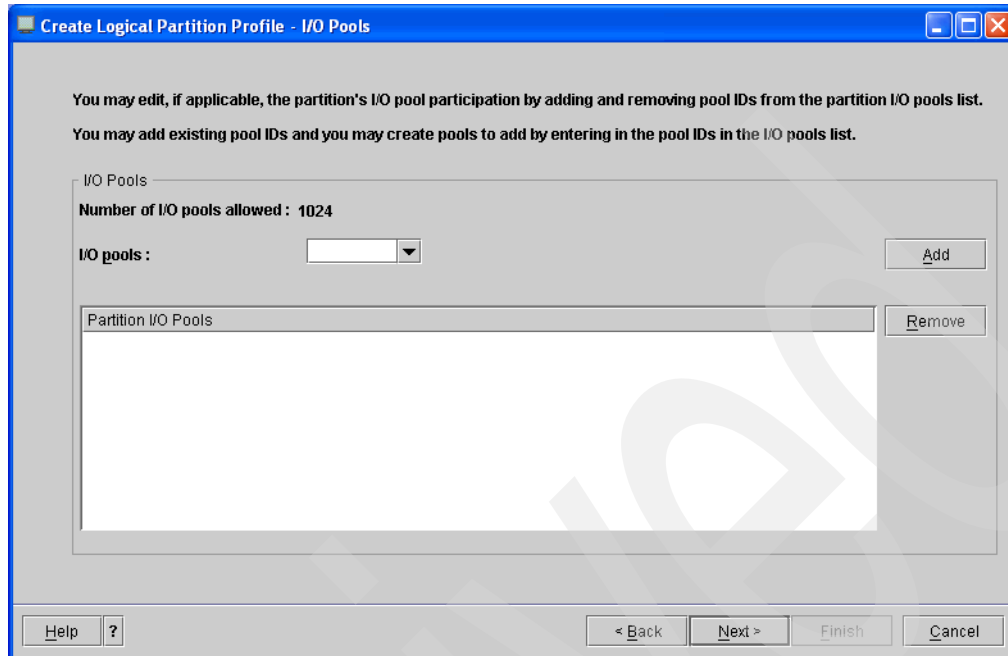


Figure 3-67 Creating a new partition I/O pools

13. The Virtual I/O Adapters window appears as shown in Figure 3-68. Since this partition is using physical resources for storage, the allocation of virtual I/O adapters is optional, however, the allocation of those adapters is supported. Some examples of configurations that may want to allocate virtual I/O adapters when native storage is being used:

- Requirement to be able to access the CD devices in i5/OS (that is, physical storage adapter is for disk only, a physical CD device for installation is not available).
- Requirement to be able to access the tape devices in i5/OS
- Usage of the Linux partition as a firewall between the physical network and partitions on the virtual network.

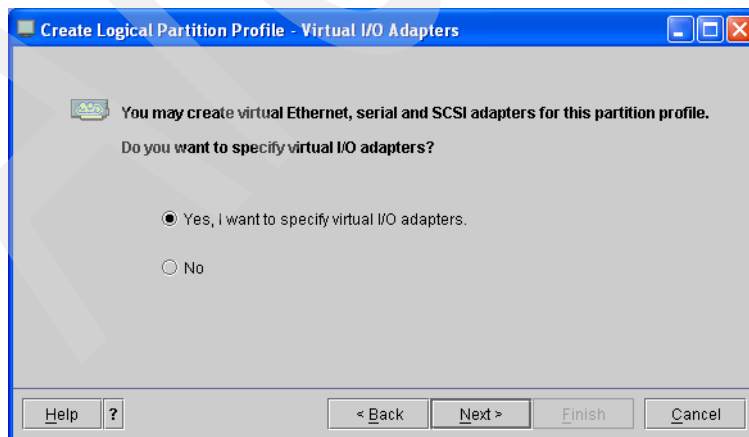


Figure 3-68 Creating a new Partition - Virtual I/O Adapters

For this example, no virtual I/O adapter will be allocated, therefore, click **No** and then click **Next**.

Note: Selecting No for virtual I/O adapters does not prevent the addition of virtual I/O adapters to the profile after it has been created. The No selection at this point in the profile definition simply means that you do not want to specify virtual I/O adapters at this time. The partition profile that results from this creation will have two virtual I/O adapters created (virtual Serial Server) and additional virtual I/O adapters can be created through modification of the partition profile.

14. A Power Controlling Partition is a partition that is allowed to make power requests (activate, shutdown) for a partition via the Hypervisor. Figure 3-69 shows the definition screen for the Power Control Partition. Since the partition will have its own I/O it is not reliant on an i5/OS component to be available for storage, therefore, a Power Controlling Partition does not need to be defined, simply click **Next** without specifying a Power Control Partition.

Note: Since a Power Control Partition is not specified, all power requests for the partition will be accomplished through the HMC.

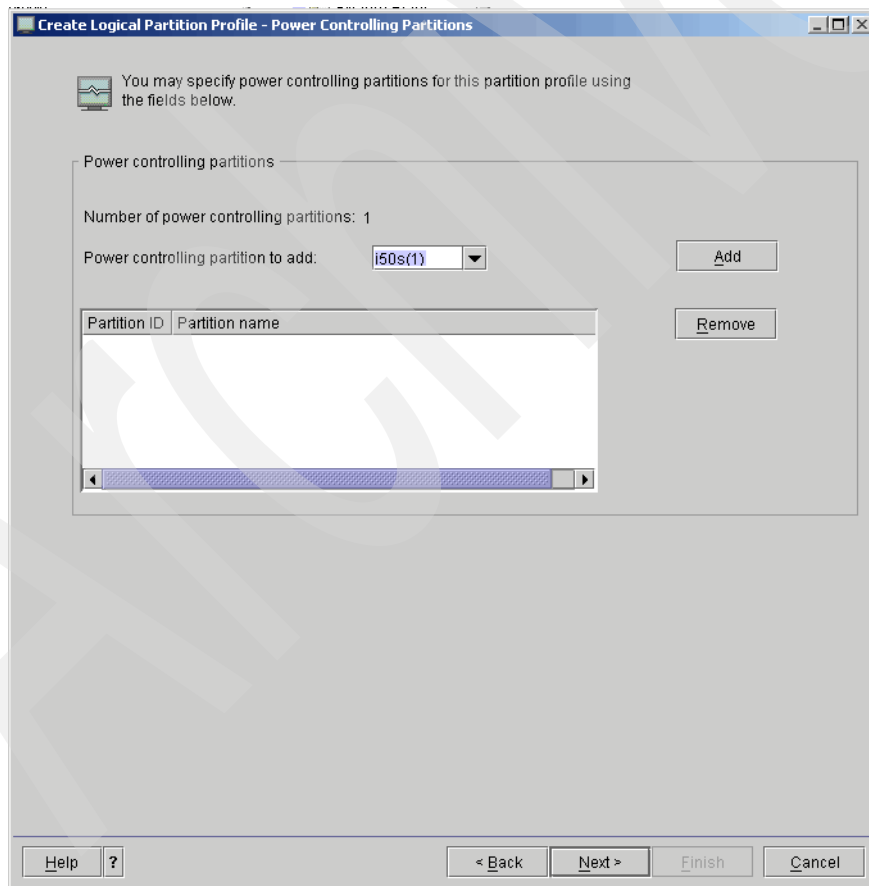


Figure 3-69 Creating a new Partition - Power Controlling Partitions

15. Optional Settings as shown in Figure 3-70 allows for the enablement of connection monitoring as well as startup options. For a Linux partition with native storage, the boot mode needs to be set to System Management Services (SMS) in order to be able to

change the device boot order of the partition to boot from attached storage. Select **System Management Services** under Boot modes and then click the **Next** button.

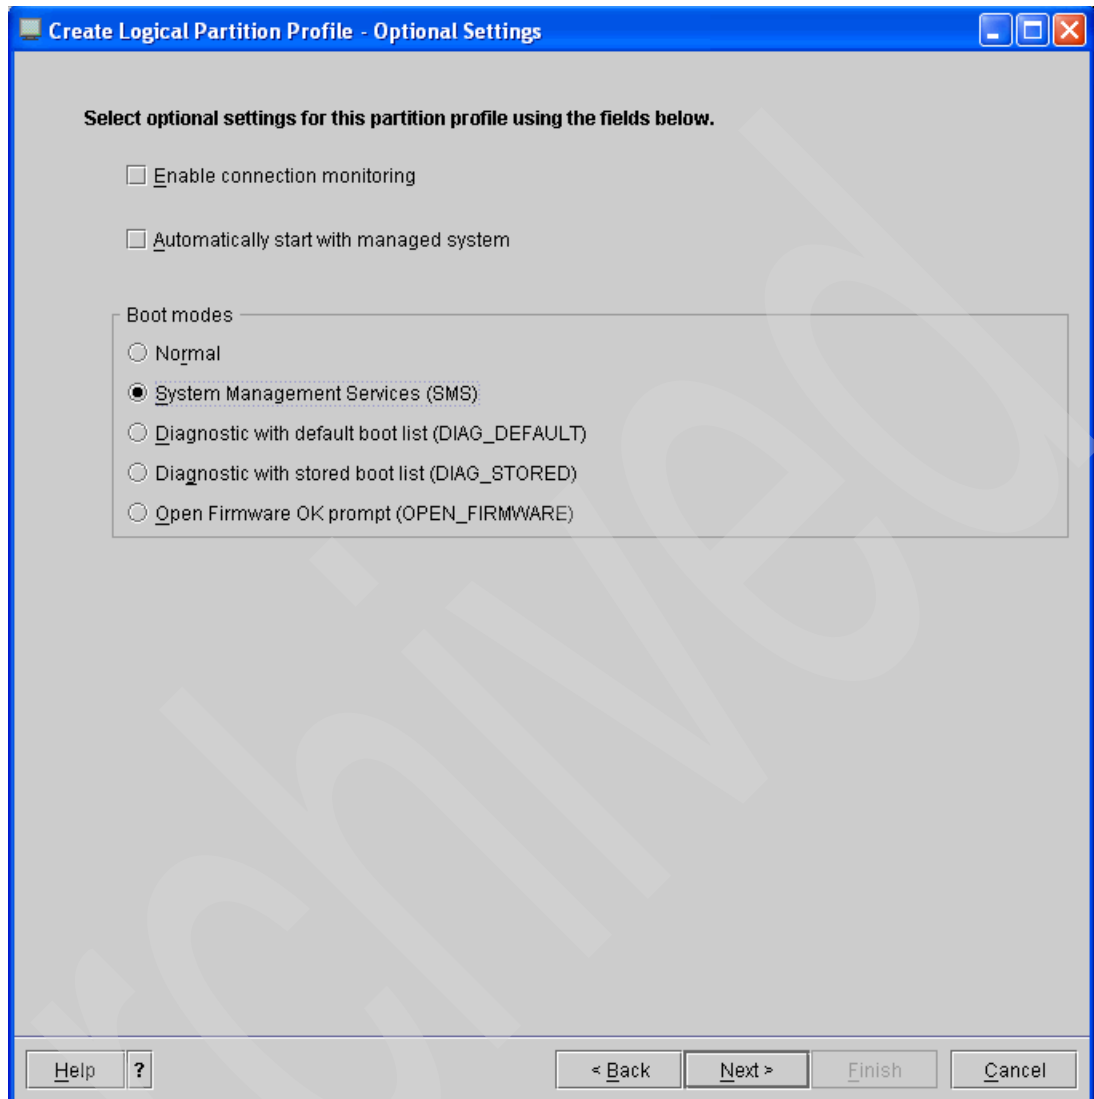


Figure 3-70 Create a new Partition - Optional Settings

16. The partition definition is now complete and a summary of the definition is shown in Figure 3-71.

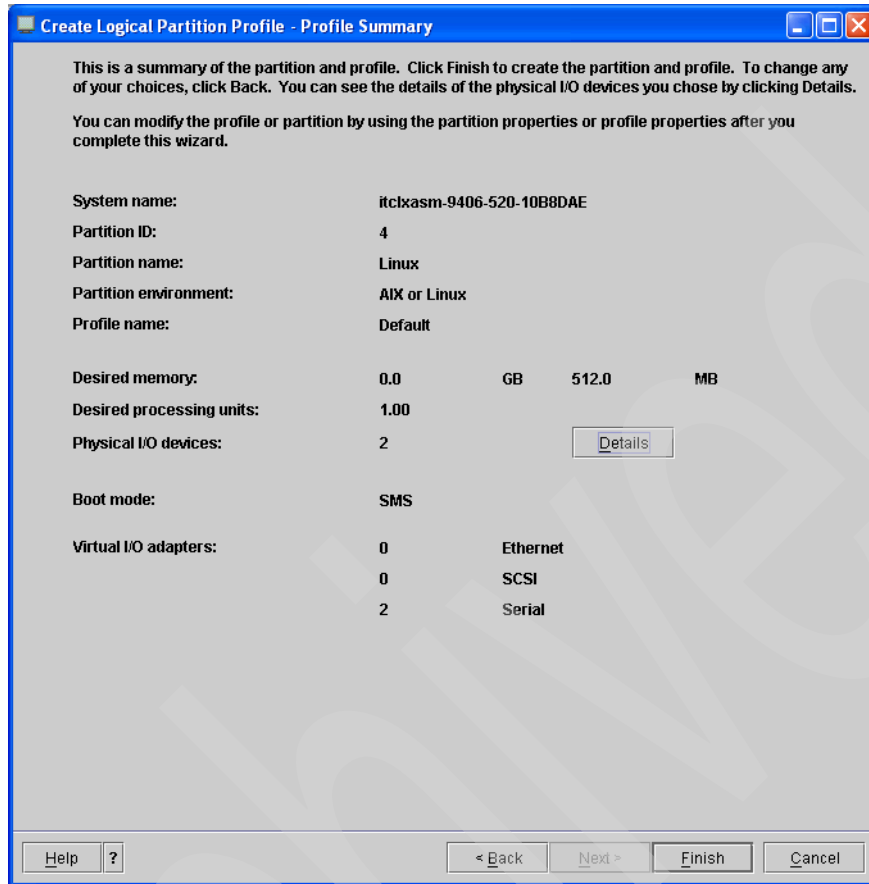


Figure 3-71 Partition profile summary

17. Click **Finish** to cause the partition to be created. A working dialog box appears as shown in Figure 3-72.




Figure 3-72 Working Dialog

This completes the definition of the logical partition. The next step will be the preparation of the native storage devices as described in Chapter 4, “Creating the storage environment” on page 117.

Archived



Creating the storage environment

This chapter will cover creating the storage environment for Linux. It provides information about the definition of virtual storage as well as native storage.

4.1 Supporting virtual I/O

Virtual I/O is the predominate storage solution for Linux on the System i platform and revolves around storage being hosted by an i5/OS partition and made available to a Linux partition. You should have already created the Linux partition with a virtual SCSI client adapter that points to a virtual SCSI server adapter created in the i5/OS partition. This section will walk you through creating the i5/OS components to support virtual I/O.

4.1.1 Network server description (NWS)

The Network Server Description defines a number of parameters for the Linux environment including the startup location as well as startup parameters.

Perform the following steps to define the Network Server Description:

1. Start the Create Network Server Description using the CRTNWS command as shown in Figure 4-1.

```
                                Create Network Server Desc (CRTNWS)

Type choices, press Enter.

Network server description . . . > LINUX           Name
Resource name . . . . . > *AUTO                 Name, *NONE, *AUTO
Network server type:
  Server connection . . . . . > *GUEST          *IXSVR, *ISCSI, *GUEST...
  Server operating system . . . > *LINUXPPC      *WIN32, *LINUX32...
Online at IPL . . . . . *YES                    *YES, *NO
Vary on wait . . . . . *NOWAIT                 *NOWAIT, 1-15 minutes
Shutdown timeout . . . . . 15                  2-45 minutes
Partition . . . . . Linux

Partition number . . . . . *NONE               Number, *NONE
Code page . . . . . *LNGVER                   *LNGVER, 437, 850, 852,
857...
Server message queue . . . . . *JOBLOG         Name, *JOBLOG, *NONE
Library . . . . .                               Name, *LIBL, *CURLIB

                                                                More...
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel  F13=How to use this display
F24=More keys
```

Figure 4-1 Create network server description

- **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 Virtual SCSI server adapter for the partition.

Note: The value of *AUTO for the Resource name indicates that the system will determine the i5/OS resource for the Server SCSI adapter by looking at the partition definition of the Linux partition and determining the first client/server pairing for the SCSI adapter. For environments that have more than one client/server SCSI pairing, the actual resource name will need to be used for the subsequent (that is, not the first) adapters. The resource name can be determined through use of the i5/OS command (WRKCFGSTS *CMN) and review of the location information for the CTL resources.

- **Network server type** → **Server connection:** This indicates the type of server that this network server is being created for. Since it is being created for a Linux guest partition, enter *GUEST.
- **Network server type** → **Server operating system:** This indicates the operating system type which in the case of Linux in a logical partition is set to *LINUXPPC.

Note: Prior to V5R4, the Network server type was a single field and was set to *GUEST for Linux in a logical partition.

- **Online at IPL:** This field is typically left set at the default of *YES. It indicates whether or not this network server should be varied on when i5/OS is started. By leaving the value at *YES, when i5/OS is started the Linux environment will also be started.
- **Vary on wait:** This field is typically left set at the default value of *NOWAIT. It indicates the amount of time to wait after start of i5/OS before vary on of this Network Server. It is typically used when the operating system (in this case Linux) needs to wait for another operating system/server to start (such as when a Linux file server is integrated with an Windows Active Directory® Server running on an IXS/IXA).
- **Shutdown timeout:** This field is typically left set at the default value of 15. It indicates how long i5/OS will wait on a vary off before forcing the network server into a vary off state. Typically, when a network server is varied off it will send a signal to the operating system indicating that it should shut-down. Once the operating system is shutdown the network server will go into a vary off state. This value only comes into play when the shutdown of the operating system fails.
- **Partition:** Name of the logical partition for Linux. This name must match the name given in the definition of the Linux partition. In this case where the partition has been created via the HMC, you must ensure that the partition name matches the mix of upper- and lower-case characters used in the partition definition.
- **Partition number:** Unique number identifier for the partition. This field cannot be specified if the Partition name field has been specified. Typically for the Network Server definition the name of the partition is used and this field is left at the default value of *NONE.
- **Code page:** Leave at the default value of *LNGVER.

Note: The Code page value has no bearing on language support for the operating system. The code page indicates the code page that should be used when i5/OS communicates with the hypervisor of the System i platform.

2. After you complete the first page of parameters, press the **PgDn** key shown in Figure 4-2.

```

                                Create Network Server Desc (CRTNWS D)

Type choices, press Enter.

TCP/IP port configuration:
Port . . . . . *NONE          *NONE, 1, 2, 3, 4...
Internet address . . . . .
Subnet mask . . . . .
Maximum transmission unit . . . . . Number
Gateway address . . . . .
      + for more values

TCP/IP route configuration:
Route destination . . . . . *NONE
Subnet mask . . . . .
Next hop . . . . .
      + for more values

TCP/IP local host name . . . . . *NWS D

More...

F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 4-2 Network Server Description - Page 2

For most Linux partitions, page 2 of the Network Server Description is not changed. The values on this page are used to define the network parameters for the operating system that will be using the resources provided by this Network Server Description. At present, the values entered on this page are made known to Linux but they are not actually used to change the network settings.

Press **PgDn** to see the third page of the Network Server Description as shown in Figure 4-3.

```

                                Create Network Server Desc (CRTNWS D)

Type choices, press Enter.

TCP/IP local domain name . . . . *SYS

TCP/IP name server system . . . . *SYS
      + for more values
Restricted device resources . . . *NONE      Name, *NONE, *ALL...
      + for more values
Synchronize date and time . . . . *TYPE      *TYPE, *YES, *NO
IPL source . . . . .              *STMF      *NWSSTG, *PANEL, *STMF, A...
IPL stream file . . . . .         '/qopt/SLES100.001/SUSEBOOT/INST64'

IPL parameters . . . . .         'VNC=1
Power control . . . . .          *YES       *YES, *NO
Authority . . . . .              *CHANGE   Name, *CHANGE, *ALL, *USE...

More...

F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys

```

Figure 4-3 Network Server Description - Page 3

- **TCP/IP local domain name:** Leave as default value of *SYS.
- **TCP/IP name server system:** Leave as default value of *SYS
- **Restricted device resources:** Typically this value is left at the default of *NONE. This field is used to indicate if any I/O owned by the i5/OS partition should not be made available to the Linux partition. By default, when a Network Server is created against a Linux partition, the CD/DVD and Tape devices owned by i5/OS as well as any Network Storage Spaces (refer to 4.2, “Native storage” on page 126) are made available as Virtual I/O devices to the Linux partition. In addition to specify the name(s) of specific devices to restrict from the Linux partition, other possible values for this field include:
 - ***ALL:** Restrict all I/O devices from the Linux partition with the exception of those Network Storage devices that are explicitly linked to the Network Server.
 - ***ALLTAPE:** Restrict all tape devices from the Linux partition.
 - ***ALLOPT:** Restrict all optical (that is, CD/DVD/DVD-RAM) devices from the Linux partition.
- **Synchronize date and time:** Leave at the default value of *TYPE.
- **IPL source:** Indicates where to look for the initial boot file. Possible settings for this field include:
 - ***NWSSTG:** indicates that the initial boot file is in the bootable disk partition (PrepBoot) of the first disk linked to the Network Server Description.

- ***STMF**: Indicates that the initial boot file is a stream file located in the Integrated File System (IFS). When the setting is *STMF, the path will be specified in the IPL stream file field.
- ***PANEL**: Indicates that the setting in the HMC for the partition should be used.
- **A, B, D**: Not applicable for Network Servers associated with Linux partitions.

Note: The installation of Linux is typically performed with IPL source set to *STMF. As part of the steps in the Linux installation process, the setting will be changed to *NWSSTG to indicate that Linux should boot of the Linux kernel written to the virtual disk by the installation program.

- **IPL stream file:** The IPL stream file is specified when the IPL source is set to *STMF. The path indicates the path to the initial boot file located in the IFS.

Note: The initial boot file for an installation of Linux is indicated in the documentation provided with the distribution. As of the publication of this book, here are some of the more common installation paths:

- ▶ Novell / SUSE SLES 10 /qopt/SLES100.001/SUSEBOOT/INST64
- ▶ Novell / SUSE SLES 9 (GA) /qopt/SU90.001/install
- ▶ Novell / SUSE SLES 9 (SP3) /qopt/SU90SP3.P01
- ▶ Red Hat RHEL 4 /qopt/RED_HAT/IMAGES/PSERIES

- **IPL parameters:** Defines parameters that are sent directly to the Linux kernel. The IPL parameters are defined in the Linux-distributor documentation.

Note: Typically, the IPL parameters are only used for installation (although they could be used for run-time as well to meet special requirements. Both Novell / SUSE SLES 9 & 10 as well as Red Hat RHEL 4 support graphic-based installations. As of the publication of this book, here is the IPL parameter that supports a graphical based installation via VNC for each distribution:

- ▶ Novell / SUSE SLES 9 VNC=1
- ▶ Novell / SUSE SLES 10 VNC=1
- ▶ Red Hat RHEL 4 VNC

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 or lower case.

- **Power control:** Indicates if this is a power controlling Network Server. A power controlling Network Server is one that sends power requests to the Hypervisor when a vary operation (such as vary on or vary off) is performed on the Network Server. Put another way, when a power controlling Network Server is varied on, the associated logical partition will be activated and when the power controlling Network Server is varied off, the associated Logical partition will be shutdown.

For most Linux implementations on the System i platform, the setting will be left at the default of *YES. A setting of *NO is used when there are multiple Network Servers defined for the same logical partition.

Note: For the power-control requests sent by the vary operation to be accepted by the Hypervisor, the Power Controlling Partition setting of the logical partition for Linux must be set to the name of the i5/OS partition that will issue the power request.

- **Authority:** Leave set as the default value of *CHANGE.

Press **PgDn** to see the fourth (and last) page of the Network Server Description as shown in Figure 4-4.

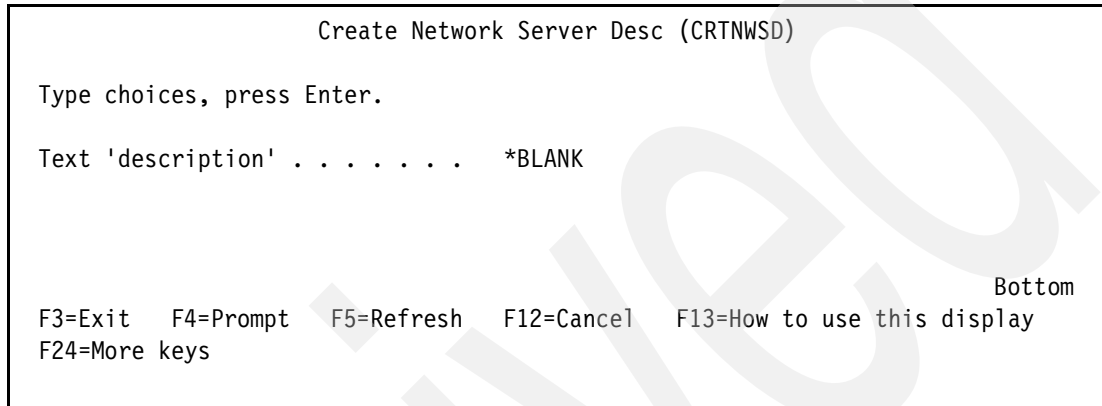


Figure 4-4 Network Server Description - Page 4

- **Text description:** Enter a description of the usage of the Network Server.
3. After you complete the fields, press **Enter** to create the Network Server Description.

4.1.2 Network Storage Space

The Network Server Storage Space (NWSSTG), also referred to as a virtual disk, provides the disk resource for the Linux partition. The Network Server 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 as shown in Figure 4-5.

```
                                Create NWS Storage Space (CRTNWSSTG)

Type choices, press Enter.

Network server storage space . . > LINUX           Name
Size . . . . . > 6124                             *CALC, 1-1024000 megabytes
From storage space . . . . . *NONE                 Name, *NONE
Format . . . . . > *OPEN                           *NTFS, *FAT, *FAT32, *OPEN...
Auxiliary storage pool ID . . . 1                  1-255
ASP device . . . . .                               Name
Text 'description' . . . . . Storage Space for Linux Partition

                                                                Bottom
F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys
```

Figure 4-5 Create Network 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 MB to 1 TB. 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 10 and Red Hat RHEL 4, a size of 6 GB (6124) 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.

- **From storage space:** This field is used when making a copy of an existing storage space. Since this storage space is being created for a new installation, leave the value set at the default of *NONE.
- **Format:** Indicates the type of “disk” partition to place on the Network Server Storage Space. A value of *OPEN indicates that no disk partition will be placed on the storage space. This is the value to use for Linux storage spaces as the Linux installer will put partitions on the disk.
- **Auxiliary storage pool ID:** Indicates the storage pool where the IFS files that represent the storage space will be placed. Most times this value is left at the default value of 1.

- **ASP device:** Indicates the name of the ASP that the storage space will be placed into. Most times this value is left blank.
 - **Text 'description':** A description of the usage of the storage space.
2. After the fields have been completed press **Enter** to create the Network Server Storage Space. At this point the space requested is reserved in the IFS and the Network Server Storage Space is available for use.
 3. 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, and press the **F4** key as shown in Figure 4-6.

```

Add Server Storage Link (ADDNWSSTGL)

Type choices, press Enter.

Network server storage space . . > LINUX      Name
Network server description . . . > LINUX      Name
Dynamic storage link . . . . . *NO          *NO, *YES
Access . . . . . *UPDATE          *UPDATE, *READ, *SHRUPD
Drive sequence number . . . . . *CALC       1-64, *CALC, *QR
Storage path number . . . . . *DFTSTGPTH    1-4, *DFTSTGPTH, *MLTPTHGRP

Bottom
F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys

```

Figure 4-6 Add Server Storage Link

- **Network server storage space:** The name of the Network Server Storage Space to be linked to the Network Server.
- **Network server description:** The name of the Network Server to which the storage space is linked. The Network Server is the component that provides access to the storage spaces, via the virtual SCSI adapters that are linked to the Network Server.
- **Dynamic storage link:** Leave as *NO (the default).
- **Access:** Typically left at the default of *UPDATE. The possible values are:
 - ***UPDATE:** The storage space can be linked to a single network server and the associated operating system will have complete read/write access to the resource.
 - ***READ:** The storage space can be linked to multiple network servers and the associated operating systems will have concurrent read-only access to the resource.

- ***SHRUPD:** The storage space can be linked to multiple network servers, however, only one of the associated operating systems can have access to the resource at any given time.
- **Drive sequence number:** Typically left at the default of *CALC. Indicates the ordering of the drives as seen by the associated operating system.
- **Storage path number:** Leave at the default of *DFTSTGPTH.

4.2 Native storage

The Hardware Management Console (HMC) will be used to start the Linux environment when Linux is using native storage. In order to start the installation program, the boot order of the partition has to be changed.

Note: When installing Linux to direct attached storage, it is possible to still use the virtual I/O support for, among other things, access to the installation source. If a CD drive (or virtual CD library) from an i5/OS partition is going to be used for the installation source then you should ensure that the following values are set in the associated Network Server Description:

- ▶ IPL source *PANEL

The *PANEL value for IPL source indicates that the boot settings of the partition should be used to determine the IPL source.

- ▶ IPL stream file *NONE

The *NONE value for IPL stream file indicates that there is no stream file in the IFS to boot from.

- ▶ Power control *NO

The *NO value for Power control indicates that power requests should not be sent to the Hypervisor (to power on/off the logical partition) when a corresponding vary operation is performed on the Network Server.

1. In the HMC, right-click the Linux partition and select **Activate**.

2. When the Activate Logical Partition dialog is displayed as shown in Figure 4-7, click **Advanced**.

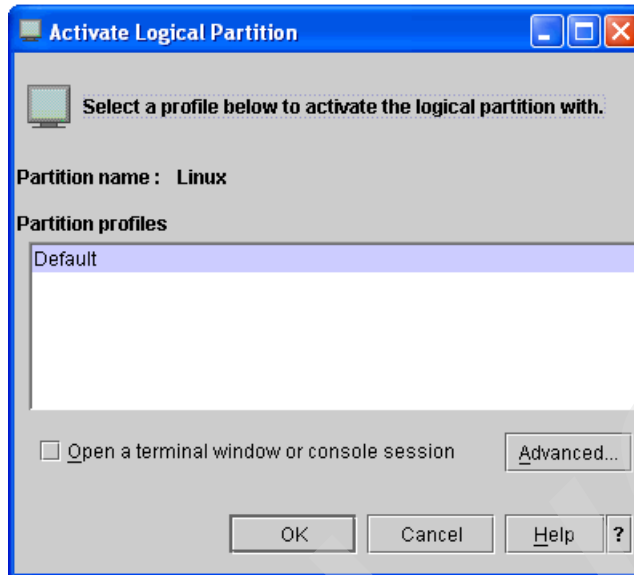


Figure 4-7 Activate Logical Partition dialog box

3. On the Advanced dialog (Figure 4-8), select **SMS** from the Boot mode drop-down list and then click **OK**.

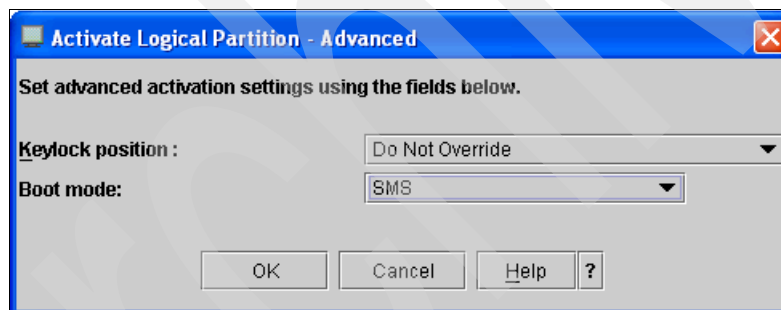


Figure 4-8 Activate Logical Partition - Advanced

4. In the Activate Logical Partition window, click **OK**. The partition will be started and the SMS menu is displayed as shown in Figure 4-9.

Note: If you have not previously accessed the Linux console, then the check box for Open a terminal window or console session should be selected in Figure 4-8.

If a virtual SCSI client/server pairing is being used to provide access to the CD/DVD drive in i5/OS then the associated network server should be varied on prior to activation of the partition.

```
Version SF240_219
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
-----
Main Menu
1.  Select Language
2.  Setup Remote IPL (Initial Program Load)
3.  Change SCSI Settings
4.  Select Console
5.  Select Boot Options
-----
Navigation Keys:
                                     X = eXit System Management
Services
-----
Type menu item number and press Enter or select Navigation key:
```

Figure 4-9 SMS Menu

- 5. Type 5 (Select Boot Options) and press **Enter**.
- 6. In the Multiboot screen shown in Figure 4-10, type 1 (Select Install/Boot Device) and press **Enter**.

```
Version SF240_219
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
-----
Multiboot
1.  Select Install/Boot Device
2.  Configure Boot Device Order
3.  Multiboot Startup <OFF>
-----
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen
                                     X = eXit System Management
Services
-----
Type menu item number and press Enter or select Navigation key:
```

Figure 4-10 Multiboot menu

7. In the Select Device Type menu shown in Figure 4-11, type 3 (CD/DVD) and press **Enter**.

```
Version SF240_219
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
-----
Select Device Type
1.  Diskette
2.  Tape
3.  CD/DVD
4.  IDE
5.  Hard Drive
6.  Network
7.  List all Devices

-----

Navigation keys:
M = return to Main Menu
ESC key = return to previous screen      X = eXit System Management
Services

-----

Type menu item number and press Enter or select Navigation key:
```

Figure 4-11 Select device type menu

8. In the Select Media Type menu shown in Figure 4-12, type 1 (SCSI) and press **Enter**.

```
Version SF240_219
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
-----
Select Media Type
1.  SCSI
2.  SSA
3.  SAN
4.  IDE
5.  ISA
6.  List All Devices

-----

Navigation keys:
M = return to Main Menu
ESC key = return to previous screen      X = eXit System Management
Services

-----

Type menu item number and press Enter or select Navigation key:
```

Figure 4-12 Select media type menu

At this point, a list of the possible SCSI devices to boot from will be displayed as shown in Figure 4-13. The number of devices displayed will be dependent on the number of SCSI devices available to the partition.

```
Select Media Adapter
1.          U9406.520.10B8DAE-V2-C2-T1  /vdevice/v-scsi@30000002
2.          U5095.001.103323C-CB1-C07-T1
/vpci@80000002000000a/pci@2,2/pci1069,b166@1/scsi@0
3.          U5095.001.103323C-CB1-C07-T2
/vpci@80000002000000a/pci@2,2/pci1069,b166@1/scsi@1
4.          U5095.001.103323C-CB1-C07-C1-T1
/vpci@80000002000000a/pci@2,2/pci1069,b166@1/scsi@ff
5.    List all devices

-----

Navigation keys:
M = return to Main Menu
ESC key = return to previous screen      X = eXit System Management
Services

-----

Type menu item number and press Enter or select Navigation key:
```

Figure 4-13 Select media adapter

Note: If virtual I/O is being used to access the installation source from an i5/OS CD device, then the media adapter to select will have a /vdevice/v-scsi indication as shown in option 1 in Figure 4-13.

9. Type the number of the appropriate SCSI adapter (you may need to try several times to get the right one) and press the **Enter** key.

Note: When a client/server SCSI relationship exists between Linux and i5/OS partitions and a Network Server Description has been created against the server SCSI adapter, all i5/OS CD/DVD resources (as well as tape and any linked network storage spaces) are placed on the resulting SCSI bus and made known to Linux. Each of the CD/DVD (that is, OPT##) devices in i5/OS will show up as optical devices in Linux and, therefore, it is possible that you will have to go through the process of selecting the CD device several times until the correct one is selected.

10. From the List of SCSI devices as shown in Figure 4-14, type the number of that represents the CD/DVD drive that the installation media has been mounted on and press **Enter**.

```
Version SF240_219
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
-----
Select Device
Device Current Device
Number Position Name
1.      1      SCSI CD-ROM
          ( loc=U9406.520.10B8DAE-V2-C2-T1-W802000000000000-L0 )
2.      -      SCSI CD-ROM
          ( loc=U9406.520.10B8DAE-V2-C2-T1-W812000000000000-L0 )
-----
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen      X = eXit System Management
Services
-----
Type menu item number and press Enter or select Navigation key:
```

Figure 4-14 Select Device

11. In the Select Task menu shown in Figure 4-15, type 2 (Normal Mode Boot) and press **Enter**.

```
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
-----
Select Task

SCSI CD-ROM
( loc=U9406.520.10B8DAE-V2-C2-T1-W802000000000000-L0 )

1. Information
2. Normal Mode Boot
3. Service Mode Boot
-----
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen      X = eXit System Management
Services
-----
Type menu item number and press Enter or select Navigation key:
```

Figure 4-15 Select Task

12. When prompted to exit as shown in Figure 4-16, type 1 (Yes) and press **Enter**.

```
Version SF240_219
SMS 1.6 (c) Copyright IBM Corp. 2000,2005 All rights reserved.
-----
Are you sure you want to exit System Management Services?
1.  Yes
2.  No
-----
Navigation Keys:
                                     X = eXit System Management
Services
-----
Type menu item number and press Enter or select Navigation key:
```

Figure 4-16 Exit menu

Note: If the wrong CD/DVD device was selected in step 9 above then the system will display the initial LPAR start screen followed by a line that indicates List all Devices. If that occurs, press **Enter** to return to the SMS menu and try a different CD/DVD device.

4.2.1 Starting Rescue Mode

At this point, the installation program from the CD will be started and the Welcome screen (Figure 4-17) will be displayed.

```
Welcome to SuSE Linux Enterprise 10 (PPC)!

Type "install"  to start the YaST installer on this CD/DVD
Type "slp"      to start the YaST install via network
Type "rescue"   to start the rescue system on this CD/DVD

Welcome to yaboot version 10.1.14-r716.SuSE
booted from
'/vdevice/v-scsi@30000002/disk@8120000000000000:1,\suseboot\yaboot.ibm'
Enter "help" to get some basic usage information
boot:
```

Figure 4-17 Welcome screen

1. Type rescue and press **Enter** to start the Linux rescue system.

Note: The rescue selection will cause a Linux kernel to be loaded from the CD into memory. This kernel contains the necessary drivers to be able to access the physical disk adapter/devices and also contains the IBM Power Redundant Array of Independent Disks (RAID) configuration utility that is required to format the disks.

1. In the keyboard map selection screen shown in Figure 4-18, select the keyboard/language you want for the installation, then press **Enter**.

```
Choose a keyboard map.  
YaST will offer additional keyboard tables later.
```

- 1) Belgian
- 2) Ceske
- 3) Dansk
- 4) Deutsch
- 5) English (UK)
- 6) English (US)
- 7) EspaÃ±ol
- 8) FranÃ§ais
- 9) Hellenic
- 10) Italiano
- 11) Japanese
- 12) Magyar
- 13) Nederlands
- 14) Norsk
- 15) Polski
- 16) PortuguÃas Brasileiro
- 17) PortuguÃas
- 18) Russian
- 19) Slovak
- 20) Slovene
- 21) Suomi
- 22) Svensk

```
>
```

Figure 4-18 Keyboard map selection

At this point the rescue kernel will be loaded from the CD and the rescue login prompt will be displayed as shown in Figure 4-19.

```
Rescue login:
```

Figure 4-19 Rescue login prompt

2. At the login prompt, type `root` and press **Enter**.

A password is not required. The rescue prompt will be displayed as shown in Figure 4-20.

```
Rescue login: root
Rescue:~ #
```

Figure 4-20 Rescue prompt

4.2.2 Formatting the disks

To prepare the physical disks for use by Linux, perform the following steps:

1. The IBM Power RAID Configuration Utility will be used to prepare the physical disks for use by Linux. Preparation of the disk involves the formatting of the disk to a format recognized by Linux and the building of a RAID set. When the configuration utility menu is displayed as shown in Figure 4-21, type 6 (Work with disk configuration) and press **Enter** to display the current disk configuration.

```
IBM Power RAID Configuration Utility
```

```
Select one of the following:
```

1. Display hardware status
2. Work with disk arrays
3. Work with disk unit recovery
4. Work with SCSI bus configuration
5. Work with driver configuration
6. Work with disk configuration
7. Download microcode
8. Analyze log

```
Selection:
e=Exit
```

Figure 4-21 IPRConfig - Main Menu

- The Disk Configuration screen shown in Figure 4-22 will display the current configuration of the disks. Disks with a status of R/W Protected are disks that have a format unrecognizable to Linux. The disks will have to be reformatted before they can be used.

```

Change Disk Configuration

Type option, press Enter.
  1=Change Disk Configuration

OPT Name  PCI/SCSI Location      Description      Status
-----
-----
          0000:c8:01.0/0:0:4:0    Advanced Function Disk  R/W Protected
          0000:c8:01.0/0:0:5:0    Advanced Function Disk  R/W Protected
          0000:c8:01.0/0:0:6:0    Advanced Function Disk  R/W Protected
e=Exit   q=Cancel   r=Refresh   t=Toggle

```

Figure 4-22 IPR Config - Disk Configuration

Note: Disks that have been formatted for use by i5/OS will have a format that is unrecognizable to Linux (this is indicated by a R/W Protected status). Those disks will have to be reformatted using the process given here.

- Press **e** (exit) to return to the main menu of the IPR configuration utility.
- To format the disks for use in Linux, type 3 (Work with disk unit recovery) and press **Enter**.
- On the Work with Disk Unit Recovery menu (Figure 4-23), type 3 (Initialize and format disk) and press **Enter**.

```

Work with Disk Unit Recovery

Select one of the following:

  1. Concurrent add device
  2. Concurrent remove device
  3. Initialize and format disk
  4. Reclaim IOA cache storage
  5. Rebuild disk unit data
  6. Force RAID Consistency Check
  7. Work with resources containing cache battery packs

Selection:
e=Exit   q=Cancel

```

Figure 4-23 IPRConfig - Work with Disk Unit Recovery

- A list of disks will be displayed as shown in Figure 4-24. Enter 1 (select) next to each disk unit to be formatted (probably all of the disk units displayed) and then press **Enter**.

```

Select Disks for Initialize and Format

Type option, press Enter.
  1=Select

OPT Name  PCI/SCSI Location      Vendor  Product ID  Status
-----
          0000:c8:01.0/0:0:4:0    IBMAS400 XCPR036    R/W Protected
          0000:c8:01.0/0:0:5:0    IBMAS400 XCPR036    R/W Protected
          0000:c8:01.0/0:0:6:0    IBMAS400 XCPR036    R/W Protected
e=Exit   q=Cancel   t=Toggle

```

Figure 4-24 IPRConfig - Initialize and Format Disks

- A prompt appears to confirm the format (Figure 4-25). Press **c** to confirm the format.

```

Confirm Initialize and Format Disks

Press 'c' to confirm your choice for 1=Initialize and format.
  q=Return to change your choice.

OPT Name  PCI/SCSI Location      Vendor  Product ID  Status
-----
  1        0000:c8:01.0/0:0:4:0    IBMAS400 XCPR036    R/W Protected
  1        0000:c8:01.0/0:0:5:0    IBMAS400 XCPR036    R/W Protected
  1        0000:c8:01.0/0:0:6:0    IBMAS400 XCPR036    R/W Protected
c=Confirm  q=Cancel   t=Toggle

```

Figure 4-25 IPRConfig - Confirm Disk Format

- At this point the format will begin and a status screen will be displayed (Figure 4-26). The amount of time the format takes is dependent on the number of disk units and the capacity of each disk.

```

Initialize and Format Status

You selected to initialize and format a disk

                               5% Complete

```

Figure 4-26 IPRConfig - Disk Format Progress

- Once the disk format is complete the Work with Disk Unit Recovery menu will be redisplayed. Type **e** (exit).
- To verify that the disk are now usable by Linux, type **6** (Work with disk configuration) on the Power RAID Configuration Menu and press **Enter**.

11. Note that the status of the disks now shows Zeroed which indicates that the disks are formatted and ready for use by Linux as shown in Figure 4-27.

```
Change Disk Configuration

Type option, press Enter.
  1=Change Disk Configuration

OPT Name  PCI/SCSI Location      Description      Status
-----  -
-----  -
          0000:c8:01.0/0:0:4:0    Advanced Function Disk  Zeroed
          0000:c8:01.0/0:0:5:0    Advanced Function Disk  Zeroed
          0000:c8:01.0/0:0:6:0    Advanced Function Disk  Zeroed

e=Exit   q=Cancel   r=Refresh   t=Toggle
```

Figure 4-27 IPConfig - Disk Configuration

12. Type e to return to the Power RAID Configuration Utility menu.

4.2.3 Building the RAID Set

At this point, the disks are formatted and you can exit the IPRconfig utility and install Linux using each disk drive as an individual unit that Linux file systems can be placed on. Typically, most installations that use physical disk for Linux will build a RAID set across the disks to facilitate data protection and improved I/O throughput.

Perform the following steps to establish a RAID set across the disks formatted in the previous section:

1. In the Power RAID Configuration Utility menu, type 2 (Work with disk arrays) and press **Enter**.

2. From the Work with Disk Arrays menu shown in Figure 4-28, type 1 (Display disk array status) to display any disk arrays that are currently configured against the storage adapters allocated to the Linux partition.

```
Work with Disk Arrays

Select one of the following:

    1. Display disk array status
    2. Create a disk array
    3. Delete a disk array
    4. Add a device to a disk array
    5. Format device for RAID function
    6. Format device for JBOD function (512)
    7. Create a hot spare
    8. Delete a hot spare
    9. Force RAID Consistency Check

Selection:

e=Exit  q=Cancel
```

Figure 4-28 IPRConfig - Work with Disk Arrays

3. Notice that in the Display Disk Array Status screen shown in Figure 4-29, there are currently no disk arrays configured. A disk array will be built through this process.

```
Display Disk Array Status

Type option, press Enter.
  1=Display hardware resource information details

OPT Name  PCI/SCSI Location  Description  Status
-----
No devices found

e=Exit  q=Cancel  r=Refresh  t=Toggle
```

Figure 4-29 IPRConfig - Display Disk Array Status

4. Press **e** to exit the display and return to the Work with Disk Arrays menu.

5. Type 2 (Create a disk array) and press **Enter**. The Create a Disk Array menu will be displayed as shown in Figure 4-30.

```

Create a Disk Array

Select the adapter.

Type choice, press Enter.
  1=create a disk array

OPT Name  PCI/SCSI Location      Vendor  Product ID  Status
-----
          0000:c8:01.0/0:         IBM    5703001     Operational

e=Exit   q=Cancel   t=Toggle

```

Figure 4-30 IPRConfig - Create a Disk Array

6. A list of storage adapters that can have Linux RAID sets built on them is displayed. Type 1 (Create a disk array) next to the storage adapter and press **Enter**.
7. The Select Disk Units for Disk Array menu is displayed as shown in Figure 4-31. This display lists all the disks that are attached to the storage controller that have a format that is known to Linux. Type a 1 (Select) next to each disk that you want in the RAID set (probably all of the disks listed) and then press **Enter**.

```

Select Disk Units for Disk Array

Type option, press Enter.
  1=Select

OPT Name  PCI/SCSI Location      Vendor  Product ID  Status
-----
          0000:c8:01.0/0:0:4:0    IBMAS400 XCPR036     Active
          0000:c8:01.0/0:0:5:0    IBMAS400 XCPR036     Active
          0000:c8:01.0/0:0:6:0    IBMAS400 XCPR036     Active

e=Exit   q=Cancel   t=Toggle

```

Figure 4-31 IPRConfig - Select Disk Units for Disk Array

- A menu appears for specification of the parameters of the RAID set as shown in Figure 4-32. Typically, these settings are left at the default. Once you have set the parameters, press **Enter**.

```

Select Protection Level and Stripe Size

Default array configurations are shown. To change
setting hit "c" for options menu. Highlight desired
option then hit Enter

c=Change Setting

Protection Level . . . . . : RAID 5
Stripe Size . . . . . : 64 k
Queue Depth (default = 12). . . . . : 12

Press Enter to Continue

e=Exit  q=Cancel

```

Figure 4-32 IPRConfig - RAID Set Parameters

- A confirmation screen appears as shown in Figure 4-33 that lists the RAID controllers as well as the selected disks. Review the devices listed and if correct, press **Enter** to confirm the creation of the disk array.

```

Confirm Create Disk Array

Press Enter to continue.
q=Cancel to return and change your choice.

OPT Name  PCI/SCSI Location      Vendor  Product ID  Status
-----
1         0000:c8:01.0/0:         IBM     5703001     Operational
1         0000:c8:01.0/0:0:4:0    IBMAS400 XCPR036    Active
1         0000:c8:01.0/0:0:5:0    IBMAS400 XCPR036    Active
1         0000:c8:01.0/0:0:6:0    IBMAS400 XCPR036    Active

q=Cancel  t=Toggle

```

Figure 4-33 IPRConfig - Create Disk Array Confirmation Screen

10. At this point, the system will build the RAID set and a status of progress will be displayed as shown in Figure 4-34. The amount of time it takes to build the RAID set is dependent on the number and capacity of the disks in the RAID set.

```

Create Disk Array Status

You selected to create a disk array

                                2% Complete

e=Exit  Return to menu, current operations will continue.
  
```

Figure 4-34 IPRConfig - Create Disk Array Status

11. Once the build of the RAID set is complete the Work with Disk Arrays menu will be re-displayed. Type 1 (Display disk array status) and press **Enter** to verify the build of the disk array. Note that a disk array is now displayed as shown in Figure 4-35.

```

Display Disk Array Status

Type option, press Enter.
  1=Display hardware resource information details

OPT Name  PCI/SCSI Location      Description              Status
-----
sdb       0000:c8:01.0/0:255:0:0  RAID 5 Disk Array       Active
          0000:c8:01.0/0:0:4:0   RAID 5 Array Member     Active
          0000:c8:01.0/0:0:5:0   RAID 5 Array Member     Active
          0000:c8:01.0/0:0:6:0   RAID 5 Array Member     Active

e=Exit  q=Cancel  r=Refresh  t=Toggle
  
```

Figure 4-35 IPRConfig - Display Disk Array Status

12. The `sdb` in the name column is the Linux device name given to the RAID set. This is the device that the Linux installer will create and format file systems on. Press **e** (Exit) to return to the Work with Disk Arrays menu.

13. Press **e** (Exit) to return to the Power RAID Configuration Utility.

14. Press **e** (Exit) to exit the Power Raid Configuration Utility and return to the rescue mode prompt.

Archived



Creating the network support

The method for establishing networking for a Linux partition will vary based upon whether the partition has a physical or virtual network card configured. This chapter will present information about the configuration of networking for Linux on the System i platform.

5.1 Networking with physical interfaces

As discussed elsewhere in this book, direct attached (or native) network adapters can be allocated to the logical partition (LPAR) and used by the Linux operating system. Configuration of networking in such an environment is simply a case of providing the necessary network parameters (TCP/IP address, subnet mask, gateway/router) during the installation or post-installation via administration tools provided with the distribution.

5.2 Networking with virtual interfaces

This section 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 through a physical network connection in the i5/OS partition.

There are a number of methods for making a Linux partition on a virtual local area network (LAN) accessible to an external network including:

- ▶ Proxy Address Resolution Protocol (ARP)
- ▶ Network Address Translation (NAT)

5.2.1 Adding virtual Ethernet adapter to i5/OS partition

Regardless of which method is being used to forward traffic between a physical and virtual network, a virtual LAN connection has to be established for the i5/OS partition that will be forwarding the network traffic.

Note: These instructions are provided in the event that a virtual network adapter has not previously been created for the i5/OS partition. If the adapter has been created, then you can skip these steps.

Adding virtual adapter to partition profile

1. From the HMC, right-click the partition profile and select **Properties**. The properties window appears as shown in Figure 5-1. Click the **Virtual I/O Adapters** tab.

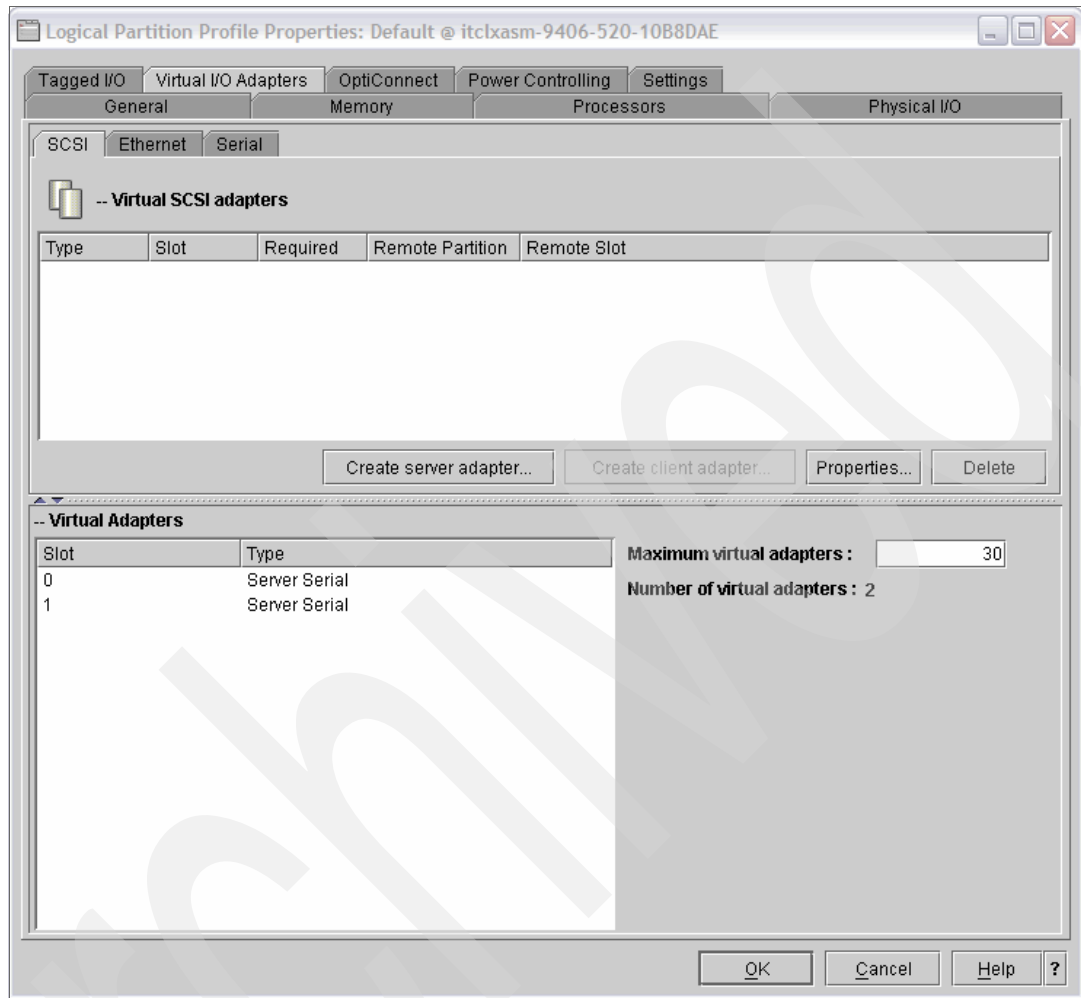


Figure 5-1 Partition Profile Properties

2. Click the **Ethernet** tab to display a list of the already configured Ethernet adapters as shown in Figure 5-2.

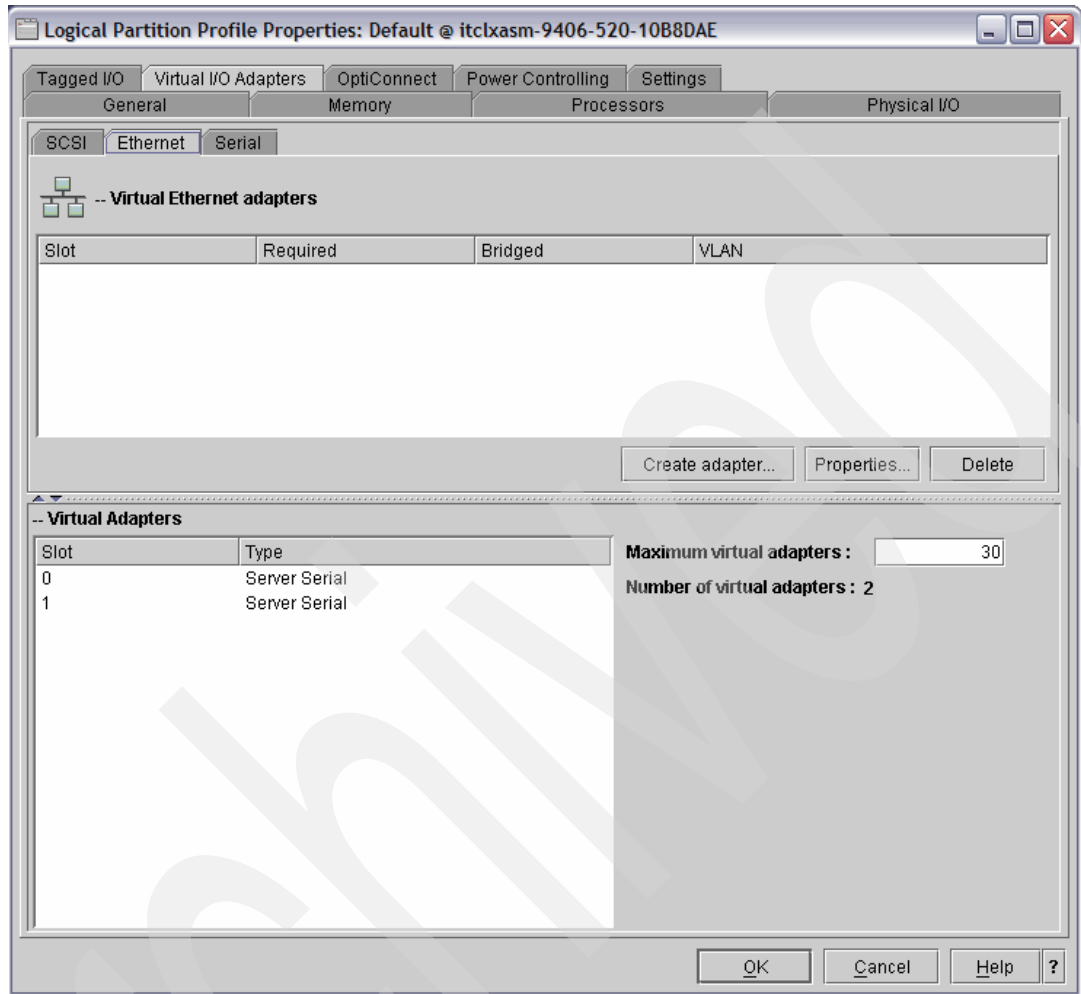


Figure 5-2 Virtual Ethernet Adapters

3. To create a new virtual Ethernet adapter, click **Create adapter**.
4. On the Virtual Ethernet Adapter properties page shown in Figure 5-3, leave the slot number as provided and enter the Virtual LAN number that corresponds to the virtual LAN that the i5/OS partition should have a connection to. This should be the same virtual LAN

that the partitions that need to have their traffic forwarded to have a virtual LAN connection configured for.

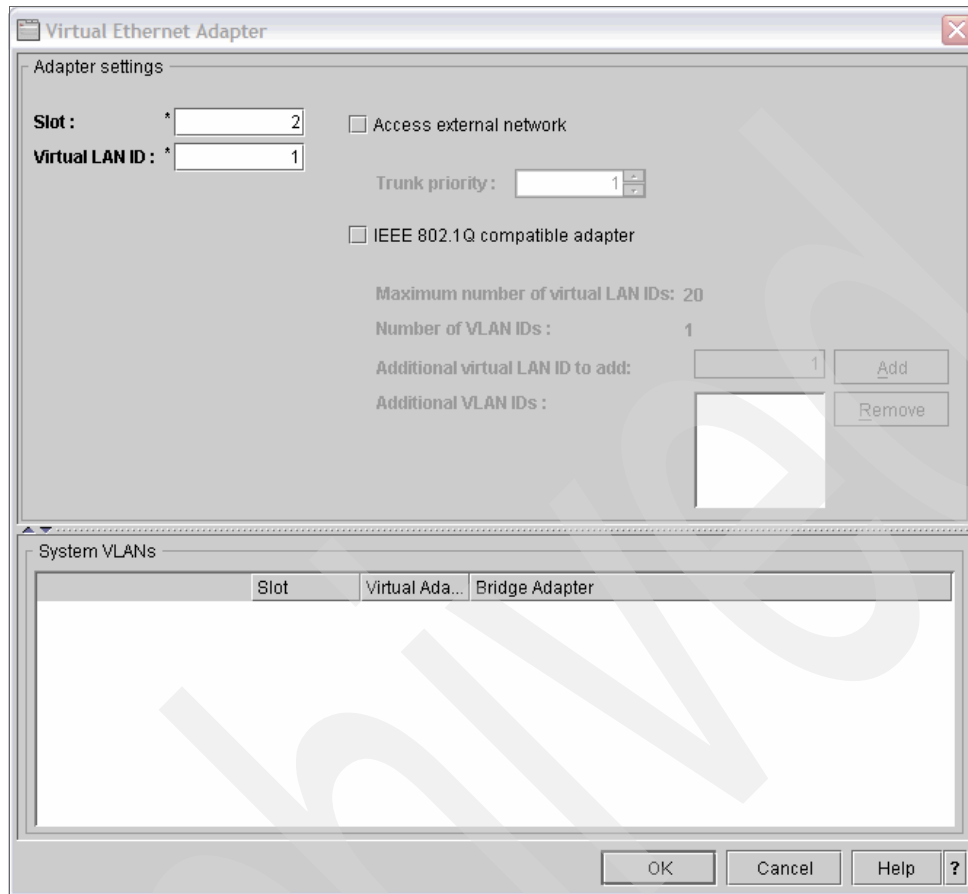


Figure 5-3 Virtual Ethernet Adapter - Properties Page

- Once the fields have been completed, click **OK**.

At this point, the virtual Ethernet adapter has been created (in the partition profile) and the partition profile page will be re-displayed as shown in Figure 5-4, reflecting the new adapter.

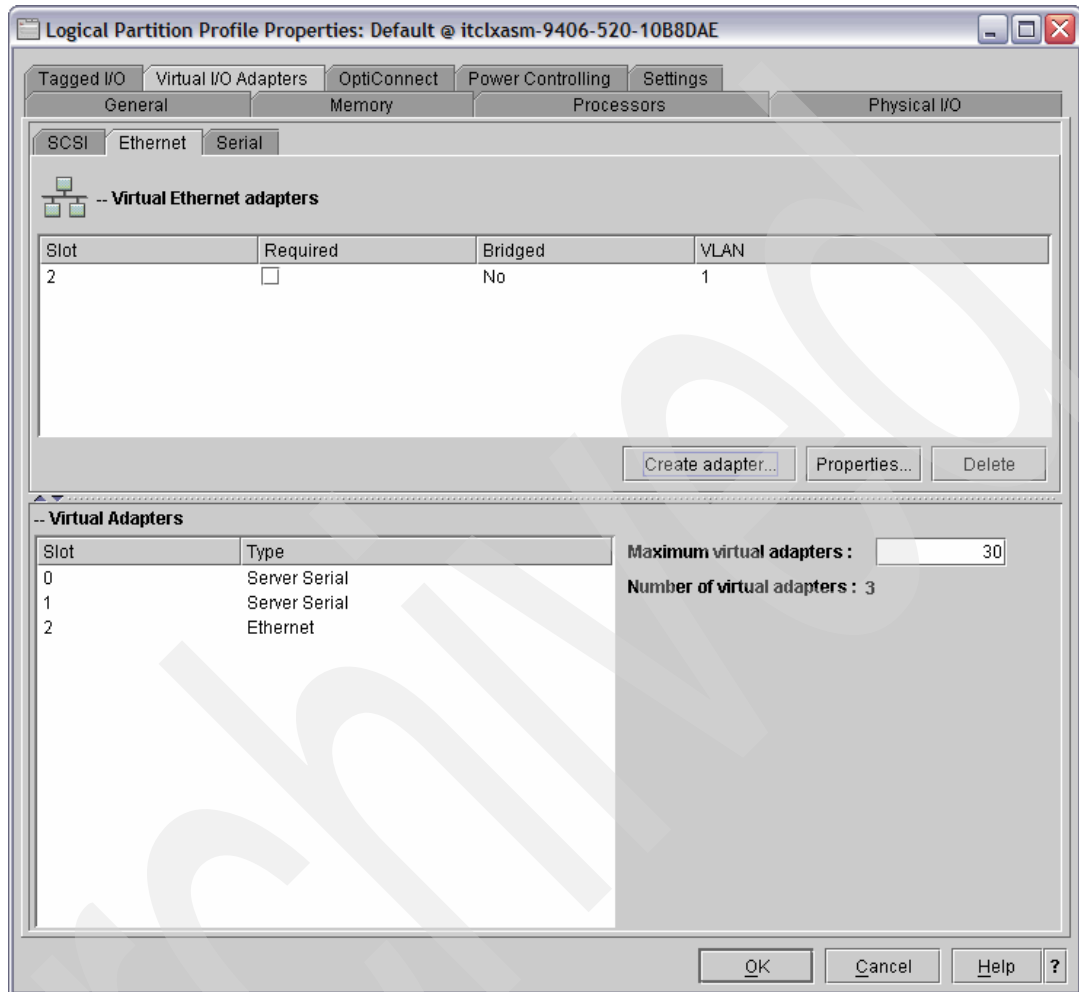


Figure 5-4 Partition Profile Properties - Virtual I/O Adapters

- Click **OK** to complete the update to the partition profile.

Dynamic addition of the adapter to the i5/OS partition

Adding the adapter to the partition profile means that the next time the partition is activated from the Hardware Management Console (HMC) it will have the newly-configured adapter. To use the adapter now, it also needs to be added to the partition through the use of Dynamic Logical Partitioning.

- From the HMC, right-click the i5/OS partition (not the partition profile) and select **Dynamic Logical Partitioning** → **Virtual Adapter Resources** → **Add/Remove**.

The Dynamic Logical Partitioning screen for virtual adapters appears as shown in Figure 5-5.

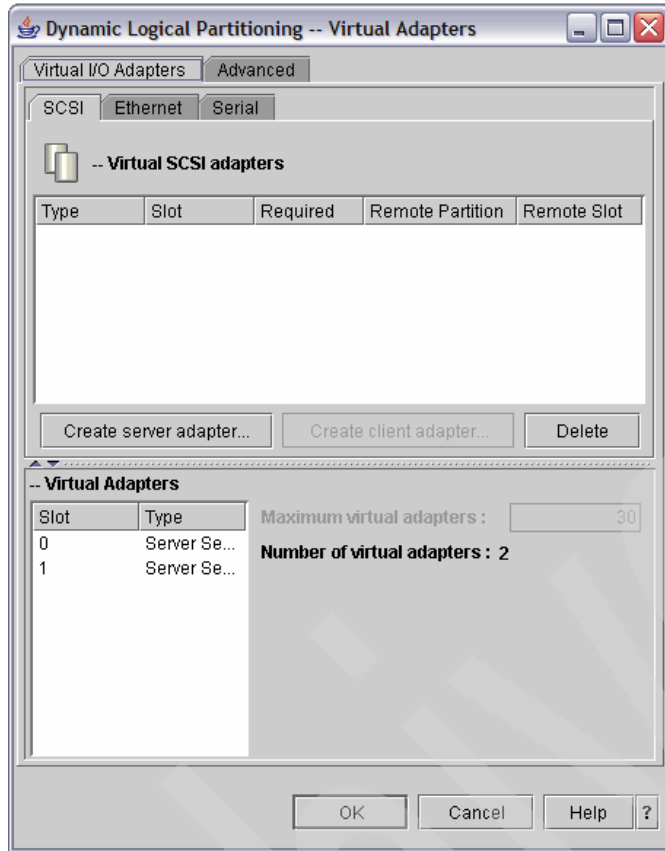


Figure 5-5 Dynamic Logical Partitioning - Virtual Adapters

2. Click the **Ethernet** tab and then click **Create Adapter**.

3. Complete the fields of the Virtual Ethernet Adapter properties as shown in Figure 5-6 and click **OK**.

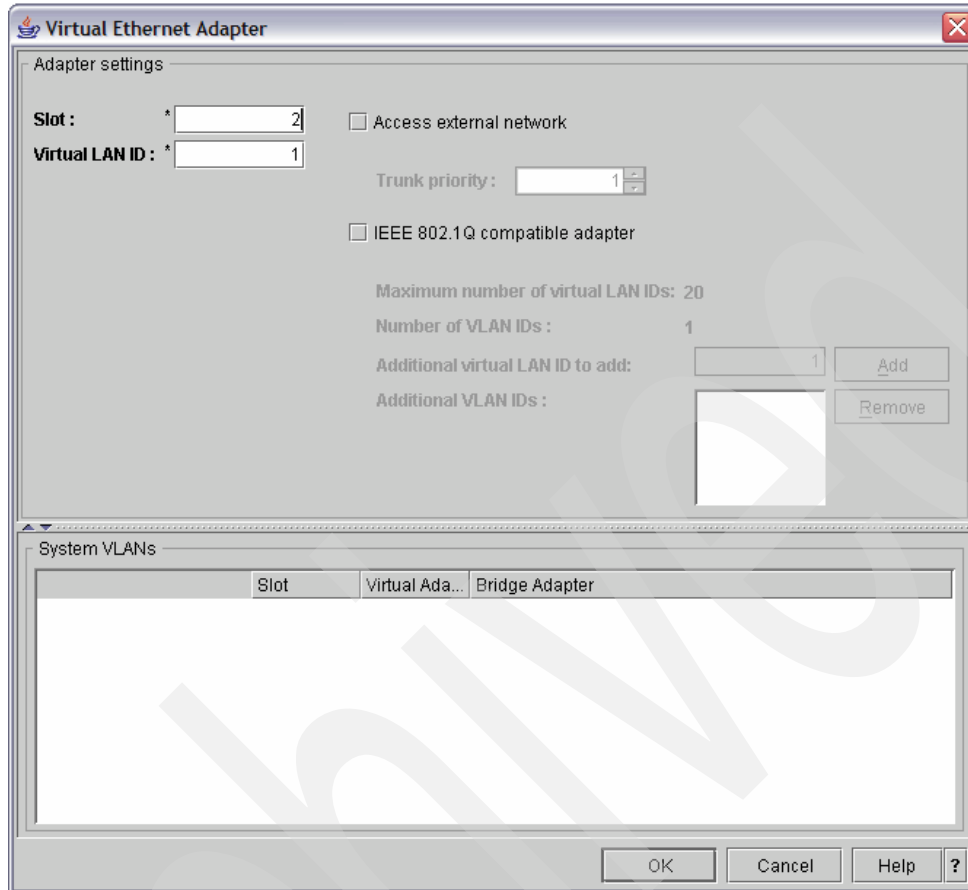


Figure 5-6 Virtual Ethernet Adapter - Properties

The Virtual Ethernet adapter is now displayed in the list of adapters for the partition as shown in Figure 5-7.

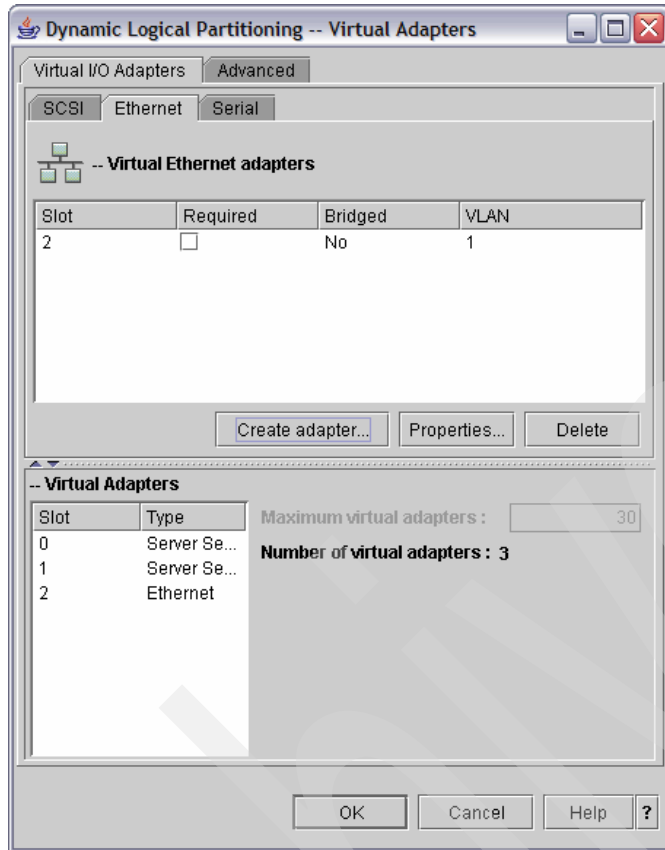


Figure 5-7 Dynamic Logical Partitioning - Virtual Adapters

4. Click **OK** to update the partition. At this point the adapter is added to the running partition as shown in Figure 5-8.




Figure 5-8 Dynamic Logical Partitioning - Working dialog box

5.2.2 Creating the Ethernet Line Description in i5/OS

Once the Virtual Ethernet adapter has been added to the i5/OS partition an Ethernet Line Description will need to be created in i5/OS. A review of the communication hardware resources will help to identify the resource for the line description.

1. From a 5250 session, issue the Work with Hardware Resources (WRKHDWRSC) command to review the communication resources as shown in Example 5-1.

Example 5-1 Work with Hardware Resources - command syntax

WRKHDWRSC *CMN

The above command will result in a display of the communication resources allocated to the i5/OS system as shown in Figure 5-9.

```
Work with Communication Resources                               System:  S10B8DAE
Type options, press Enter.
  5=Work with configuration descriptions  7=Display resource detail

Opt Resource      Type Status      Text
CMB01          2844 Operational Combined function IOP
  LIN14        2793 Operational Comm Adapter
    CMN23      2793 Operational Comm Port
    CMN24      2793 Operational Comm Port
  LIN01        2849 Operational LAN Adapter
    CMN01      2849 Operational Ethernet Port
CMB03          268C Operational Combined function IOP
  LIN02        6B03 Operational Comm Adapter
    CMN02      6B03 Operational Comm Port
  LIN03        6B03 Operational Comm Adapter
    CMN03      6B03 Operational Comm Port
  LIN04        268C Operational LAN Adapter
    CMN36      268C Operational Ethernet Port
CMB02          2844 Operational Combined function IOP
  LIN15        2849 Operational LAN Adapter

F3=Exit  F5=Refresh  F6=Print  F12=Cancel

More...
```

Figure 5-9 Work with Communication Resources

The virtual Ethernet adapter will be a CMNxx resource with a type of 268C.

- To identify the specific resource, use the display resource detail (7) option on resources that meet the aforementioned criteria. The details for the selected resource will be displayed as shown in Figure 5-10.

```

                                Display Resource Detail
                                System:  S10B8DAE

Resource name . . . . . : CMN36
Text . . . . . : Ethernet Port
Type-model . . . . . : 268C-001
Serial number . . . . . : 00-00000
Part number . . . . . :

Location:  U9406.520.10B8DAE-V1-C2-T1

Logical address:
SPD bus:
  System bus                255
  System board              0
  System card                0
More...

Press Enter to continue.

F3=Exit  F5=Refresh  F6=Print  F12=Cancel

```

Figure 5-10 Resource detail

The Location field indicates the mapping of the i5/OS resource to the LPAR resource. In this example, the V1 indicates the partition number (1) and the C2 indicates the slot number (2) of the virtual Ethernet adapter.

- Once the resource has been identified, the Ethernet line description can be created. Issue the Create Line Description Ethernet (CRTLINETH) command and press **F4** (as shown in Example 5-2).

Example 5-2 Create Line Description (Ethernet) - syntax

```
CRTLINETH
```

The Create Line Description definition screen appears as shown in Figure 5-11.

```

Create Line Desc (Ethernet) (CRTLINETH)

Type choices, press Enter.

Line description . . . . .
Resource name . . . . .

Name
Name, *NWID, *NWS

Bottom

F3=Exit  F4=Prompt  F5=Refresh  F10=Additional parameters  F12=Cancel
F13=How to use this display  F24=More keys

```

Figure 5-11 Create Line Description (Ethernet) definition screen

4. Enter a name of your choosing for the Line description field. It is recommended that the name reflects that this is a virtual adapter as well as the virtual LAN number. As an example, for a virtual LAN adapter on virtual LAN 1, a name like VRTETH01 may be appropriate.
5. Enter the communication resource (CMNxx) identified as the resource that maps to the virtual Ethernet adapter for the Resource name field.
6. Once you have completed the first two fields, press **Enter** to display additional fields as shown in Figure 5-12.

```

Create Line Desc (Ethernet) (CRTLINETH)

Type choices, press Enter.

Line description . . . . . > VRTETH01
Resource name . . . . . > CMN36
Online at IPL . . . . . *YES
Vary on wait . . . . . *NOWAIT
Local adapter address . . . . . *ADPT
Exchange identifier . . . . . *SYSGEN
Ethernet standard . . . . . *ALL
Line speed . . . . . 10M
Duplex . . . . . *HALF

Name
Name, *NWID, *NWS
*YES, *NO
*NOWAIT, 15-180 seconds
020000000000-FEFFFFFFFFF...
05600000-056FFFFF, *SYSGEN
*ETHV2, *IEEE8023, *ALL
10M, 100M, 1G, *AUTO
*HALF, *FULL, *AUTO

Bottom

F3=Exit  F4=Prompt  F5=Refresh  F10=Additional parameters  F12=Cancel
F13=How to use this display  F24=More keys

```

Figure 5-12 Create Line Description (Ethernet) definition screen - additional fields

7. The setting of the Line speed field will be dependent on the line speed of the physical interface. The line speed should be set to match the line speed of any associated physical interface to avoid fragmentation of the network packets.
8. The setting of 'Duplex' is typically set to *FULL.
9. When you have completed the fields, press **Enter**.
10. Additional fields will be displayed. Press **Enter** again to complete the definition.
11. Once the line description has been created, vary it on as shown in Example 5-3.

Example 5-3 VRYCFG command syntax

```
VRYCFG CFGOBJ(VRTETH01) CFGTYPE(*LIN) STATUS(*ON)
```

5.2.3 Enabling datagram forwarding

For all three approaches (Proxy ARP, NAT, and TCP routing) the TCP/IP attribute datagram forwarding should be enabled. Datagram forwarding allows network packets to be forwarded between network interfaces on the system.

To ensure that datagram forwarding is enabled, issue the command shown in Example 5-4.

Example 5-4 CHGTCPA command syntax

```
CHGTCPA IPDTGFWD(*YES)
```

5.2.4 Proxy ARP

Proxy ARP is used to build a subnet within the overall network that the i5/OS partition's physical network interface is connected to. Essentially, i5/OS becomes a router for a subnet

of an address that will be assigned to the virtual LAN. An example environment that will use Proxy ARP is shown in Figure 5-13.

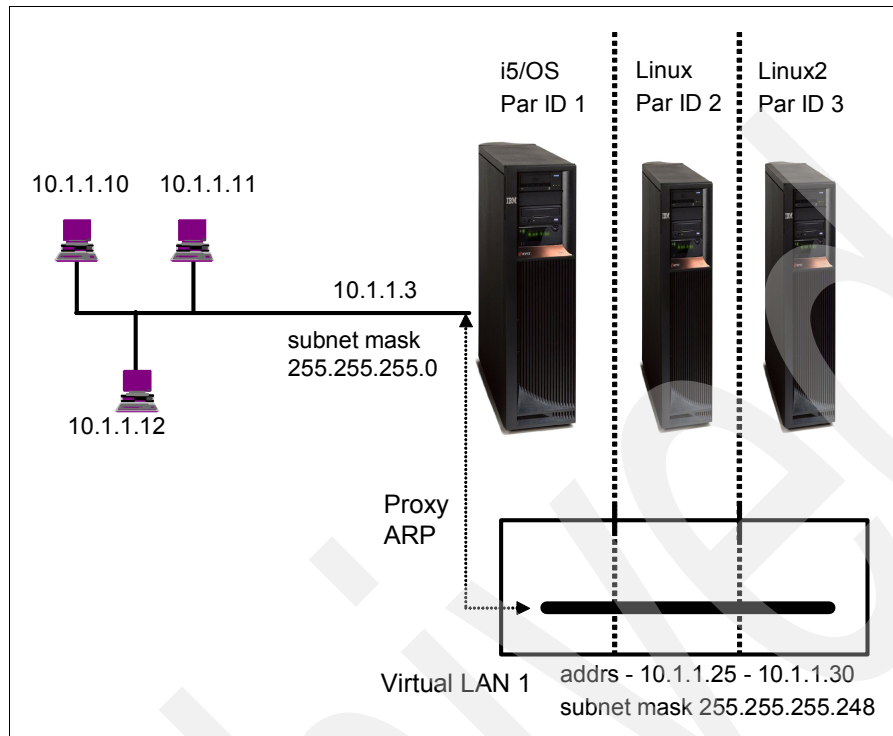


Figure 5-13 Proxy ARP example

In this example, a subnet is established that ranges from 10.1.1.24 to 10.1.1.31. This range of addresses is assigned to the virtual LAN and is associated with an interface on the physical LAN (in this case, 10.1.1.3) to act as the router for the subnet. The range of addresses for the subnet is determined by an address in the network along with the subnet mask. There are a number of rules that the subnet must meet:

- ▶ The size of the subnet must be a power of 2.
- ▶ The first address of the subnet cannot be assigned to a partition. The first address is the network address.
- ▶ The last address of the subnet cannot be assigned to partition. The last address is the broadcast address.
- ▶ The subnet must be contained wholly within the overall network that the associated i5/OS physical network is attached to.

Highlights of the Proxy ARP environment are:

- ▶ Traffic intended for partitions on the virtual LAN is *routed* to the i5/OS physical network interface.
- ▶ Traffic is then re-broadcast on the virtual LAN via the i5/OS connection on that LAN.
- ▶ The Linux operating system running in the partition “sees” the network traffic just like in any other network structure.
- ▶ Domain Name System (DNS) entries point to the address on the virtual LAN.
- ▶ Virtual LAN address range must be a subset of the physical LAN addresses.
- ▶ The virtual LAN MTU (frame size) should be equal to or less than that of the physical LAN.

Perform the following steps to set up the components in i5/OS to support Proxy ARP.

Note: These steps assume that the i5/OS partition and Linux partition have already been defined with virtual Ethernet adapters on the same virtual LAN and that the Ethernet Line description has been created and varied on.

Determining range of addresses to proxy

As indicated earlier, Proxy ARP is the establishment of a range of addresses that will be associated with a single physical interface. A subnet calculator is a good tool to help determine a range of addresses for the subnet. One example is the subnet calculator available on the Web at:

http://www.wildpackets.com/products/free_utilities/ipsubnetcalc/overview

Building on the example shown earlier, the subnet calculator can be used to establish the range of addresses as shown in Figure 5-14.

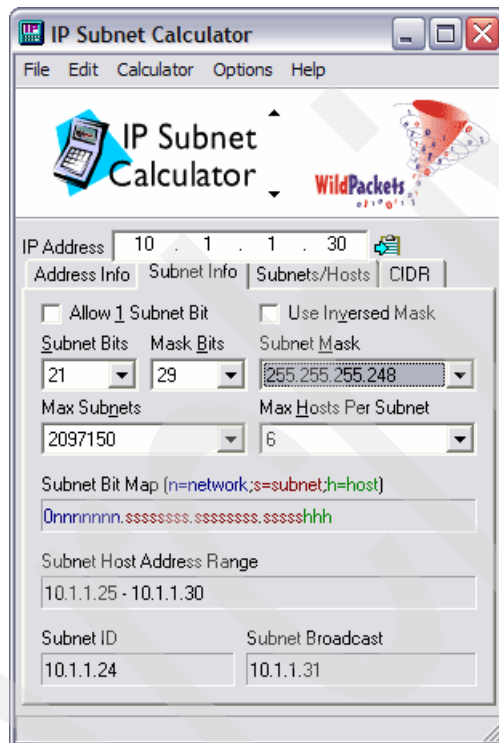


Figure 5-14 Subnet Calculator

In this example, an address was specified that should fall within the range of addresses for the subnet. Then the subnet mask was used to indicate the size of the subnet (that is, 255.255.255.248 indicates a subnet of 8 addresses.)

By specifying the above values, a subnet is defined with the following characteristics:

- ▶ Network Address 10.1.1.24
- ▶ Range of usable addresses 10.1.1.25 - 10.1.1.30
- ▶ Broadcast Address 10.1.1.31

Now that the subnet range has been defined, the next step is to create the TCP/IP interface for i5/OS on the virtual LAN.

Creating TCP/IP interface for virtual LAN adapter

After you have created the Ethernet Line Description and defined the subnet, you are ready to define the TCP/IP interface for i5/OS on the virtual LAN by performing the following steps.

Note: Typically, i5/OS will be assigned the first usable address in the subnet.

1. Issue the Add TCP/IP Interface command to create the TCP/IP interface as shown in Example 5-5.

Example 5-5 ADDTCPIFC syntax

```
ADDTCPIFC INTNETADR('10.1.1.25') LIND(VRTETH01) SUBNETMASK('255.255.255.248')
LCLIFC('10.1.1.3')
```

In Example 5-5:

- Value of INTNETADR is the address of the i5/OS interface on the virtual LAN.
- Value for LIND is the name of the Ethernet Line Description for the virtual LAN.
- Value for SUBNETMASK is the size of the subnet.
- Value for LCLIFC is the address of the physical network interface in i5/OS that will forward traffic to the virtual LAN.

You could also prompt the ADDTCPIFC command (that is, press **F4**) and complete the fields shown in Figure 5-15.

Add TCP/IP Interface (ADDTCPIFC)

Type choices, press Enter.

Internet address	10.1.1.25	
Line description	VRTETH01	Name, *LOOPBACK...
Subnet mask	255.255.255.248	
Alias name	*NONE	
Associated local interface . . .	10.1.1.3	
Type of service	*NORMAL	*MINDELAY, *MAXTHRPUT...
Maximum transmission unit . . .	*LIND	576-16388, *LIND
Autostart	*YES	*YES, *NO
PVC logical channel identifier		001-FFF
	+ for more values	
X.25 idle circuit timeout . . .	60	1-600
X.25 maximum virtual circuits .	64	0-64
X.25 DDN interface	*NO	*YES, *NO
TRLAN bit sequencing	*MSB	*MSB, *LSB

Bottom

F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

Figure 5-15 Add TCP/IP interface

- Once the TCP/IP interface has been created it can be varied on with the Start TCP/IP Interface command shown in Example 5-6.

Example 5-6 STRTCP syntax

```
STRTCPIFC INTNETADR('10.1.1.25')
```

Linux network attributes

The i5/OS setup to support Proxy ARP is now complete. When Linux is installed, the network parameters will be configured as follows:

- ▶ **TCP/IP address:** an available address within the usable range of addresses for the subnet
- ▶ **Subnet mask:** the size of the subnet (should be the same value as used for the TCP/IP interface in i5/OS for the subnet)
- ▶ **Router/Gateway:** the address of the i5/OS interface on the virtual LAN.

Figure 5-16 represents a completely defined Proxy ARP environment.




Figure 5-16 Proxy ARP - example

In the above example, a range of addresses (10.1.1.25 to 10.1.1.30) is assigned to the virtual LAN. Additionally, the address for i5/OS on the virtual LAN (10.1.1.25) is associated with the i5/OS's address on the physical LAN (10.1.1.3) which provides the forwarding of traffic between the virtual and physical network.

SUSE Linux Enterprise 10

As indicated above, configuration of networking can be accomplished during the installation process and this method is documented in Chapter 6, "Linux installation" on page 177. If you have to set or modify the network settings after the installation, you can use the yast tool.

- In Linux, start the YaST Control Center as shown in Example 5-7 and Figure 5-17.

Example 5-7 Starting the yast tool

```
yast
```



Figure 5-17 YaST Control Center

2. Use the arrow keys to select **Network Devices** → **Network Card** and then press **Enter** as shown in Figure 5-18.

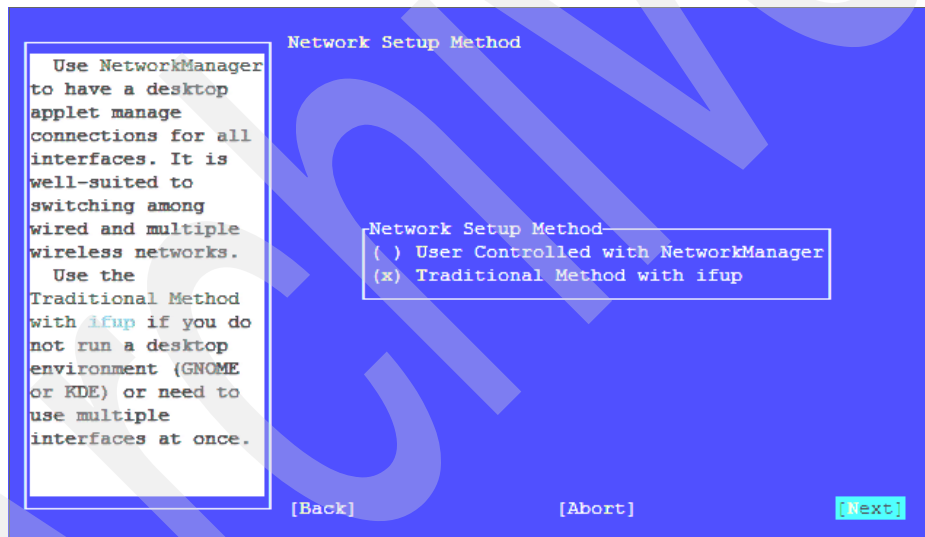


Figure 5-18 Network Setup Method

3. In the Network Setup Method screen, leave the Network Setup Method set at Traditional Method with ifup and press the **Enter** key as shown in Figure 5-19.

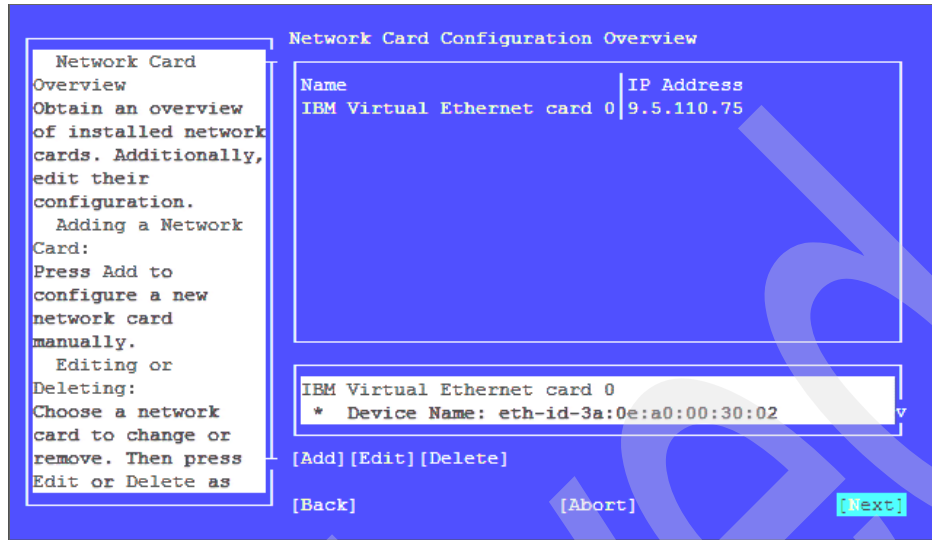


Figure 5-19 Network Card Configuration Overview

The Network Card Configuration Overview shows the network adapters that are detected in the system.

4. Press the **Tab** key until [Edit] is highlighted and then press **Enter** as shown in Figure 5-20.

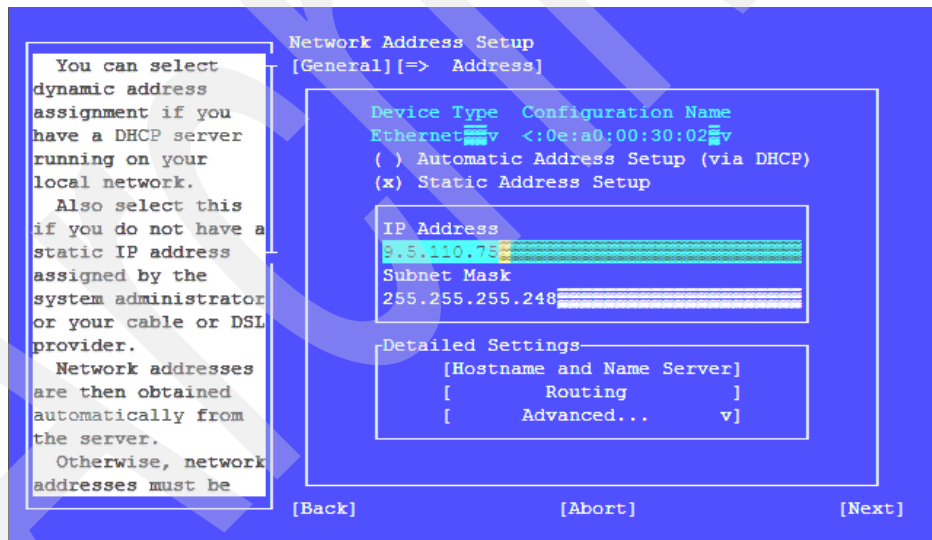


Figure 5-20 Network Address Setup

The Network Address Setup window shows the current settings for the network adapter.

5. Press the **Tab** key to position the cursor in the appropriate fields for the parameters and enter values to match the desired network settings.

- To set the route/gateway, press the **Tab** key until [Routing] is highlighted and then press **Enter** as shown in Figure 5-21.

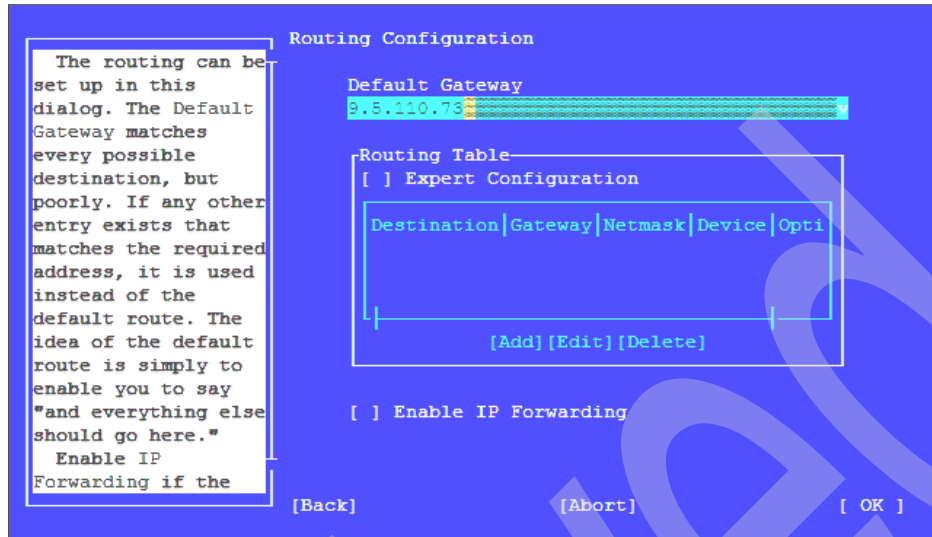


Figure 5-21 Routing Configuration

The Route Configuration window shows the current router settings.

- Press the **Tab** key to position the cursor in the Default Gateway field and set it to the appropriate value.
 - Press the **Tab** key until OK is highlighted and then press **Enter**.
 - In the Network Address Setup screen, press the **Tab** key until [Next] is highlighted and then press **Enter**.
 - In the Network Card Configuration Overview screen, press the **Tab** key until [Next] is highlighted and then press **Enter**.
- At this point the configuration settings are saved and the network adapters are restarted with the new settings.
- Back on the YaST Control Center screen, press the **Tab** key until [Quit] is highlighted and then press **Enter**.

RedHat Enterprise Linux 4

As indicated before, configuration of networking can be accomplished during the installation process and this method is documented in the Chapter 6, "Linux installation" on page 177. Perform the following steps:

- In Linux, start the text version of the network configuration program as shown in Example 5-8 and Figure 5-22.

Example 5-8 Start the network configuration program

```
system-config-network-tui
```

```

+----+ Network Configuration +----+
| What do you want to configure? |
| Ethernet                       |
| Modem                          |
| ISDN                           |
|                                 |
| +-----+ | +-----+ | | | | |
| | Configure | | | Exit | |
| +-----+ | | +-----+ |
|                                 |
+-----+
<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen

```

Figure 5-22 Network Configuration

2. In the Network Configuration screen, use the arrow key to position the cursor on Ethernet.
3. Press the **Tab** key until the cursor is positioned over [Configure] and press **Enter** as shown in Figure 5-23.

```

+-----+ Ethernet Configuration +-----+
| Name           eth0 |
| Device         eth0 |
| Use DHCP       [ ] |
| Static IP      9.5.110.76 |
| Netmask        255.255.255.248 |
| Default gateway IP 9.5.110.73 |
|                 |
| +-----+ | +-----+ | | | | |
| | Ok | | | Cancel | |
| +-----+ | | +-----+ |
|                 |
+-----+
<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen

```

Figure 5-23 Ethernet Configuration

4. In the Ethernet Configuration screen, use the **Tab** key to position the cursor in the various parameter fields and enter the appropriate values.
5. Once the network parameter values are complete, press **Tab** until the cursor is positioned over [Ok] and press **Enter**.

6. In the Network Configuration screen, press the **Tab** key until the cursor is positioned over [Exit] and press **Enter**.
7. At this point, the network settings have been saved. You can use the system startup script for network to restart the network stack with the new settings as shown in Example 5-9.

Example 5-9 Restarting the network stack

```
service network restart
```

5.2.5 Network Address Translation

Network Address Translation (NAT) is used to map a public address to a private address. With NAT, a private network is established between the Linux and i5/OS partitions and a mapping of public to private address is established. Figure 5-24 shows an example environment that will use NAT.




Figure 5-24 Network Address Translation example

In this example, the Linux partitions are assigned public address (10.1.101 and 10.1.154) which are mapped to private address (172.168.1.2 and 172.168.1.3) through the NAT support within i5/OS.

Note: The i5/OS partition's virtual LAN interface will be assigned an address within the private network range (normally the first usable address), however, the address will not be mapped to a public address.

Highlights of the NAT environment can be summarized as follows:

- ▶ Addresses for each of the LPARs are defined in the NATing (that is, i5/OS) partition. This address becomes an additional address on the physical adapter.

- ▶ Static NAT rules are defined to map the public (physical) address to the private (virtual LAN) address.
- ▶ DNS entries will point to the public (physical) addresses.
- ▶ The physical network adapter replies to the network request.
- ▶ The IP address in the IP header is rewritten (mangled) to reflect the private address.
- ▶ The routing table is checked and the network traffic is forwarded to the destination.
- ▶ Output traffic from the private LAN is routed through the NATing (that is, i5/OS) partition.

The following sections describe the setup of the components in i5/OS to support NAT.

Note: These steps assume that the i5/OS partition and Linux partition have already been defined with virtual ethernet adapters on the same virtual LAN. Additionally, these steps assume that the Ethernet Line Description for i5/OS's virtual LAN connection has been created and varied on and that the TCP/IP attribute *Datagram Forwarding* has been enabled.

Creating TCP/IP interface for i5/OS on virtual LAN

A TCP/IP interface needs to be created for the i5/OS partition on the virtual LAN. This is the interface that will be used to forward traffic between the virtual and physical networks. Four aspects of this TCP/IP interface to keep in mind:

- ▶ The TCP/IP address assigned to the interface needs to be one of the addresses in the private network. Typically the address assigned will be the first usable address in the private network.
- ▶ Unlike the private addresses for the Linux partitions, the private address assigned to the i5/OS partition will not be mapped to a public address.
- ▶ The i5/OS address on the virtual LAN will be used by the Linux partitions as their router/gateway for the Linux operating system network configuration.
- ▶ The TCP/IP interface on the virtual LAN will be associated with the TCP/IP interface on the physical LAN.

1. Issue the Add TCP/IP Interface command to create the TCP/IP interface as shown in Example 5-10.

Example 5-10 ADDTCPIFIC command syntax

```
ADDTCPIFIC INTNETADR('172.168.1.1') LIND(VRTETH01) SUBNETMASK('255.255.255.0')
LCLIFC('10.1.1.3')
```

In Example 5-10:

- Value of INETADR is the address of the i5/OS interface on the virtual LAN.
- Value for LIND is the name of the Ethernet Line Description on the virtual LAN.
- Value for SUBNETMASK is the size of the subnet.
- Value for LCLIFC is the address of the physical network interface in i5/OS that will forward traffic to the virtual LAN.

2. Once the TCP/IP interface has been created it can be varied on with the Start TCP/IP Interface command as shown in Example 5-11.

Example 5-11 STRTCPIFIC command syntax

```
STRTCPIFIC INTNETADR('172.168.1.1')
```

Creating TCP/IP interfaces for public addresses

Each partition that will have their network traffic NATed requires a TCP/IP interface created for its public address against the i5/OS physical Ethernet Line Description that will handle the network traffic. To create TCP/IP interfaces for public addresses, perform the following steps:

1. For each public address, use the Add TCP/IP Interface command to create the TCP/IP interface as shown in Example 5-12.

Example 5-12 ADDTCPIFC syntax

```
ADDTCPIFC INETADR('10.1.1.101') LIND(ETHLINC05) SUBNETMASK('255.255.255.0')
```

In Example 5-12:

- The value for INETADR is the public TCP/IP address for the partition.
 - The value for LIND is the name of the Ethernet Line Description that will handle the network traffic for the partition.
 - The value for SUBNETMASK is the subnet mask of the public network.
2. Once the TCP/IP interface has been created, start it as shown in Example 5-13.

Example 5-13 STRTCPIFC command syntax

```
STRTCPIFC INETADR('10.1.1.101')
```

Note: The above steps should be completed for each public address that will be NATed.

Creating NAT rules

The NAT rules (the rules used to translate the private addresses to public addresses and the public address to the private addresses) are created through the iSeries Navigator interface in the following steps:

1. The rules are created through use of the Rules Editor feature of iSeries Navigator. Navigate to the **Network** → **IP Policies** section of iSeries Navigator and then right-click **Packet Rules** and select **Rules Editor**.
2. In the welcome screen shown in Figure 5-25, select **Create a new packet rules file** and then click **OK**.

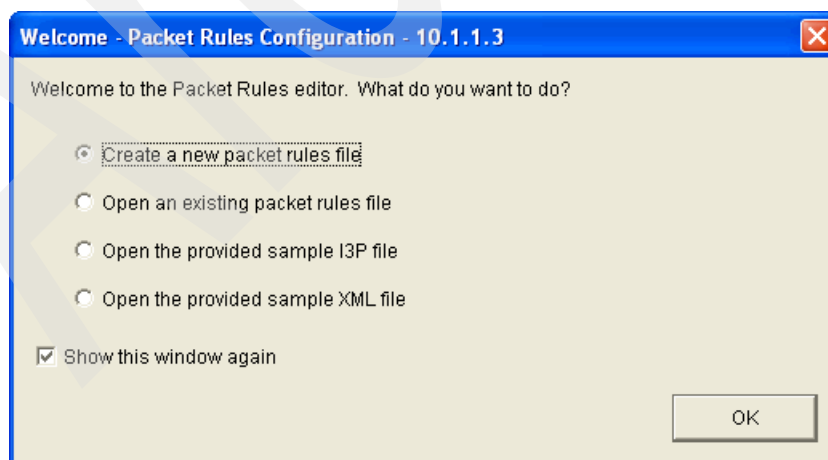


Figure 5-25 Packet Rules Configuration - Welcome Screen

3. In the Getting Started window shown in Figure 5-26, select **OK** to display the Packet Rules Editor window shown in Figure 5-27.

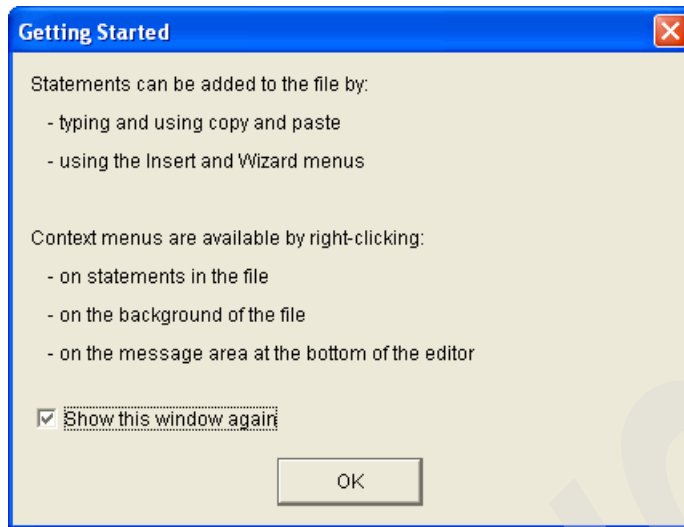


Figure 5-26 Packet Rules Editor - Getting Started




Figure 5-27 Packet Rules Editor

4. From the Packet Rules Editor select **Wizards** from the menu bar and then select **Address Translation**.

5. In the Address Translation Wizard Welcome screen shown in Figure 5-28, click **Next**.

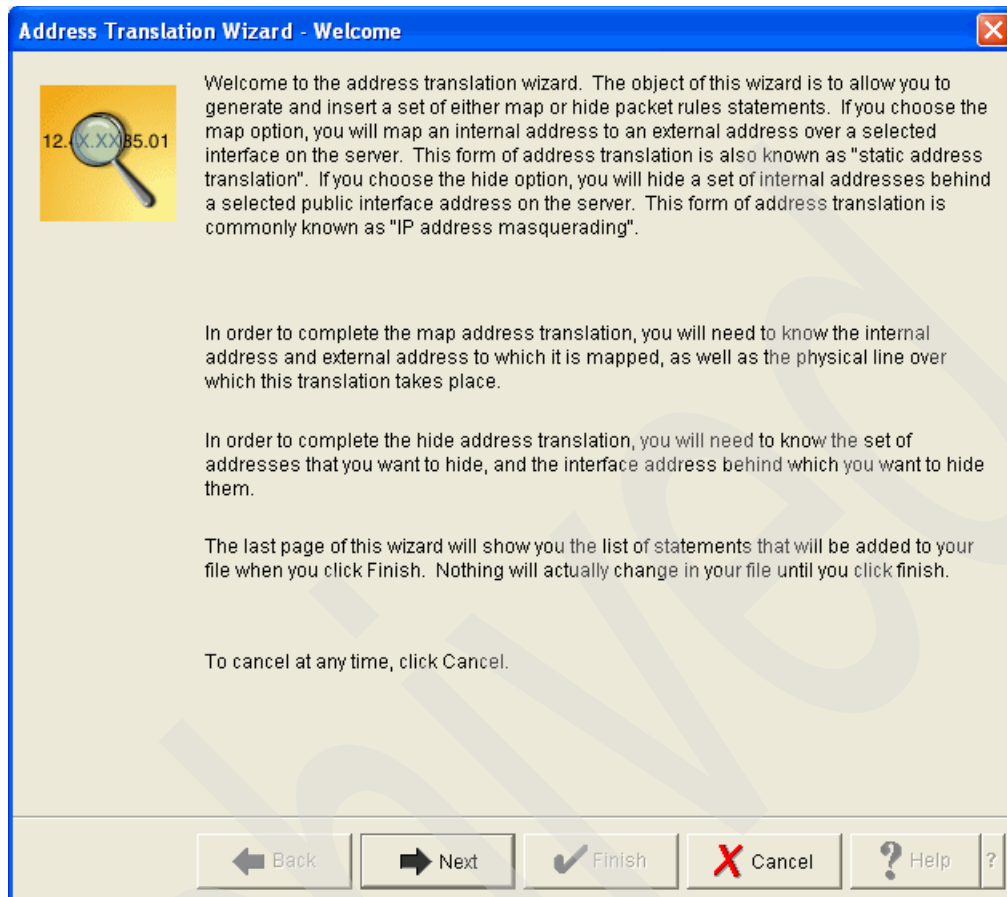


Figure 5-28 Address Translation Wizard - Welcome screen

6. For Network Address Translation, select **Map address translation** on the Address Translation Selection screen shown in Figure 5-29 and click **Next**.

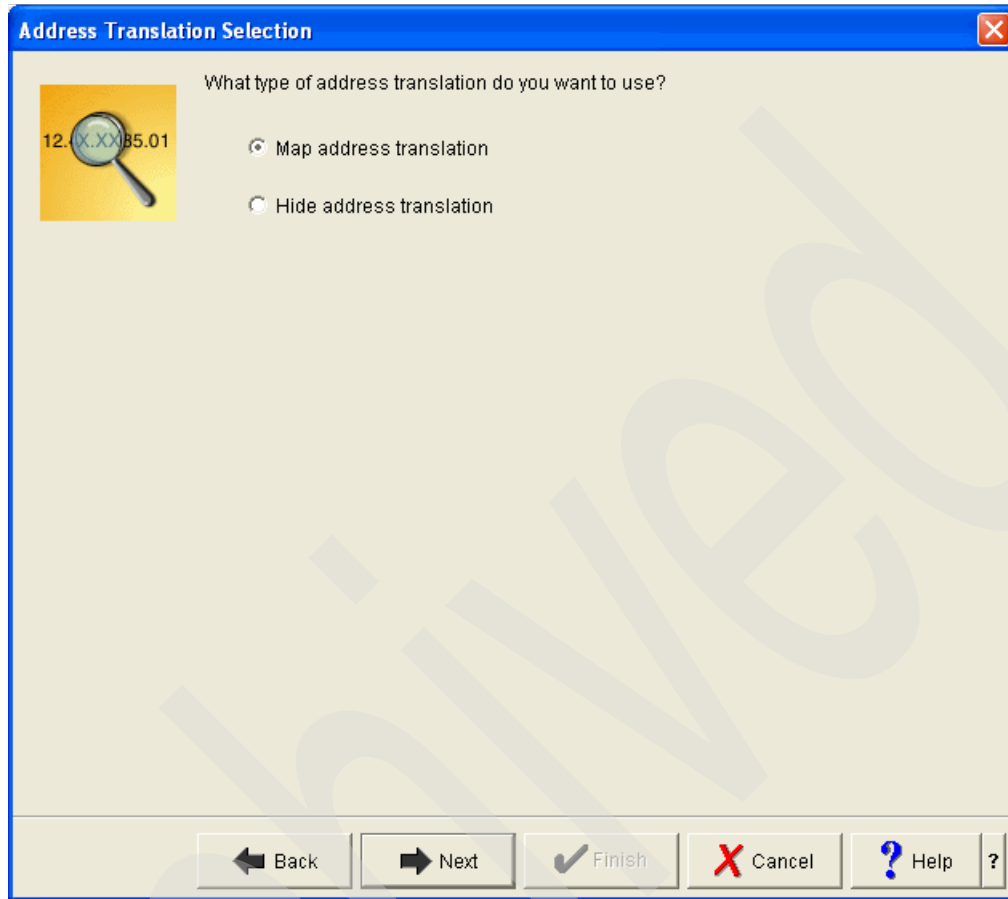


Figure 5-29 Address translation selection

7. In the Private Address screen shown in Figure 5-30, enter the IP address on the virtual LAN for the Linux partition. Once the IP address has been specified, click **Next**.

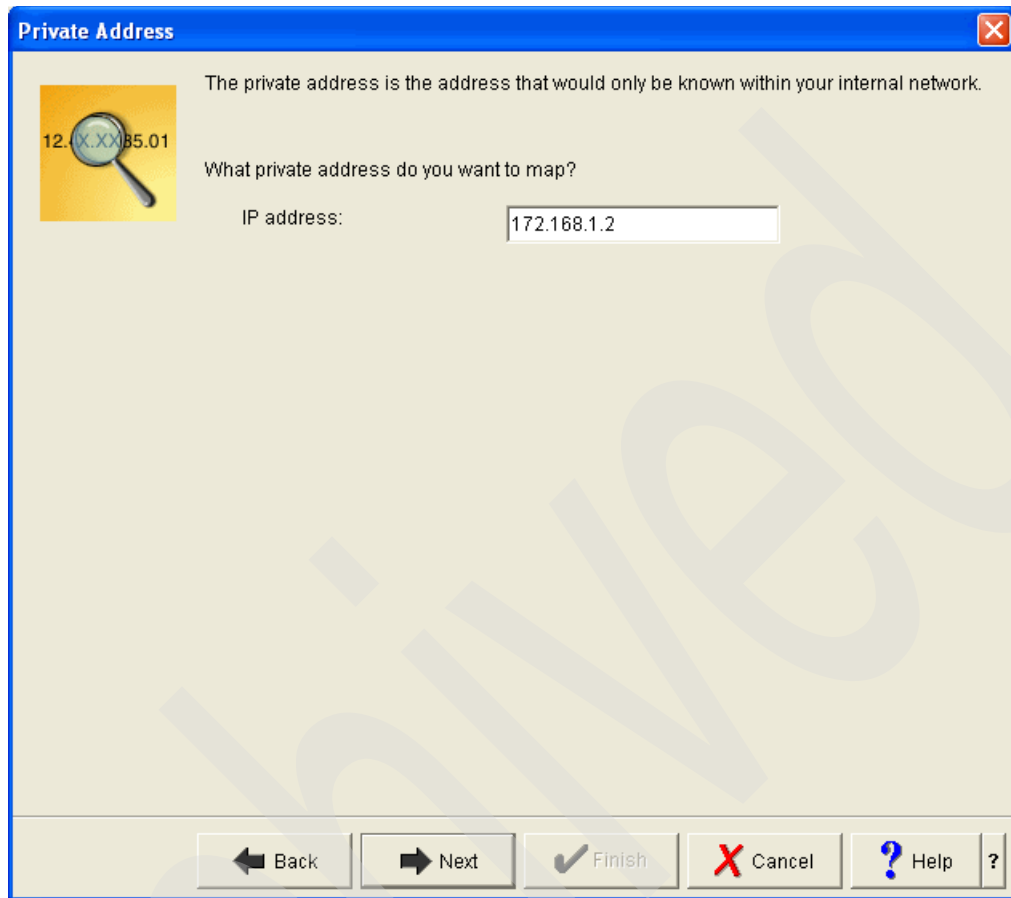


Figure 5-30 Packet Rules Editor - Private Address

8. In the Public Address screen shown in Figure 5-31, enter the IP address that systems outside of the virtual LAN will use to access the Linux partition (that is, the partitions public address). Once the IP address has been specified, click **Next**.

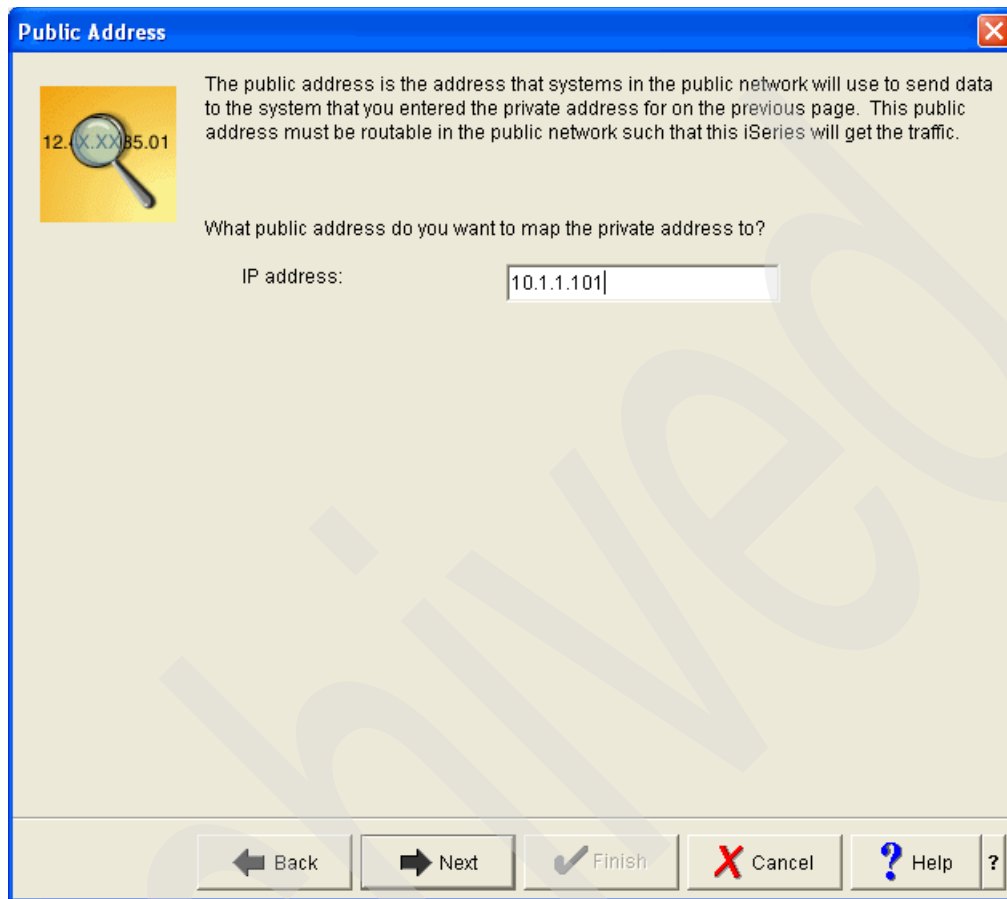


Figure 5-31 Packet Rules Editor - public address

9. In the Line screen shown in Figure 5-32, select the interface that the traffic will be seen on. The interface is the i5/OS line description for the physical network interface that is connected to the external LAN. Once the interface has been selected, click **Next**.

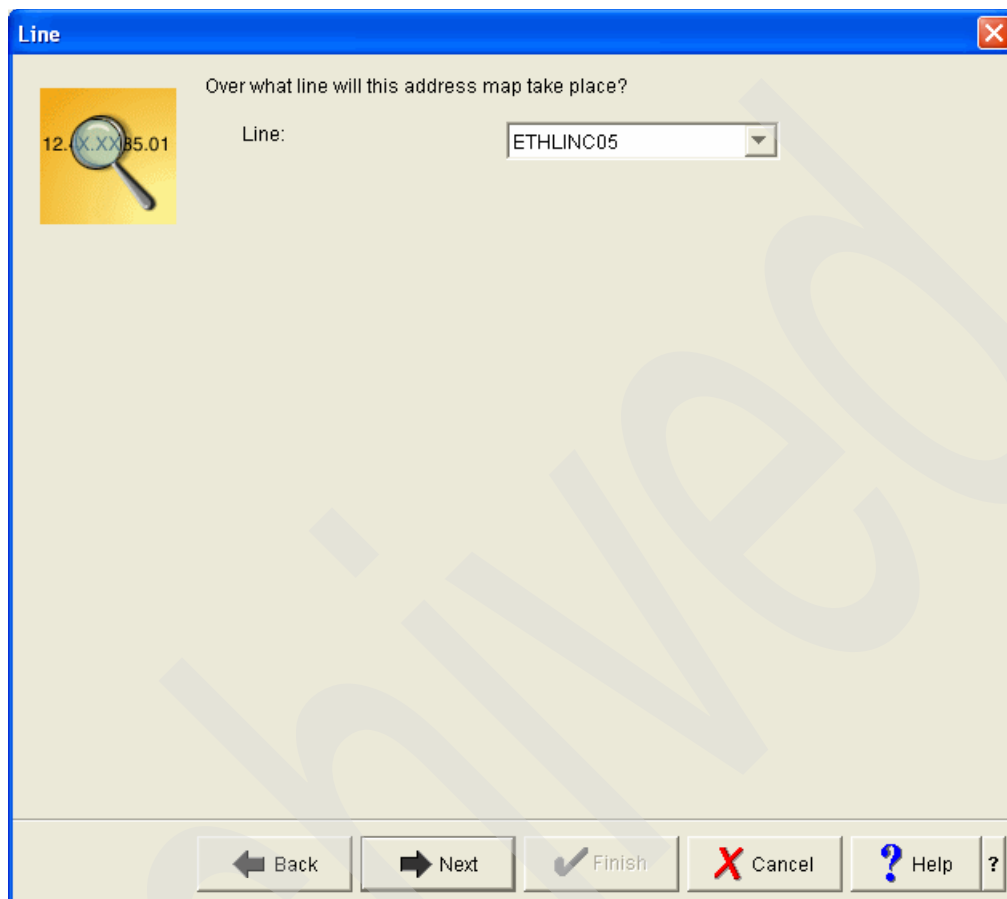


Figure 5-32 Packet Rules Editor - line selection

10. Finally, a summary screen appears as shown in Figure 5-33. Select **Finish** to generate the rules.

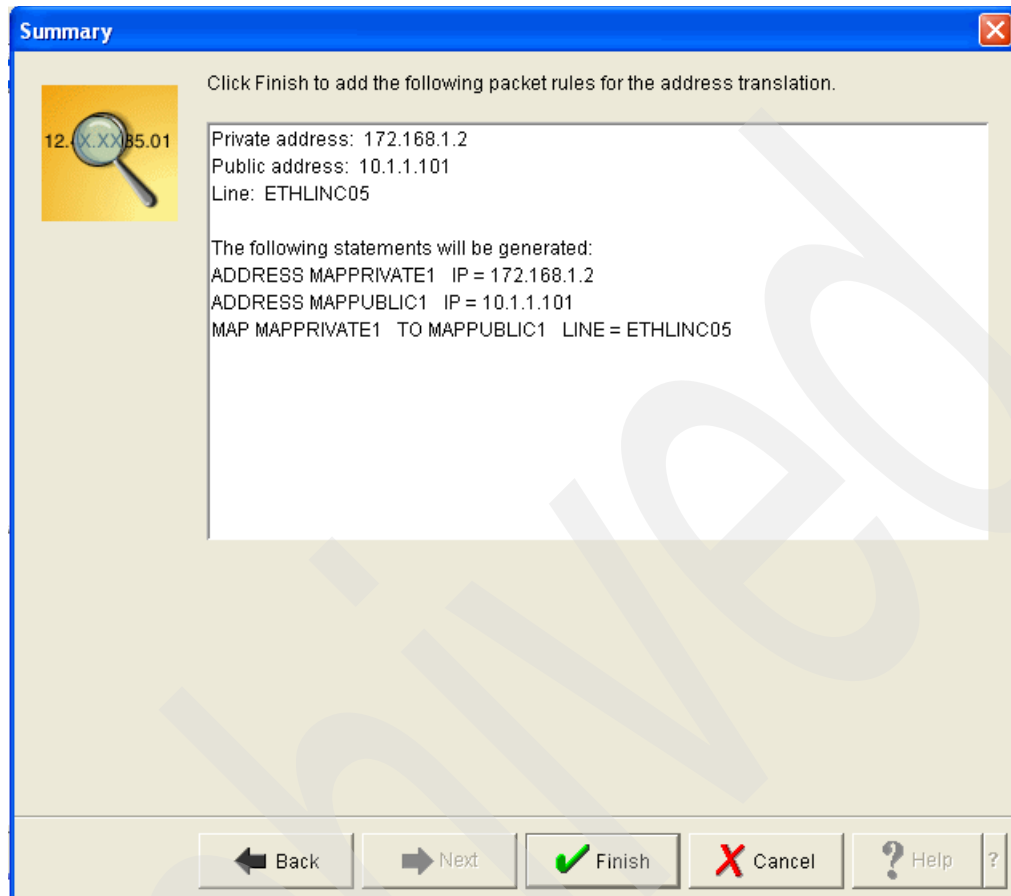


Figure 5-33 Packet Rules Editor - summary

Note: Remember to create an address translation rule for each public address that needs to be translated to a private address.

11. Once the rules have been generated, they will be displayed in the Packet Rules Editor window. Prior to activation and usage of the rules they need to be saved. To save the rules, select **File** from the menu bar and then select **Save As**.

12. In the Save File window shown in Figure 5-34, specify a file name for the packet rules. Ensure that the Type is left as I3P files and click **Save**.

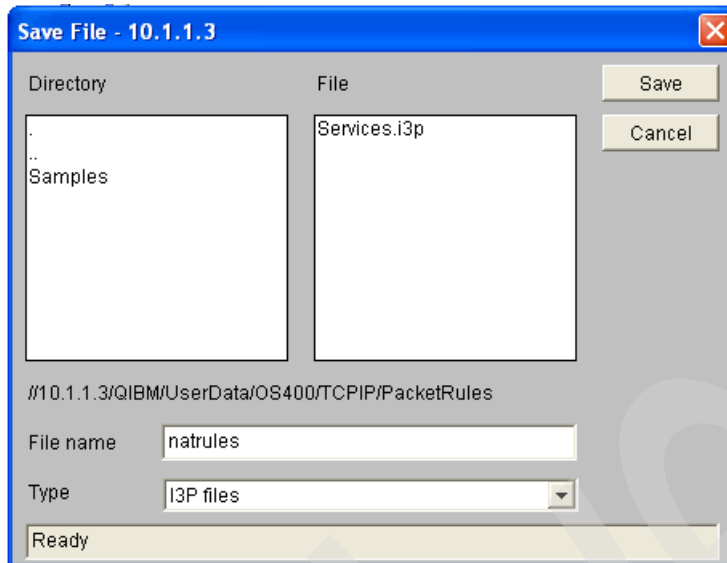


Figure 5-34 Packet Rules Editor - save file

13. To activate the newly-created rules, select **File** from the menu bar and then select **Activate Rules**.

14. In the Activate Packet Rules window shown in Figure 5-35, leave the selections at the defaults and click **OK**. As a point of reference, here are the correct settings:

- **Packet rules files:** activate both VPN-generated rules and selected file.
- **File name:** name of NAT rules just saved.
- **Interface:** activate these rules on all interfaces and all point-to-point identifiers.

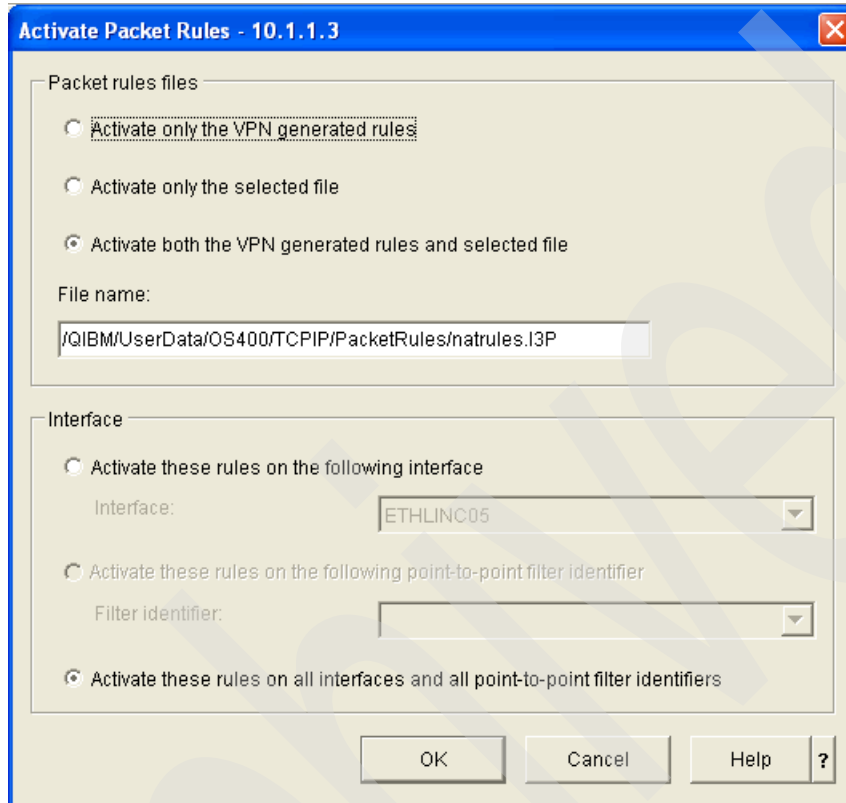


Figure 5-35 Activate packet rules

15. A message indicating that the rules have been activated will be displayed at the bottom of the Packet Rules Editor window.

Linux network settings

All that is left is to configure networking for the Linux partition. The tools used to configure networking in Linux vary based on the tool set available within the Linux distribution, however, the actual settings will be the same. The network parameters within Linux should be set as follows:

- **TCP/IP address:** The network address for the Linux partition on the virtual LAN. This is the address assigned as the “private” address in the NAT rules.
- **Subnet mask:** The subnet mask assigned to the virtual LAN (typically set to 255.255.255.0). This is the same subnet mask assigned to the i5/OS’s interface on the virtual LAN.
- **Router/Gateway:** The address of the i5/OS partition’s virtual LAN connection.

Figure 5-36 represents a completely defined NAT environment.

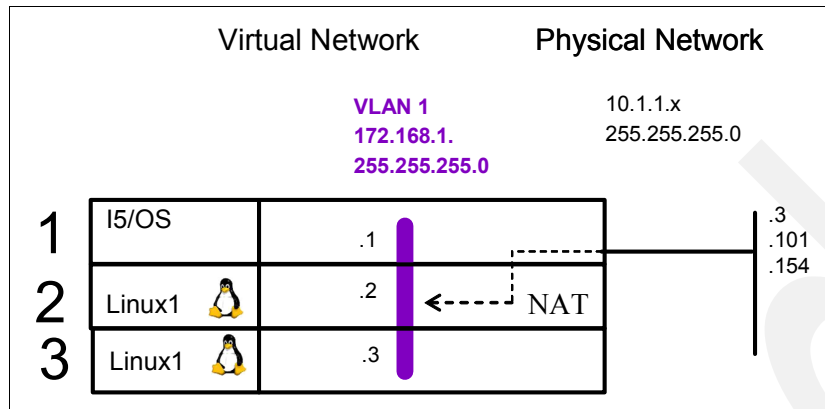


Figure 5-36 NAT - example

In Figure 5-36, a range of addresses (172.168.1.1 to 172.168.1.254) has been assigned to the virtual LAN and public addresses (10.1.1.101 and 10.1.1.154) for the partitions have been assigned to the i5/OS physical interface. Additionally, NAT rules have been established that translate the public (physical network) addresses to the private (virtual network) addresses.

5.2.6 Other approaches

In addition to Network Address Translation and Proxy ARP via i5/OS, there are other alternatives to exporting partitions connected to a virtual LAN to external network traffic. While not covered in detail in this book, these methods include:

- ▶ Proxy ARP via Linux
- ▶ Direct routing

Proxy ARP via Linux involves a Linux partition with both a physical network connection as well as a virtual network connection to other Linux partitions. This method involves the establishment of direct routing statements to the IP addresses of the partitions on the virtual LAN and like Proxy ARP in i5/OS, the Linux partition acting as the proxy responds to the IP addresses for the Linux partitions and makes the resulting traffic available on the virtual LAN.

Direct routing involves making configuration changes to the routers that the System i platform is connected to. In this method, route statements are configured in the physical router that directs traffic for the virtual LAN IP addresses to the i5/OS interface on the physical LAN. Like the Proxy ARP solutions, the DNS entries in the network will point to the addresses on the virtual LAN. When traffic for one of the IP addresses on the virtual LAN is seen by the physical router, it consults its route table for the destination and then forwards the traffic to that destination (in this case the i5/OS physical network adapter). Outbound traffic (that is, traffic going out from the partitions on the virtual LAN) direct their traffic to the i5/OS partition on the virtual LAN (the same way they do for the Proxy ARP and NAT solutions).



Linux installation

This chapter provides general information about the installation of the Novell / SUSE Linux distribution as well as the Red Hat Linux distribution. Information in this chapter includes accessing the Linux console, basic installation steps, as well as post-installation of Linux on POWER utilities.

Note: This chapter is not intended to supersede information in the distributor provided documentation and you are advised to refer to the documentation provided with the distribution for the latest installation information.

6.1 Console access

This section provides the information about console option for Linux partition.

6.1.1 Console access via the Hardware Management Console

The Linux console can be accessed through the Hardware Management Console (HMC). To access the console from the HMC (or WebSM), right-click the Linux partition and select **Open Terminal Window** as shown in Figure 6-1.

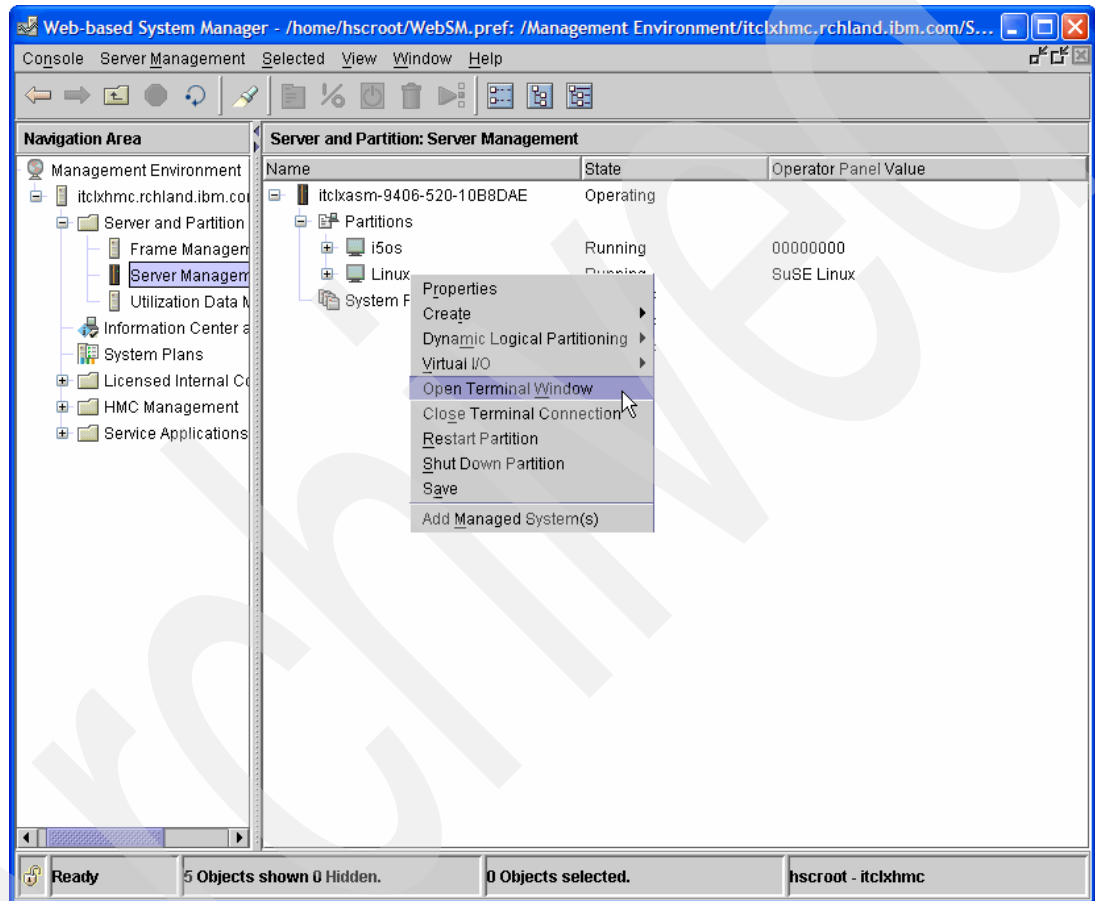


Figure 6-1 HMC - Open terminal window

A VTerm window will be displayed that you can use to interact with the Linux operating system.

6.1.2 Console access via i5/OS

In addition to access the Linux console through the HMC, the console can also be accessed through a TCP/IP enabled application running in i5/OS. Access to the Linux console through this i5/OS application requires the following:

- ▶ A Virtual Serial Client adapter in the i5/OS partition that points to the Virtual Serial Server adapter in slot 0 of the Linux partition
- ▶ A DST user ID with authority to access the console

Creating virtual serial adapter connection

Just like with the virtual SCSI server adapter created earlier in the i5/OS partition, the virtual Serial client adapter has to be created twice, once in the profile of the i5/OS partition and once dynamically in the running state of the partition by performing the following steps:

1. From the HMC, right-click the partition profile for the i5/OS partition and select **Properties**.
2. From the Profile Properties window as shown in Figure 6-2, select the Virtual I/O Adapters tab.




Figure 6-2 Partition profile properties

3. Once the virtual adapters are displayed, click the Serial tab as shown in Figure 6-3.

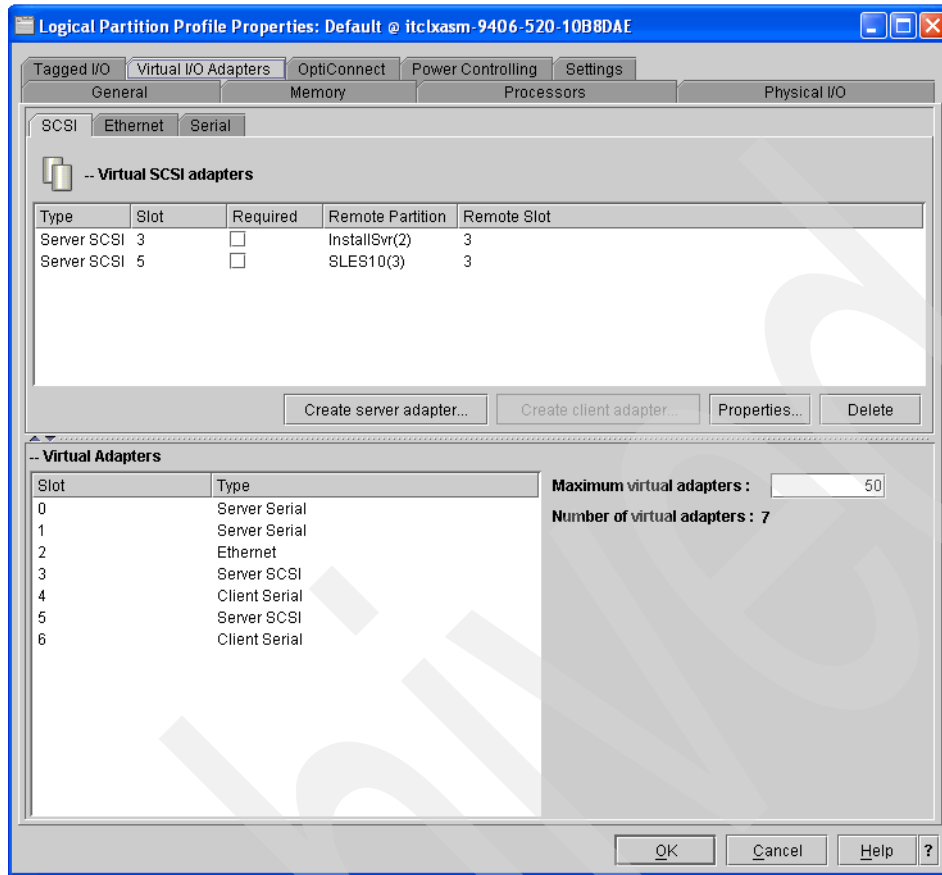


Figure 6-3 Partition Profile Properties - Virtual Adapters

4. Once the Serial tab has been selected, click **Create client adapter**.
5. Complete the fields on the Client Adapter Properties window as shown in Figure 6-4 as follows:
 - a. **Client Slot**: Assigned by the system as the next available (unused) slot number. Typically, this value will be left as provided.
 - b. **Server partition**: Select the name of the Linux partition from the drop-down list.
 - c. **Server partition slot**: Leave as the default value of 0 which is the first virtual Serial Server adapter in the Linux partition.

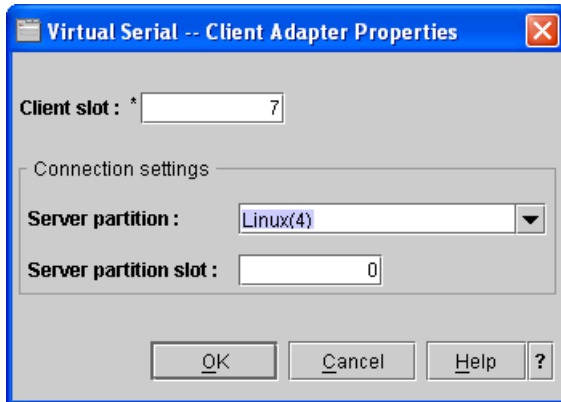


Figure 6-4 Virtual Serial - Client Adapter Properties

6. Click **OK** on the Adapter Properties window and then click **OK** on the Partition Profile Properties window.
7. To add the adapter to the running state of the partition, in the HMC, right-click the i5/OS partition (not the partition profile) and select **Dynamic Logical Partitioning** → **Virtual Adapter Resources** → **Add / Remove**.
8. On the Dynamic Logical Partitioning - Virtual Adapters window (Figure 6-5), select the Serial tab, and then click **Create client adapter**.

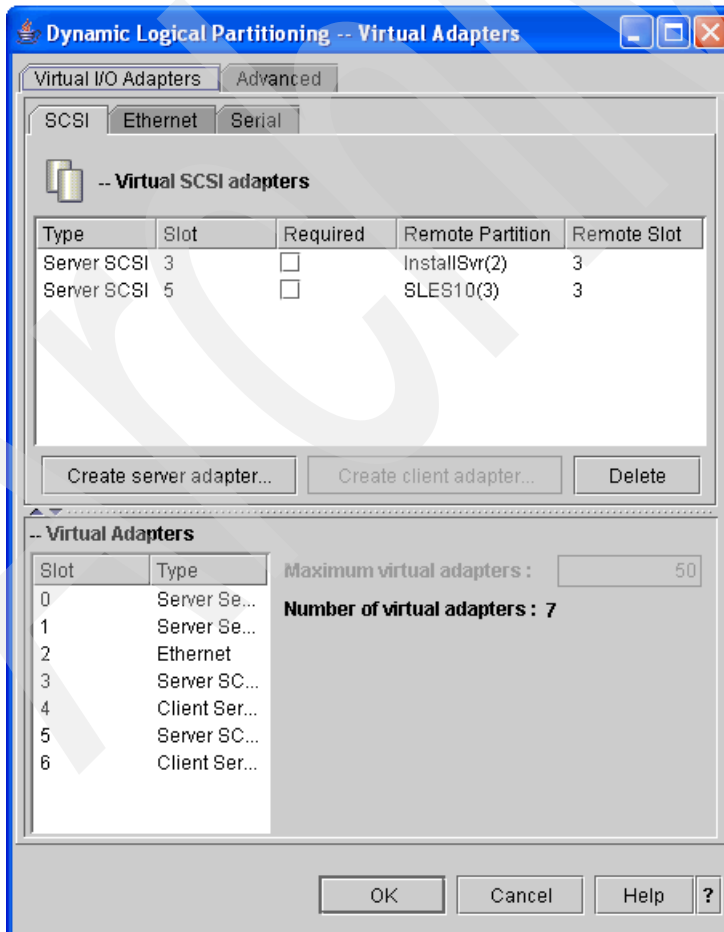


Figure 6-5 Dynamic Logical Partitioning - Virtual Adapters

9. Complete the fields of the Client Adapter Properties window to match those created above in the partition profile.
10. Once the adapter probates have been defined, click **OK** on the adapter properties window and then click **OK** on the Virtual Adapters window.

The documenting of the logical partition definition should be updated to include the addition of the Serial Client adapter in the i5/OS partition and the connection of that adapter to the Serial Server adapter in the Linux partition (Figure 6-6).

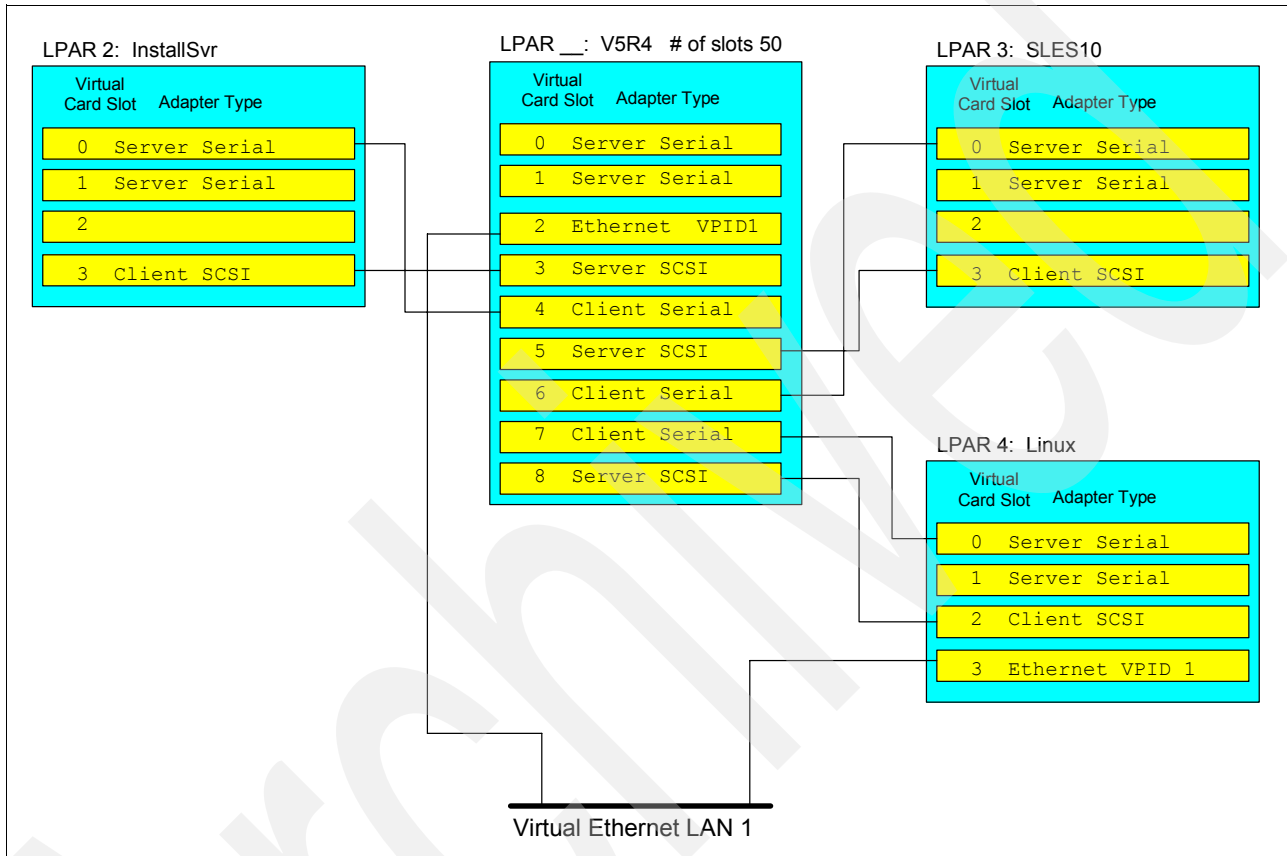


Figure 6-6 Logical Partition - Virtual Adapter Definitions

Note the addition of the Client Serial adapter to the i5/OS partition in Slot 7 and the mapping of that adapter to the Server Serial adapter in slot 0 of the Linux partition.

Setting DST user ID permissions

Once the i5/OS partition has been updated to have a serial connection to the Linux partition, you have to create a DST user ID with the requisition privileges to be able to access the virtual console.

Use the following steps to create the Dedicated System Service Tools user with the correct authorities:

1. DST users are created through System Service Tools. Type the STRSST command to start System Service Tools. When the System Service Tools menu is displayed, select option **8** 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).

- When the Work with Service Tools User IDs window is displayed as shown in Figure 6-7, enter option 1 to create a new user and enter the user name in the User ID field.

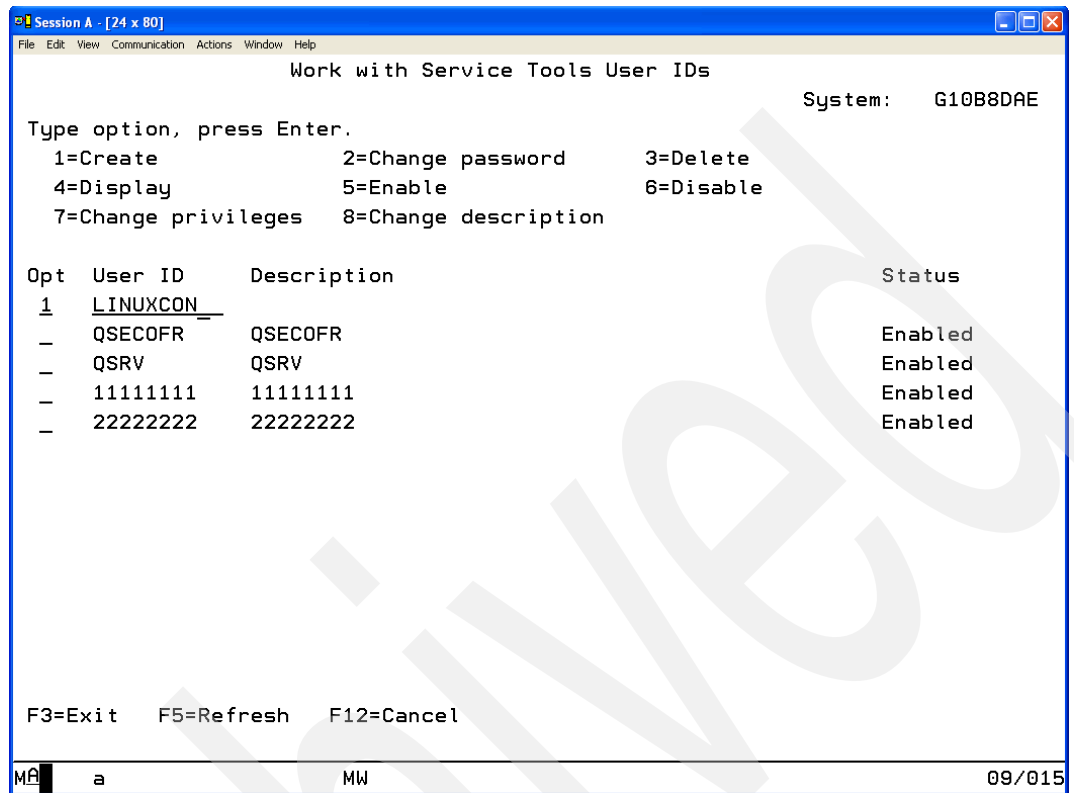


Figure 6-7 Work with Service Tools IDs

4. Press **Enter**. The Create Service Tools User ID window is displayed as shown in Figure 6-8.

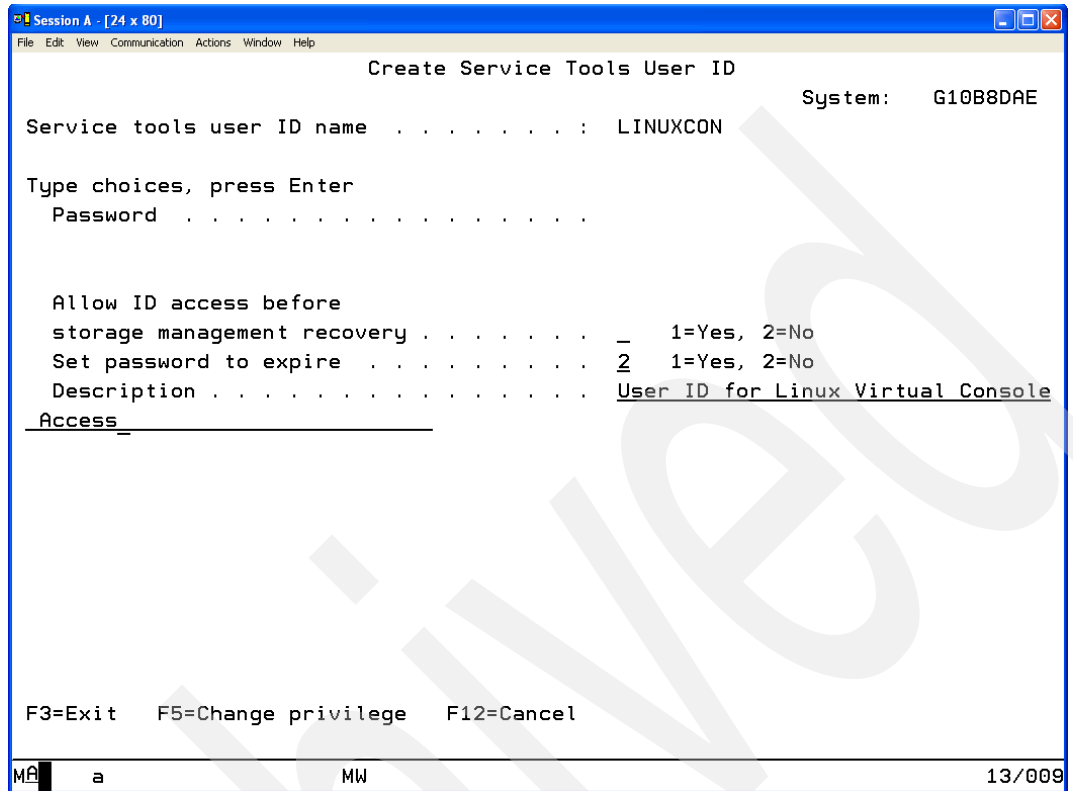


Figure 6-8 Create Service Tools User ID

In Figure 6-8:

- Password: This is the password that you will need to enter to use this Service Tools User ID.
 - Allow ID access before storage management recovery: Indicates if the user ID should be available prior to completion of the storage management recovery on an IPL of the system. Since this User ID will be used to access the Linux console it can either be set to 2 (No) or left blank.
 - Set password to expire: Typically left at the default value of 2 (No).
 - Description: A description of the usage of the User ID.
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.
 7. On the Work with Service Tools User IDs, select option 7 (Change Privileges) for the user just created.
 8. On the change Service Tools User Privileges shown in Figure 6-9, typically for a User ID that is just going to be used to access the console, the following three privileges are granted:
 - System partitions - operations
 - System partitions - administration
 - Partition remote panel key

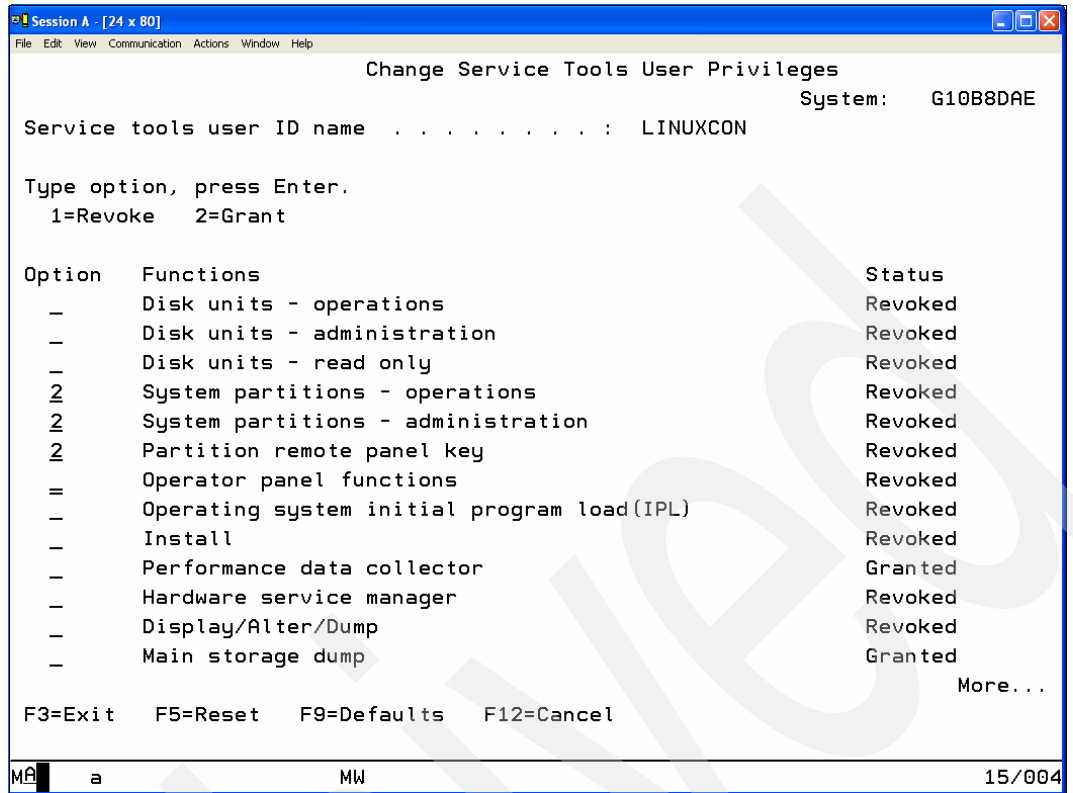


Figure 6-9 Change Service Tools User Privileges

9. After the required authority has been granted, press **Enter**.
10. Press the **F3** key to exit the Change Service Tools User Privileges menu.
11. Press the **F3** key to exit the Work with Service Tools User IDs menu.
12. Press the **F3** key to exit the Work with Service Tools User IDs and Devices menu.
13. Press the **F3** key to exit the System Service Tools menu.
14. When prompted to confirm the exit of System Service Tools, press **Enter**.

Accessing the console

Prior to stating 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:

1. Use a telnet client, such as the Windows telnet client or the PuTTY 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 **telnet 192.168.10.10 2301** command.

2. A list of Linux partitions is provided (Figure 6-10), enter the number of the partition to connect to and press **Enter**.

```
OS/400 Guest Partition Consoles
2: InstallSvr(V7-C4/C2-C0)
3: SLES10 (V7-C6/V3-C0)
4: Linux(V7-C7/V4-C0)

Enter the console partition number:
```

Figure 6-10 Virtual Console - Partition Selection

Note: If the System i platform has Windows Integration (IXS or IXA, or both) devices, then a menu will be displayed prompting you to select whether to attach to a Windows Integration devices or a guest partition. In that case, specify Guest Partition at which point the above menu will be displayed.

The information in the parenthesis () following the partition name indicates the Serial adapter pairing. The first item indicates the partition number (V#) and the slot number (C#) of the Server Serial adapter (the adapter in the i5/OS partition), while the second item indicates the partition number and slot number of the associated Client Serial adapter in the Linux partition.

3. When prompted for the i5/OS service tools user ID (as shown in Figure 6-11), type the DST user that was created for virtual console access.

```
Linux: Enter OS/400 service tools userid:
```

Figure 6-11 Virtual Console - User ID Prompt

4. When prompted for the i5/OS service tools password (as shown in Figure 6-12), type the password defined for the DST user.

```
Linux: Enter OS/400 service tools password
```

Figure 6-12 Virtual Console - Password Prompt

5. If the User ID and password were entered correctly, a connection will be established (as shown in Figure 6-13).

```
Linux: Console connecting...
Linux: Console connected.
```

Figure 6-13 Virtual Console - Console Connection Messages

6.2 Basic Linux installation steps

This section describes a basic installation scenario for the Novell / SUSE SLES 10 and Red Hat RHEL 4 Linux distributions. Refer to the distributor provided documentation for more detailed installation information.

6.2.1 Novell / SUSE SLES 10

At the time of writing of this book, SLES 10 was the latest Linux distribution available from Novell / SUSE.

Network server setting of interest

The IPL Parms field of the Network Server provides the capability to define parameters that will be passed directly to the Linux kernel. The distributor provided documentation will provide a detailed list of possible parameters. One parameter of interest is VNC=1. This parameter instructs the installation program to perform the installation in a graphical mode making use of the Virtual Network Computing server/client technology.

Installation steps

Note: The following example assumes that the Network Server has been configured to boot off of the installation kernel (typically /qopt/SLES100.001/suseboot/inst64), the IPL parameters line includes 'VNC=1', the virtual console has been accessed, and you have access to the VNC client which you downloaded from the Web at:

<http://www.realvnc.com>

A text-based installation can also be performed by omitting the VNC=1 value from the IPL parameters field of the Network Server.

1. From the hosting i5/OS partition, vary on the Network Server. After a few seconds, the portion of the installer will be executed which will collect network configuration information and start the VNC server. A prompt will be displayed instructing you to access the next portion of the installer via the VNC viewer.

The first screen displayed is the language selection screen (as shown in Figure 6-14). The language selected is the language that will be used for the remainder of the installation screens. Select the desired language, then click **Next**.

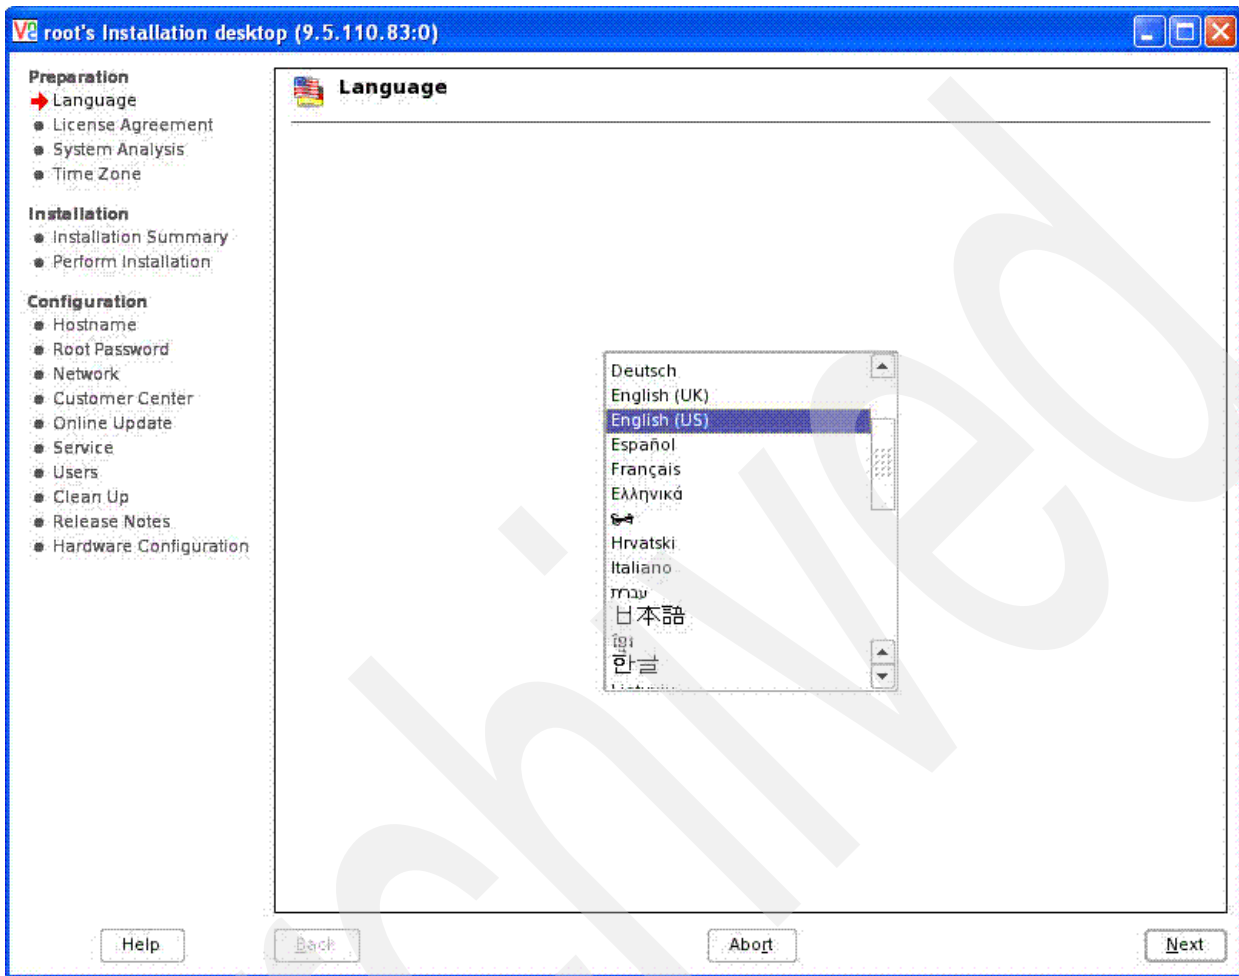


Figure 6-14 SLES 10 Language Selection Screen

- The next screen is the License Agreement screen (). Select **Yes, I Agree to the License Agreement** and then click **Next**.

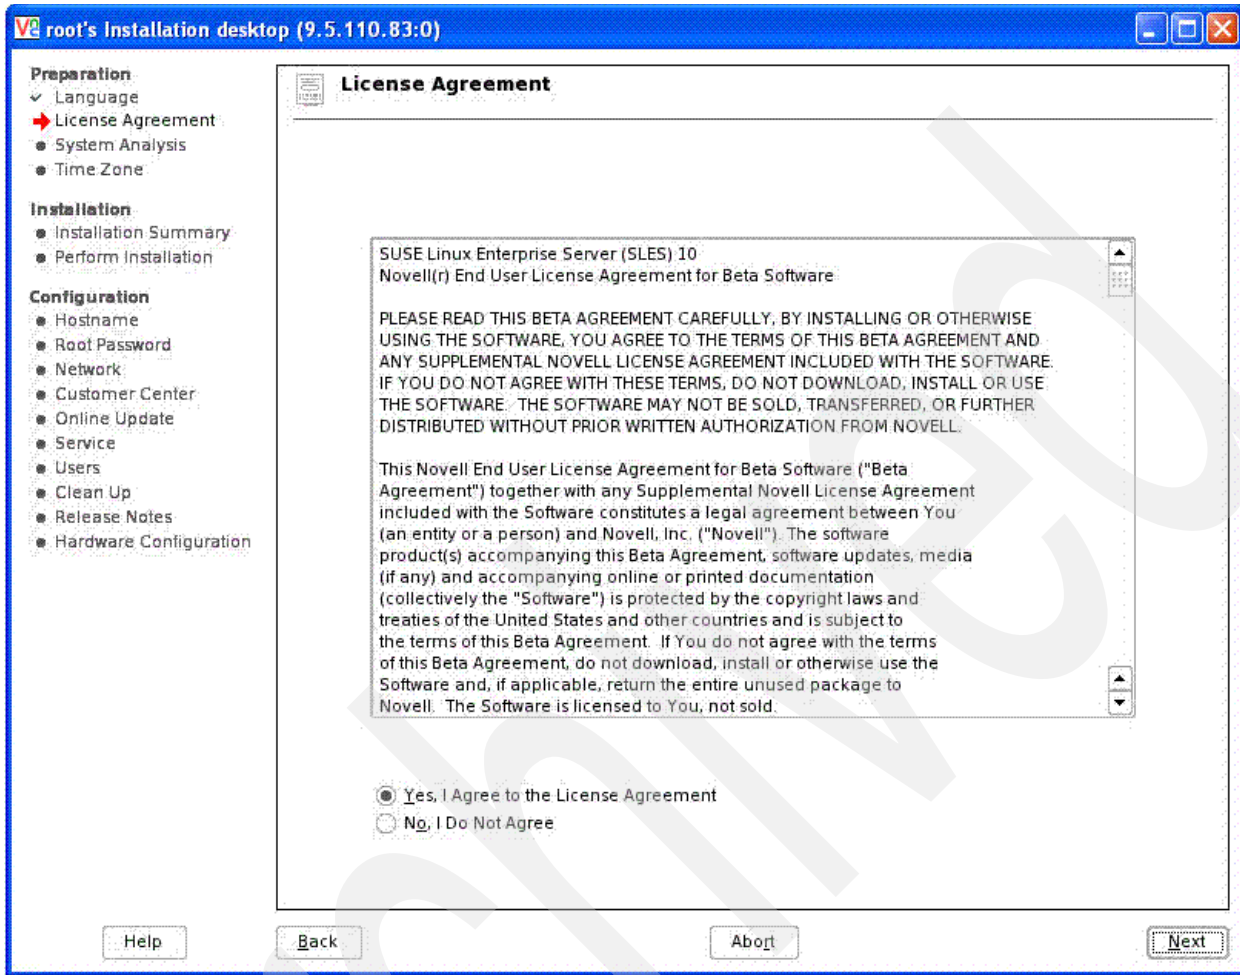


Figure 6-15 SLES 10 - License Agreement

3. The next screen is the “Installation Mode” screen (as shown in Figure 6-16). Since the installation is being performed on a new disk, the only choice available is “New Installation”. Click **Next**.

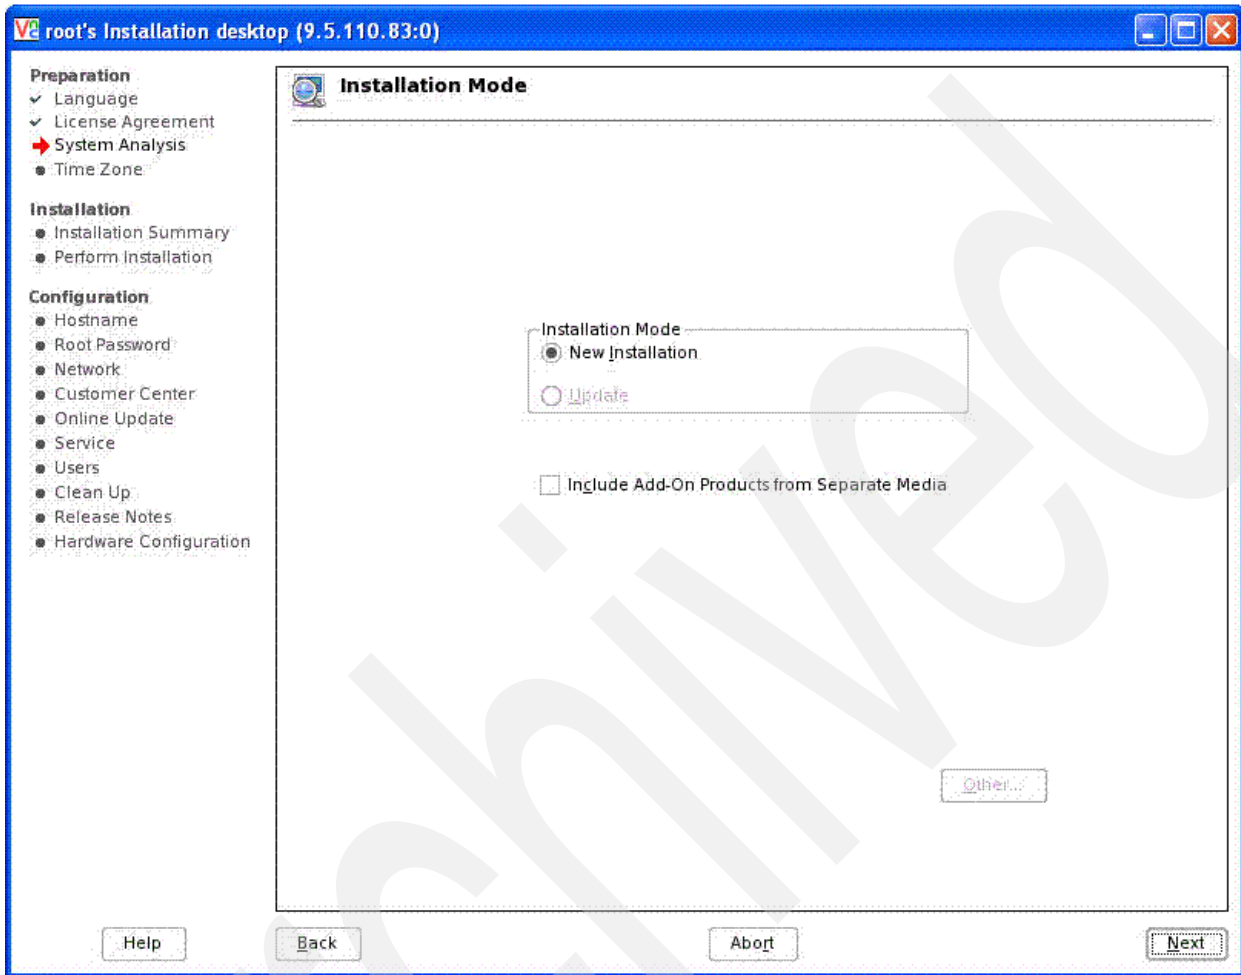


Figure 6-16 SLES 10 - Installation Mode

Note: The update capability of SLES 10 (that is, the ability to update from a SLES 9 installation to SLES 10) is described in 6.4, “Updating SLES 9 to SLES 10” on page 239.

- The “Clock and Time Zone” window (as shown in Figure 6-17) is used to select the time zone and validate/change the time. Make the appropriate selections and then click **Next**.

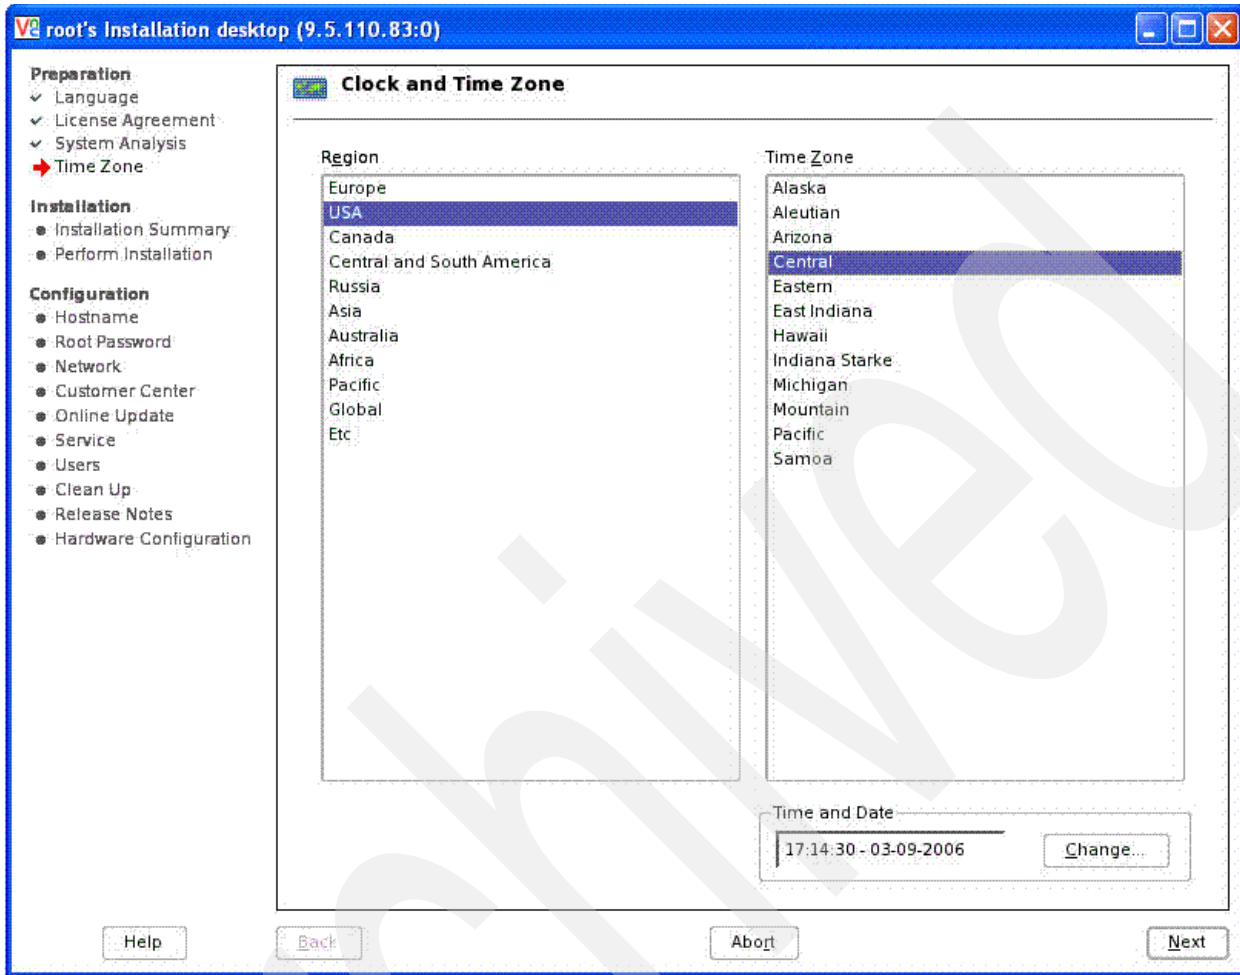


Figure 6-17 SLES 10 - Clock and Time Zone

- The next window is the “Installation Settings” window (as shown in Figure 6-18). This window essentially summarizes the disk partitioning and software selections that will be

used by the installer. You can either accept these values or change them (via the Change button). Once the settings are complete, click **Accept**.

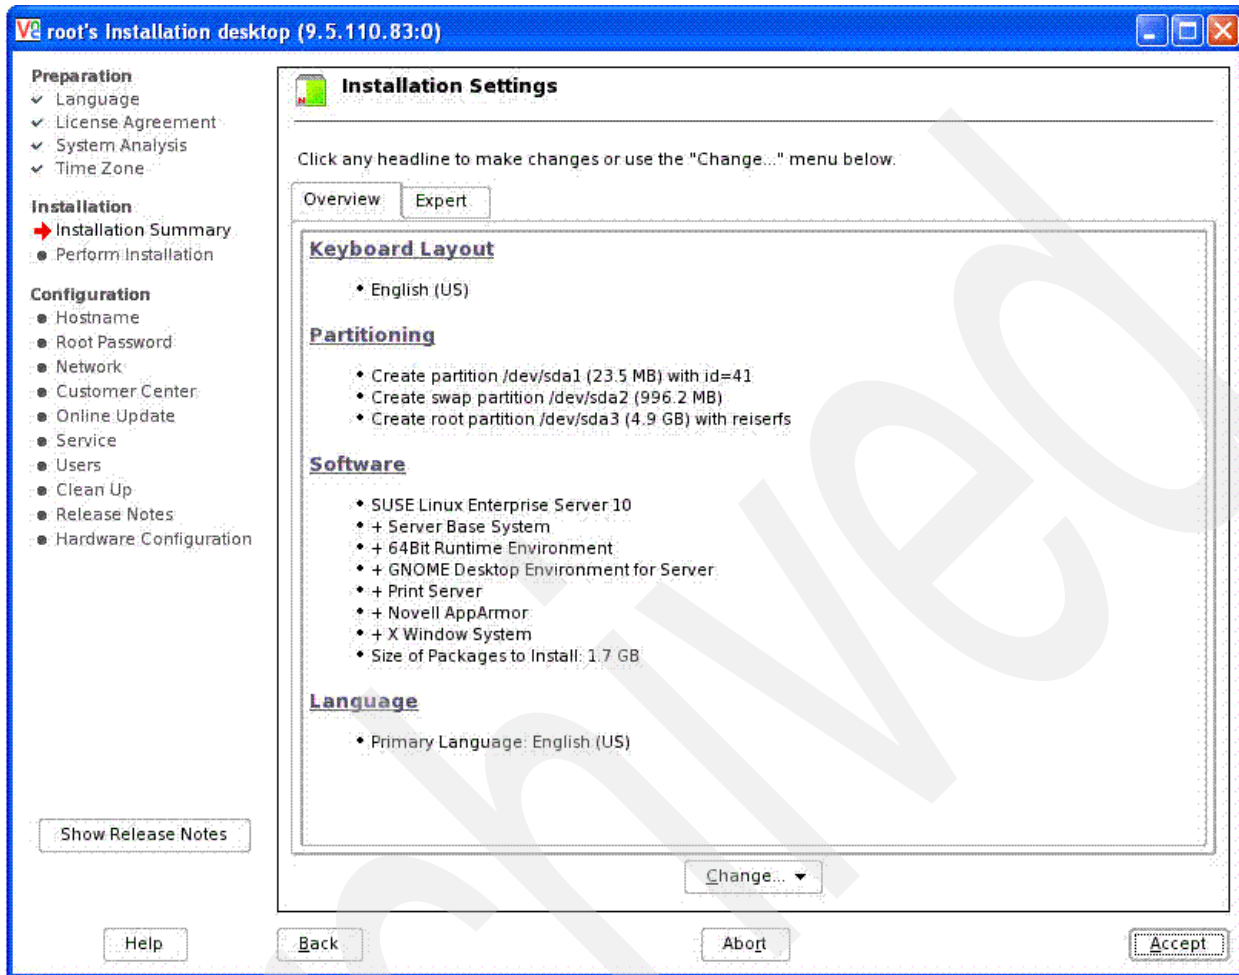


Figure 6-18 SLES 10 - Installation Settings

- At this point the installer has all of the values required to perform the installation. A confirmation window will be displayed (as shown in Figure 6-19). Click **Install** to continue the installation.

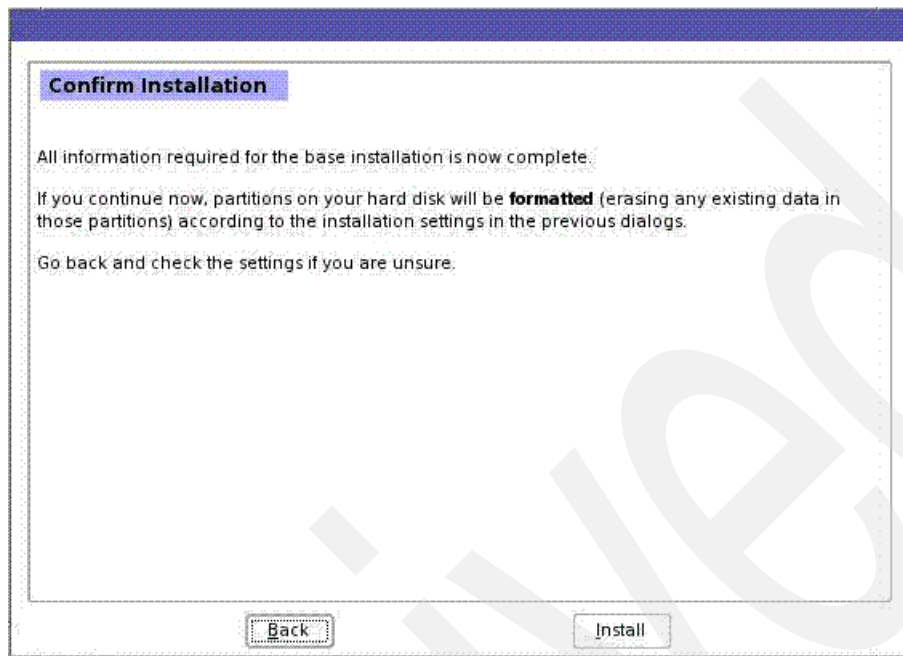


Figure 6-19 SLES 10 - Confirm Installation

- a. The first step that the installer will perform is the partitioning and formatting of the disk (as shown in Figure 6-20).

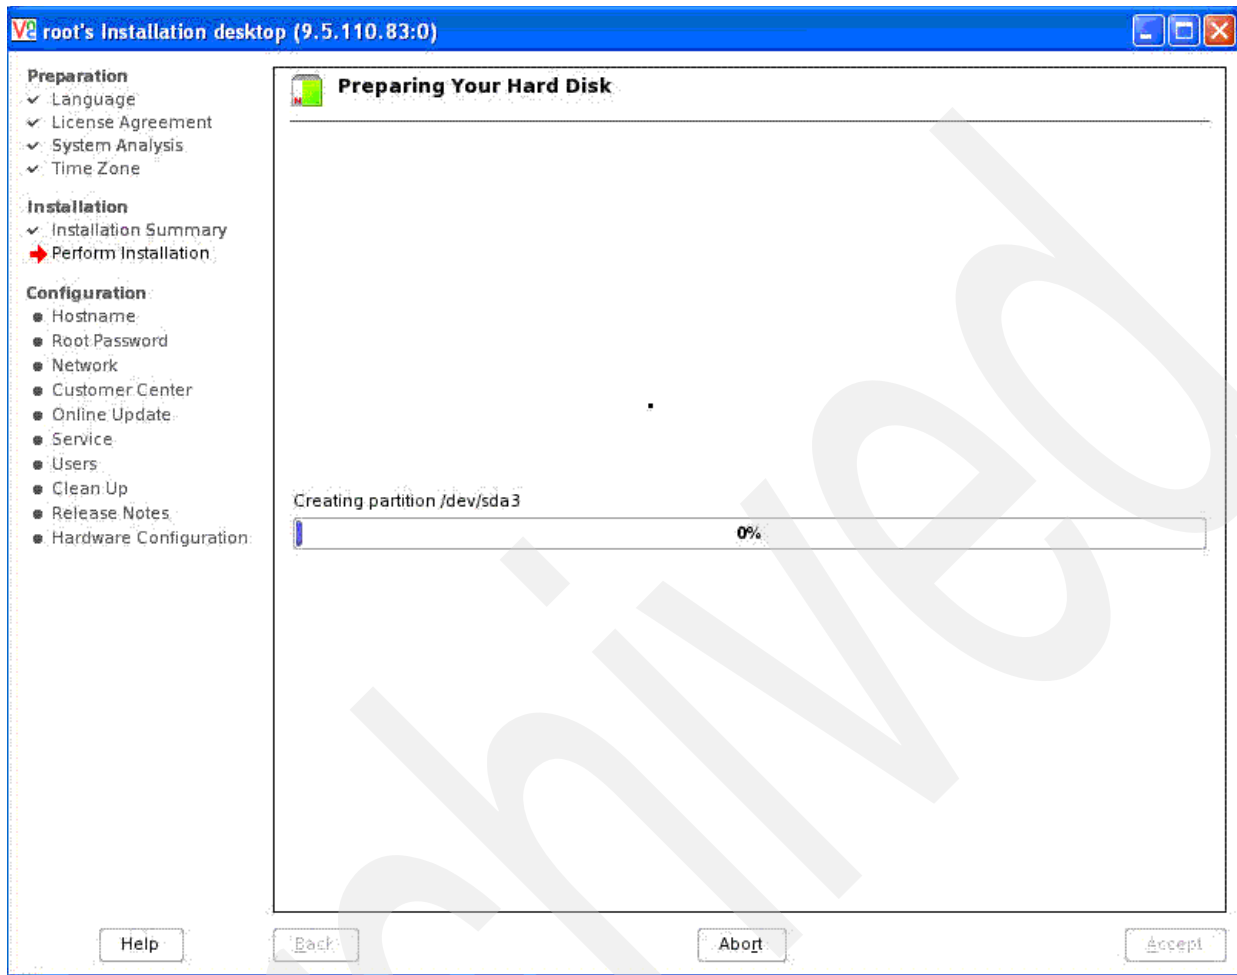


Figure 6-20 SLES 10 - Preparing Your Hard Disk

- b. Once the disk has been partitioned, the packages will be installed (as shown in Figure 6-21).

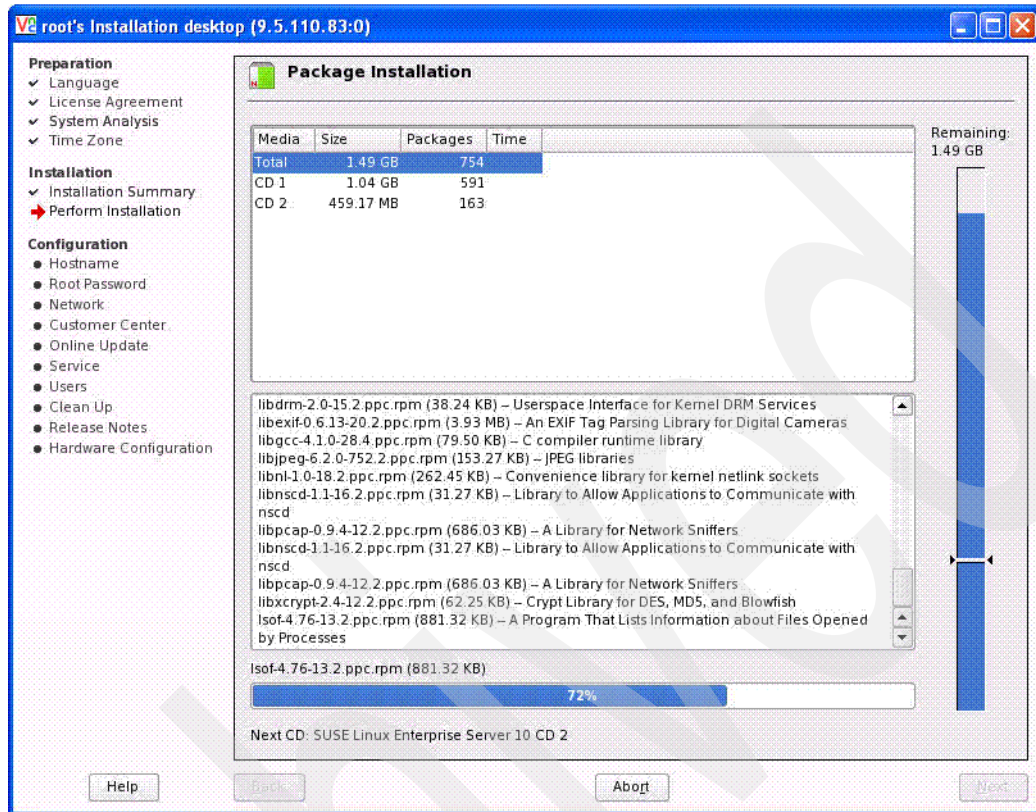


Figure 6-21 SLES 10 - Package Installation

Prompts will be displayed for switching of the installation CDs.

- c. Once the installation of packages has been completed, system settings will be updated (as shown in Figure 6-22).

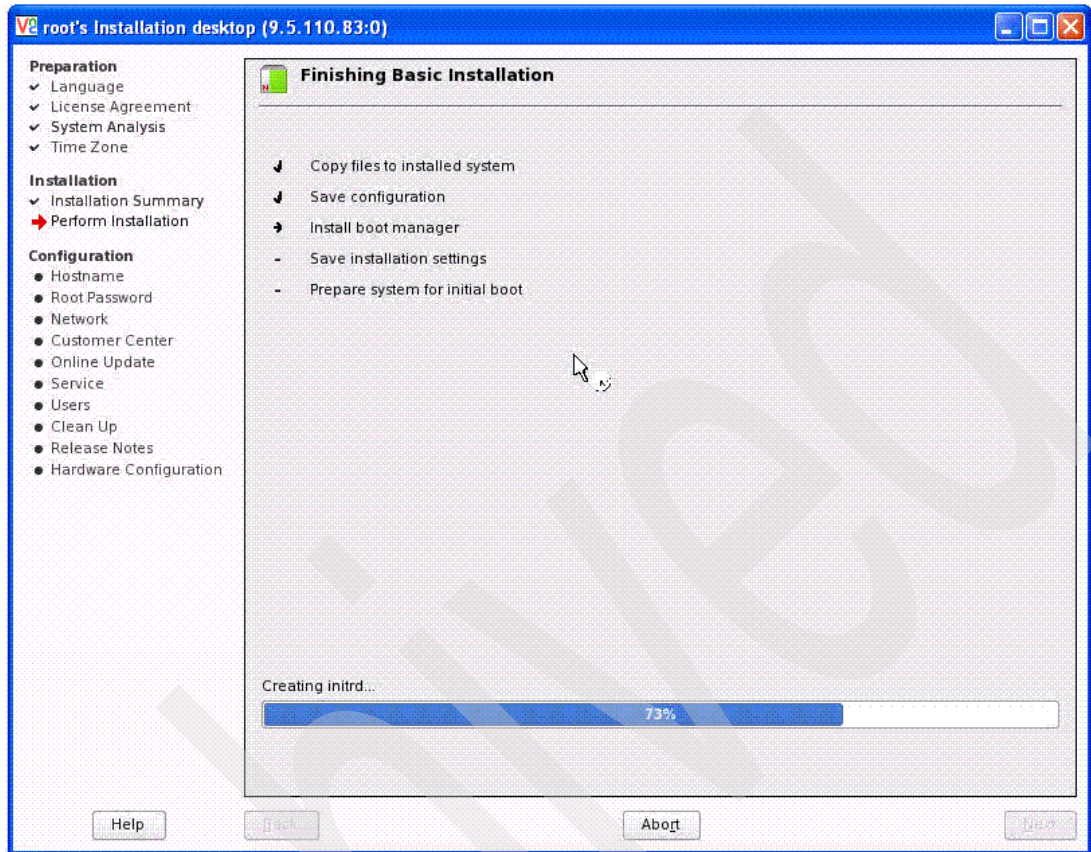


Figure 6-22 SLES 10 - Finishing Basic Installation

Once the basic installation has been completed, the Linux partition will be restarted.

When the partition is restarted it is still configured (through the Network Server Description) to boot the installation program. The partition needs to be completely shutdown so that the Network Server can be updated to boot the installed kernel (that is, the kernel installed to the Network Server Storage Space.)

7. When the Main Menu (Figure 6-23) is displayed, type 8 to select PowerOff and press **Enter**.

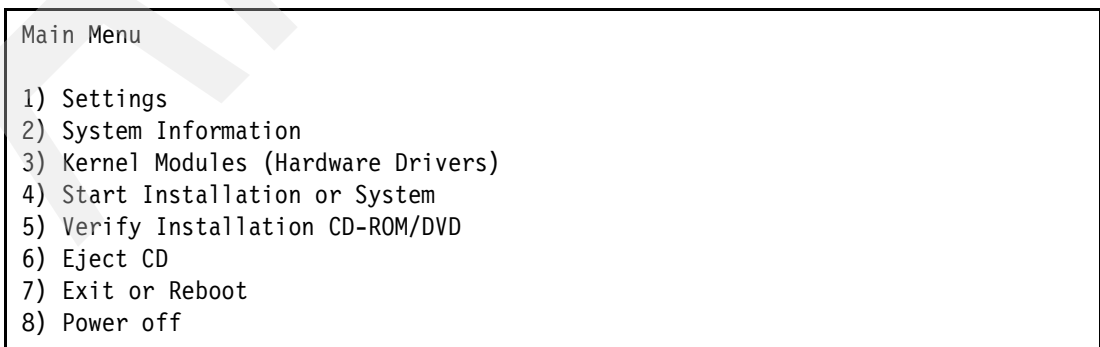


Figure 6-23 SLES 10 - Main Menu

8. When prompted for confirmation of the shutdown request (Figure 6-24), type 1 (Yes) and press **Enter**.

```
Do you want to halt the system now?
```

- ```
1) Yes
2) No
```

Figure 6-24 SLES 10 - Halt System confirmation

At this point the operating system is shutdown which also causes the logical partition to be shutdown.

Prior to update of the Network Server Description, the Network Server needs to be varied off.

9. In the i5/OS hosting partition, enter the command shown in Example 6-1.

Example 6-1 Varying off network server

```
VRYCFG CFGOBJ(SLES10) CFGTYPE(*NWS) STATUS(*OFF)
```

Here, the CFGOBJ value is the name of your network server.

There are three values that need to be modified in the Network Server Description:

- **IPLSRC**: Needs to be changed to \*NWSSTG to boot off of the kernel that was written to the virtual disk during the first phase of the installation.
- **IPLSTMF**: Needs to be set to \*NONE.
- **IPLPARM**: Needs to be set to \*NONE.

10. In the i5/OS hosting partition, enter the following command to make the necessary changes to the Network Server Description (Example 6-2).

Example 6-2 Change Network Server Description

```
CHGNWSD NWSD(SLES10) IPLSRC(*NWSSTG) IPLSTMF(*NONE) IPLPARM(*NONE)
```

The NWSD value is the name of your Network Server Description.

Once the Network Server Description has been updated, the Network Server can be varied back on to complete the installation.

11. In the i5/OS hosting partition, enter the following command to vary on the Network Server (Example 6-3).

Example 6-3 Varying off network server again

```
VRYCFG CFGOBJ(SLES10) CFGTYPE(*NWS) STATUS(*ON)
```

The CFGOBJ value is the name of your network server.

**Note:** The virtual console screen should still be available from the first stage of the installer. If you closed the console, re-activate it either through the HMC or the virtual console support via i5/OS.

Once Linux has been re-started, the next stage of the installer takes over. The first portion of the post installation steps is the detection and configuration of the network adapter.

12. When the Hardware Detection prompt (Figure 6-25) is displayed for the network adapters, click **Continue**.



Figure 6-25 SLES 10 - Network card detection

The next part of the installation installs additional packages from the installation CDs as shown in Figure 6-26.

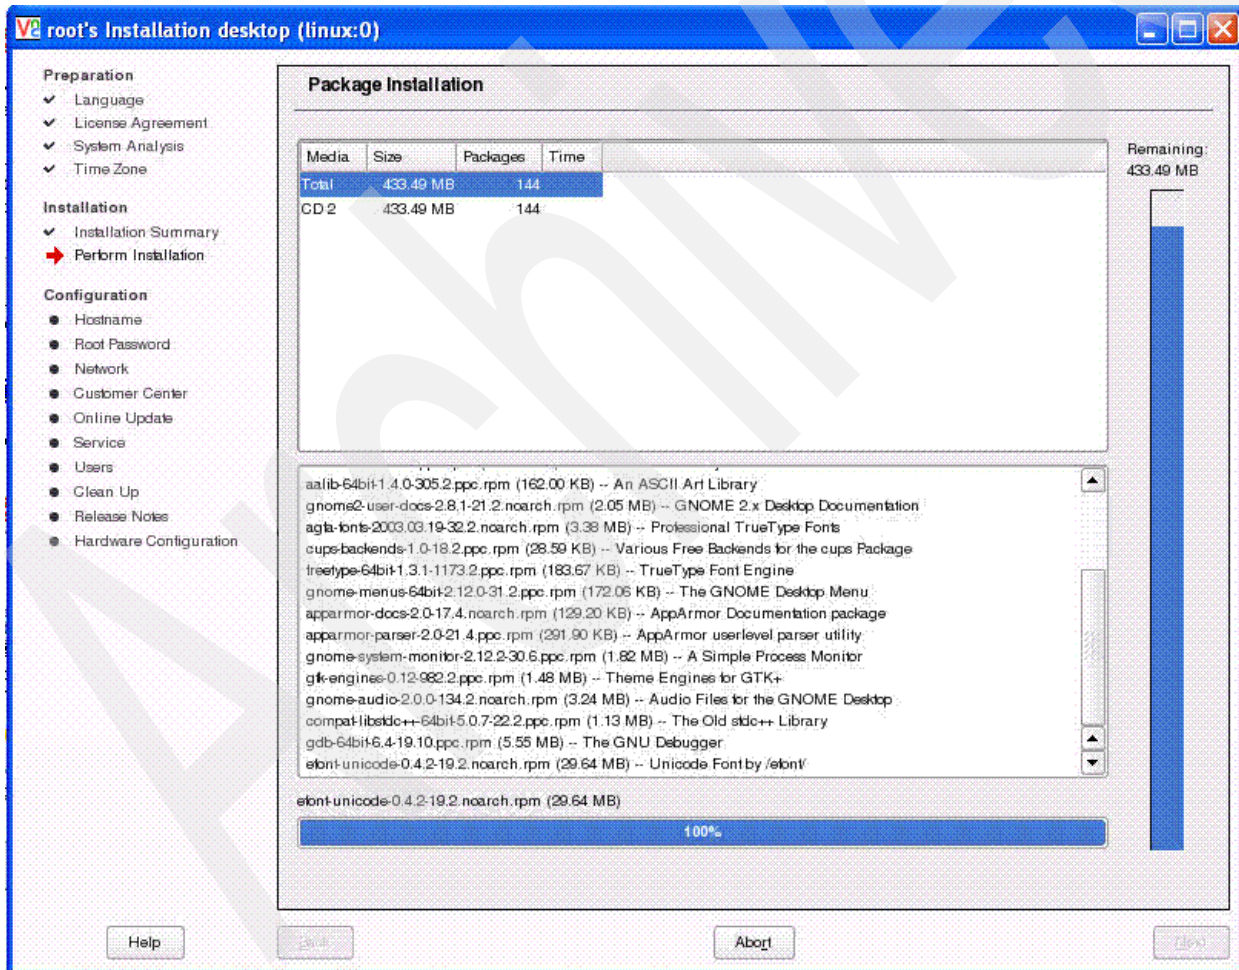


Figure 6-26 SLES 10 - Additional package installation

13. Once the additional packages have been installed, you are prompted for the host name and domain name for the system (shown in Figure 6-27). Provide the appropriate values and then click **Next**.

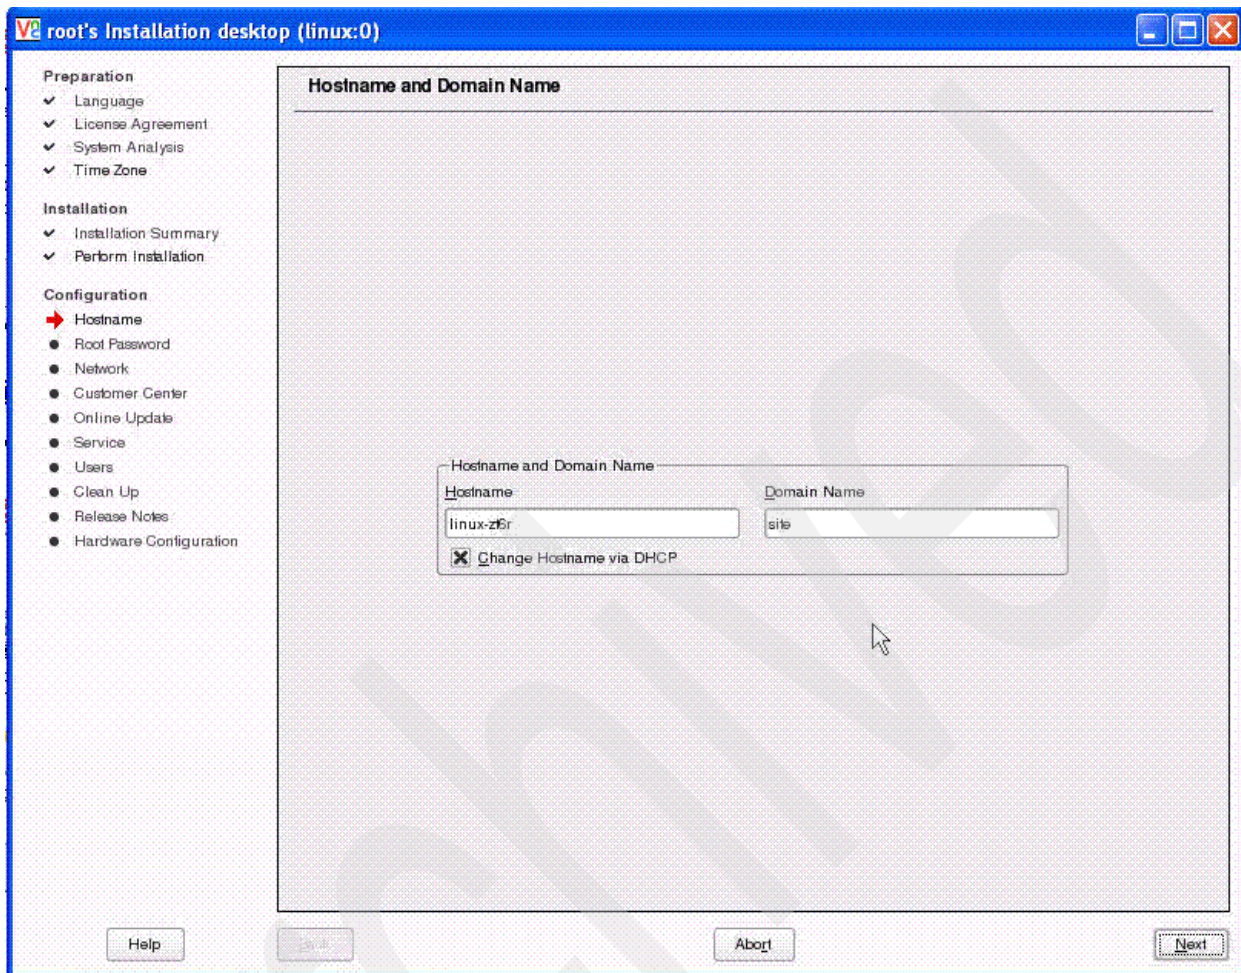


Figure 6-27 SLES 10 - Hostname and Domain Name

14. The next prompt is for the root password (as shown in Figure 6-28). The root user is the user ID that is used to perform system administration functions. Provide and confirm the password and then click **Next**.

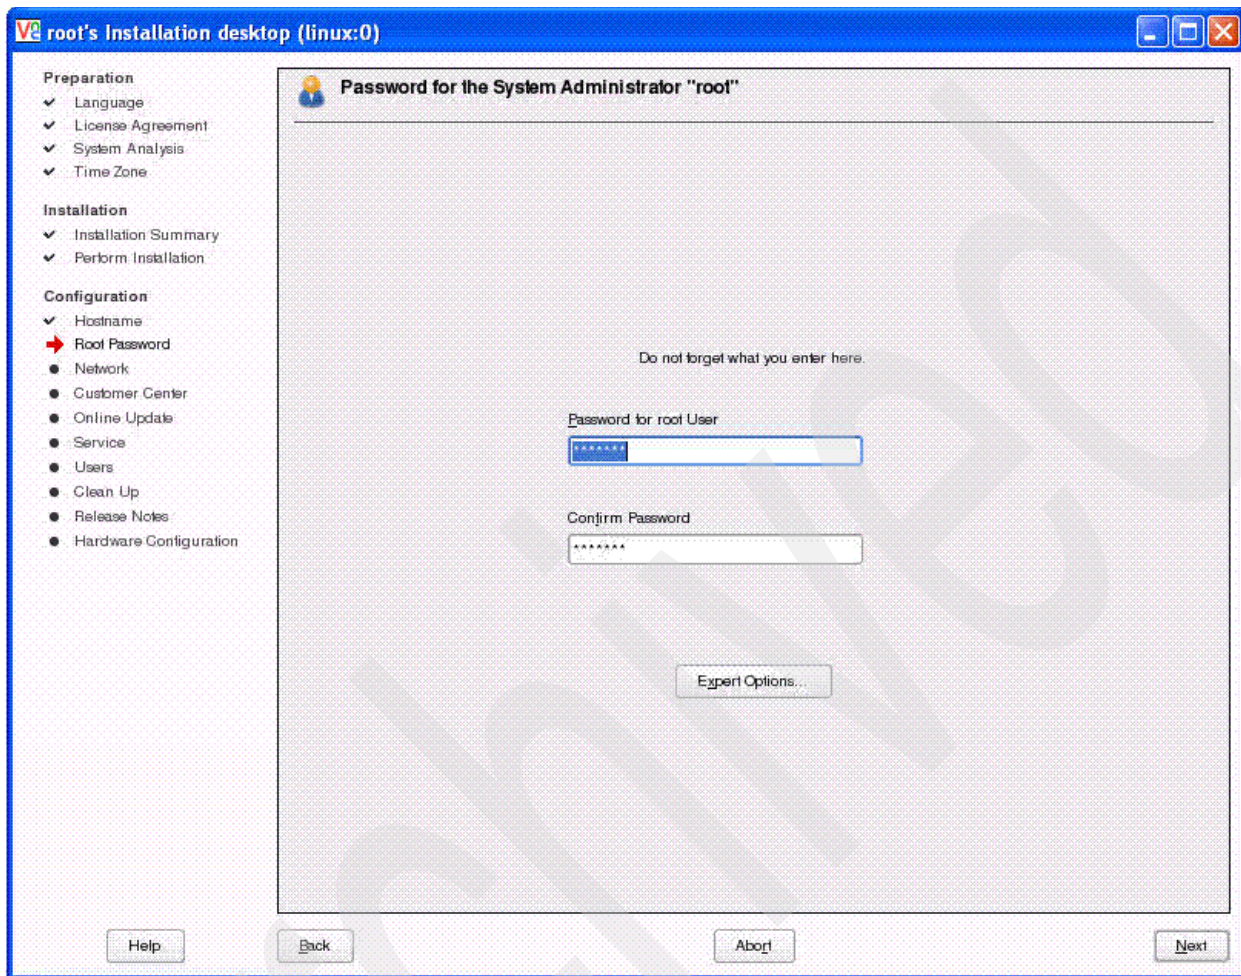


Figure 6-28 SLES 10 - Password for the System Administrator

At this point the Network Configuration screen is displayed as shown in Figure 6-29.

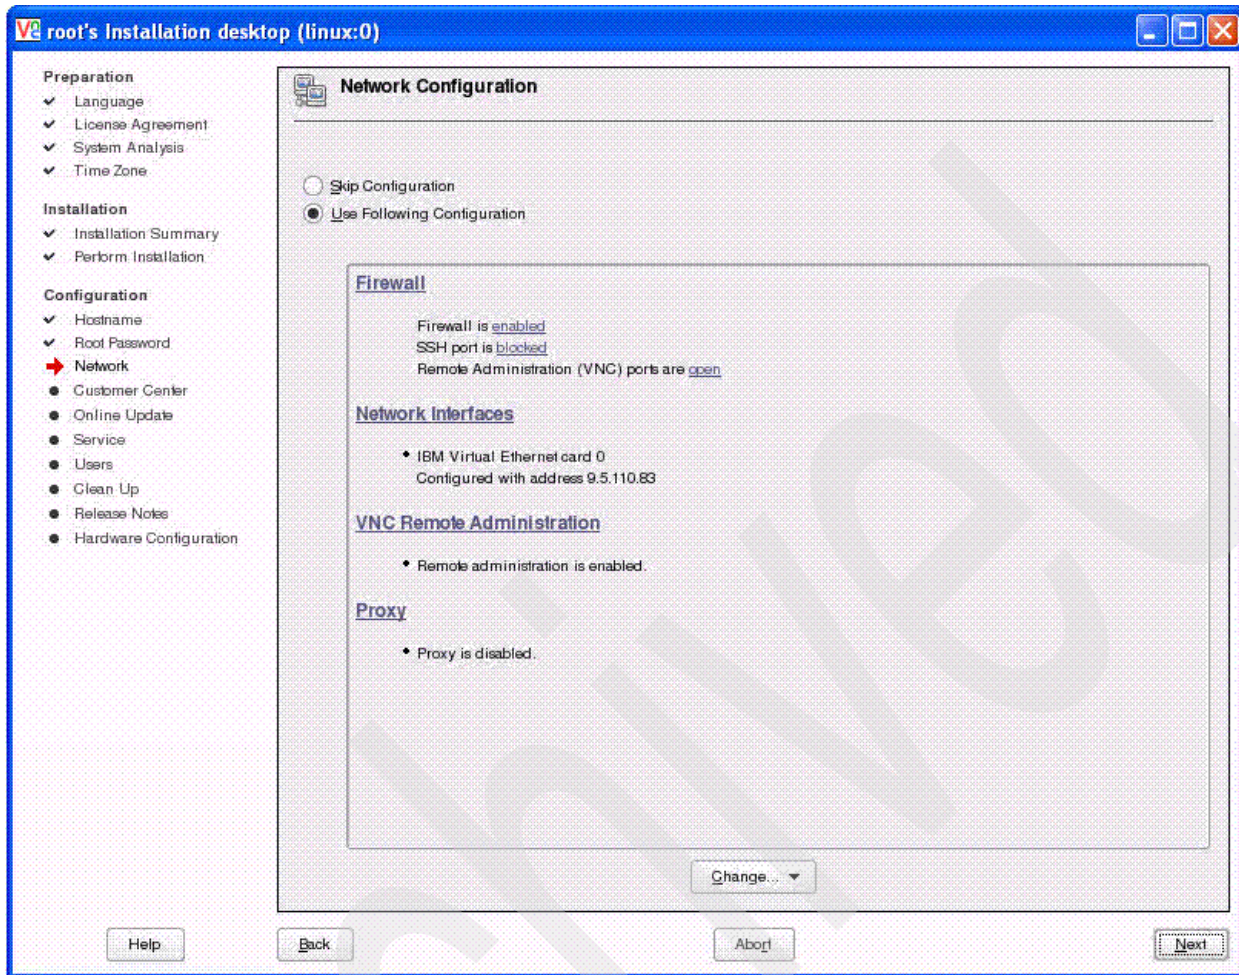


Figure 6-29 SLES 10 - Network Configuration

If you configured the network adapter as part of the first stage of the installer (that is, either for a network based install or a VNC supported installation), then those installation settings are pulled forward and displayed. This portion of the installation also allows you to configure firewall support, web proxy, as well as remote administration support via VNC.

15. Once you have selected the appropriate network configuration settings appropriate to your configuration, click **Next**.

At this point, the installer provides an opportunity to test the network settings by downloading the latest release notes and updates for the distribution, as shown in Figure 6-30.

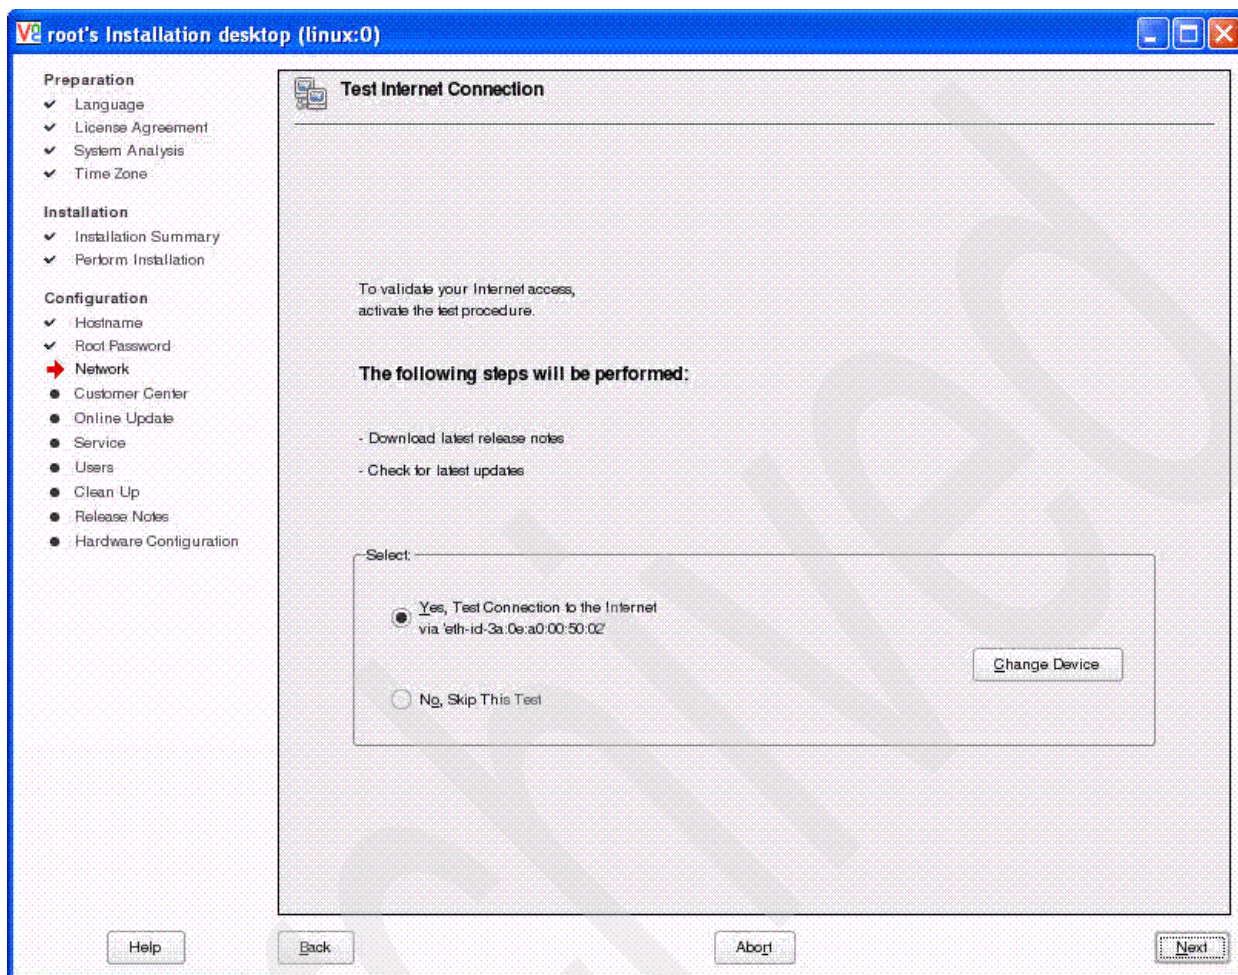


Figure 6-30 SLES 10 - Test Internet Connection

**Note:** The ability to download the release notes and updates will be dependent on the availability of an internet connection to the system.



16. Make the appropriate selection on the Test Internet Connection page and then click **Next**.

The installer now provides an opportunity to configure support for LDAP and security certificates as shown in Figure 6-31.

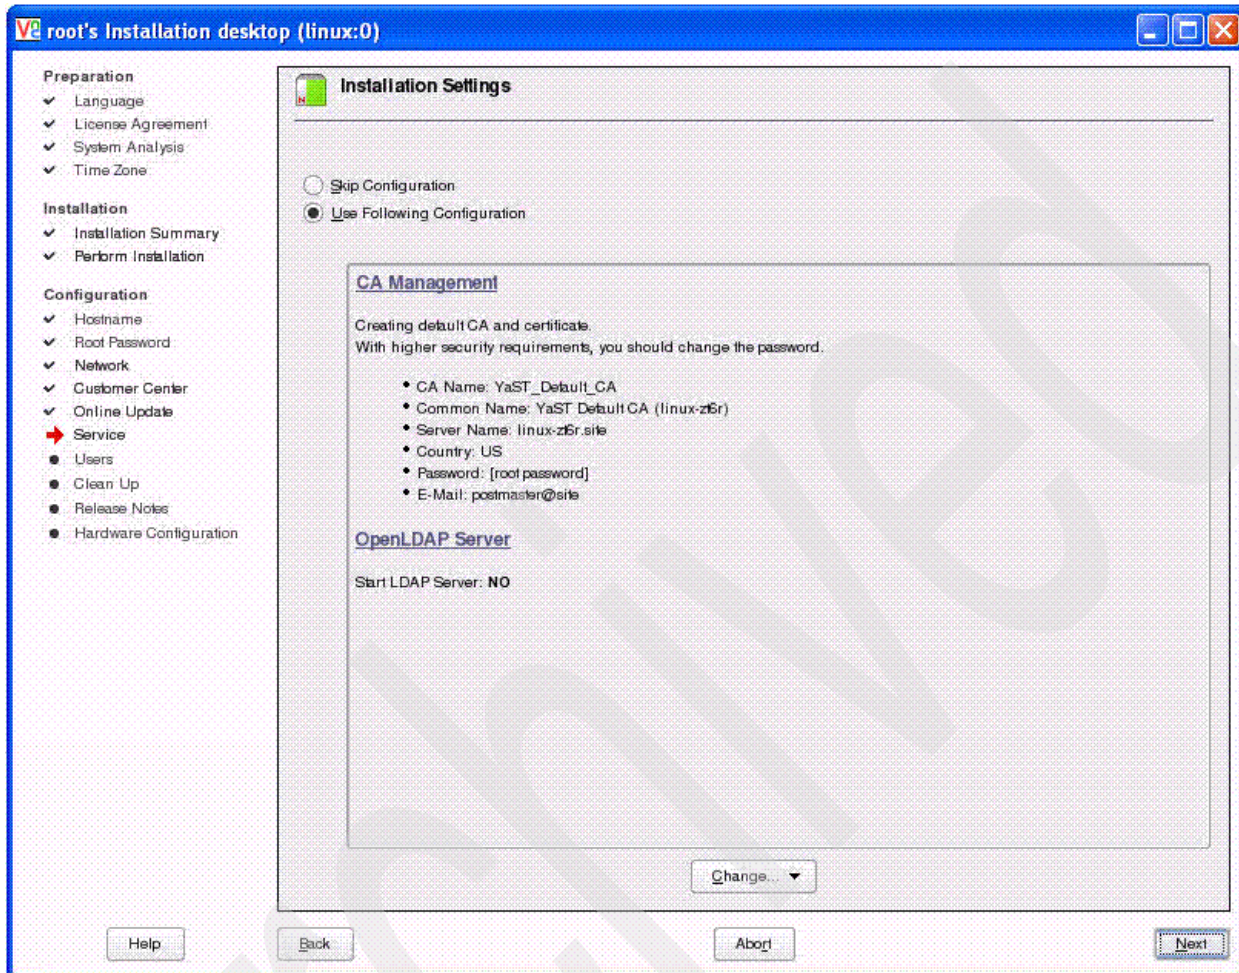


Figure 6-31 SLES 10 - Installation Settings

17. Make the appropriate selections on the Installation Settings screen for your environment and then click **Next**.

The next screen prompts for the definition of the user authentication mechanism that will be used for access to the operating system as shown in Figure 6-32.

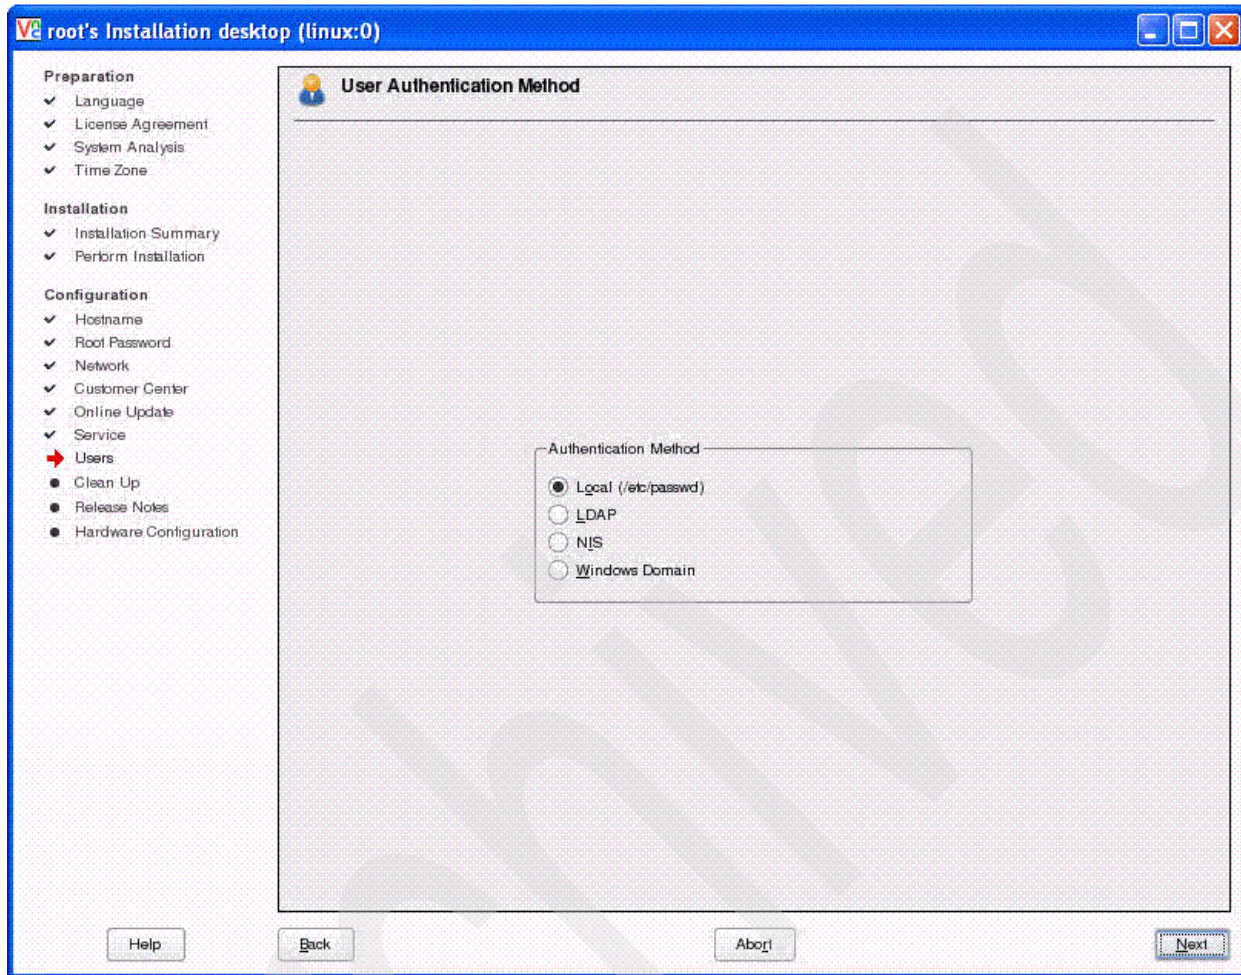


Figure 6-32 SLES 10 - User Authentication Method

These methods control how access to terminal sessions will be validated.

18. On the User Authentication Method screen, select the appropriate method for your environment (note you can change the method after the installation) and then click **Next**.

If Local is selected for the user authentication method, then a window will be displayed for creation of new users as shown in Figure 6-33.

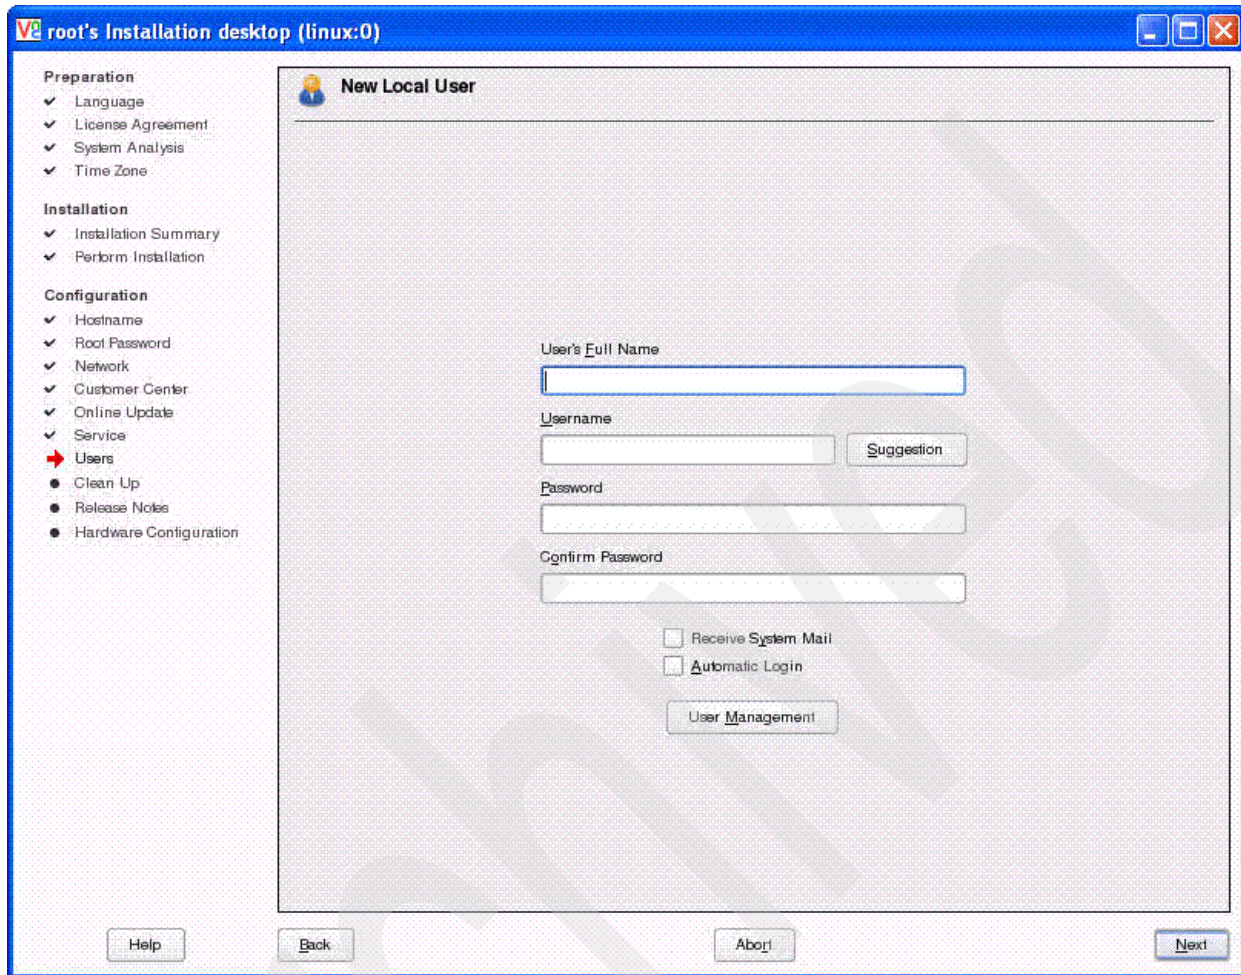


Figure 6-33 SLES 10 - New Local User

You can choose to create local users at this time or wait until after the installation is complete. Once user definition is complete, click **Next**.

Post-installation steps are almost complete. At this point, the system configuration files are updated as shown in Figure 6-34.

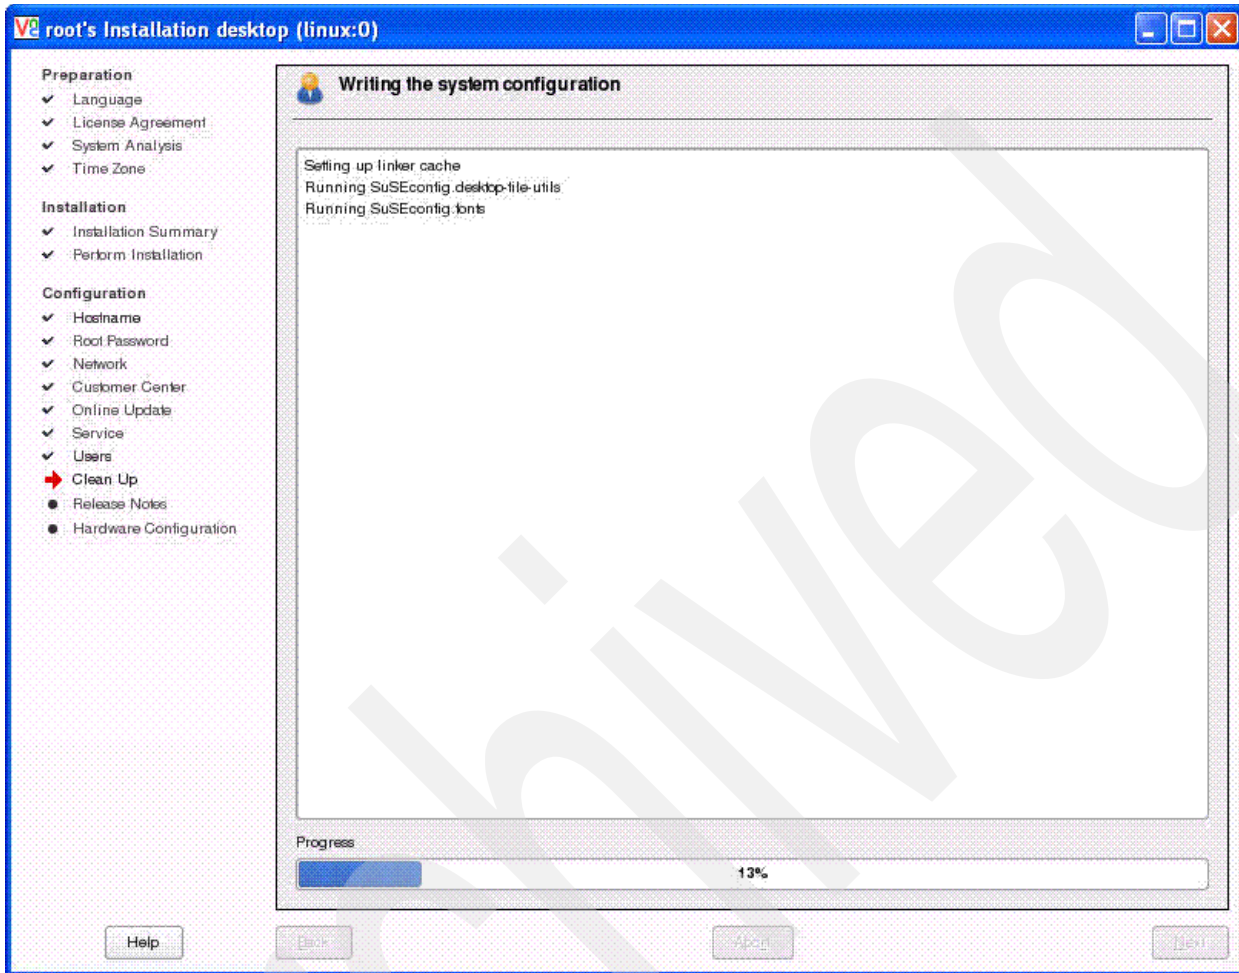


Figure 6-34 SLES 10 - Writing the system configuration

Once the system configuration has been updated, the release notes for the distribution are displayed as shown in Figure 6-35.

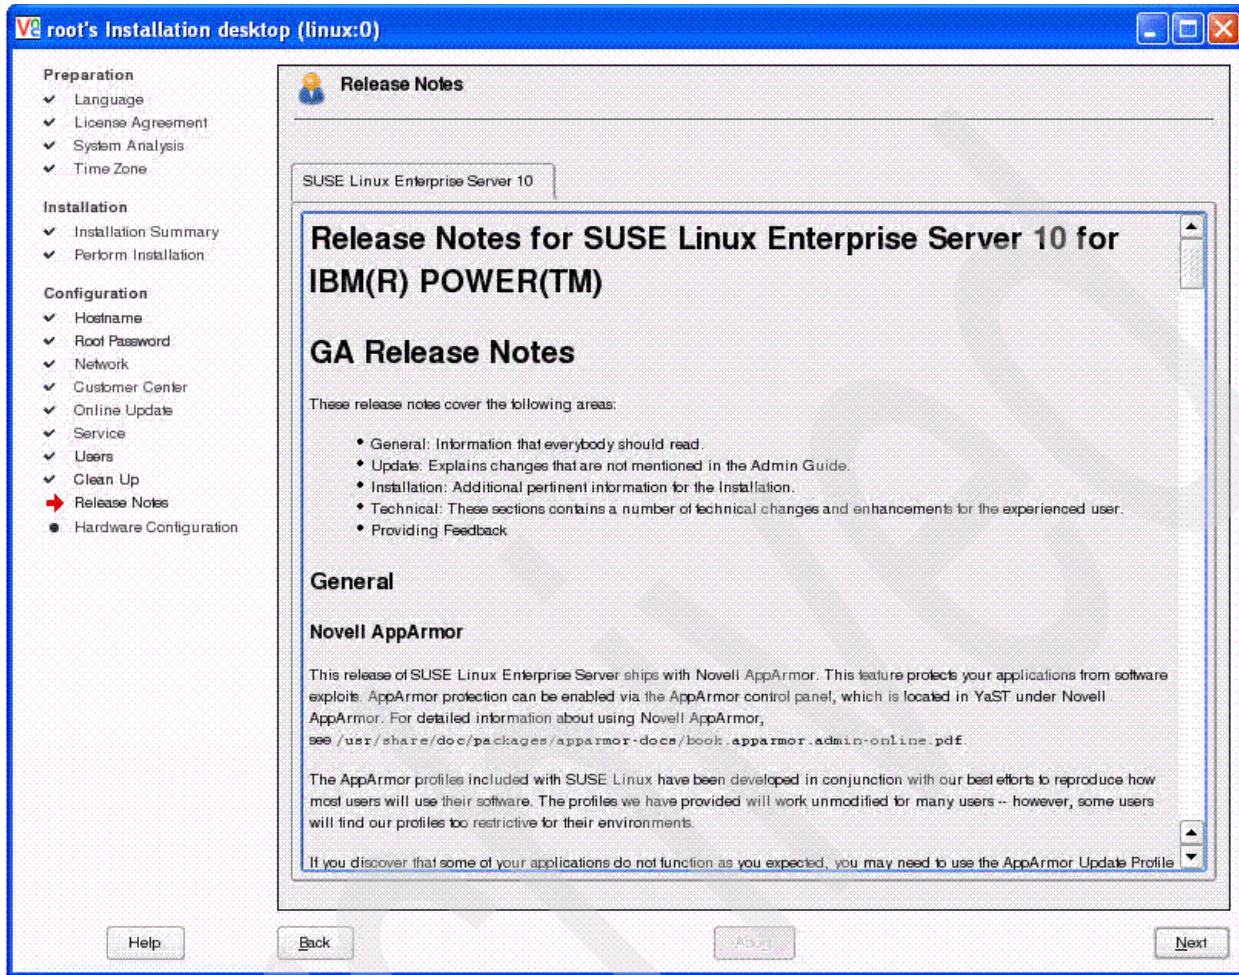


Figure 6-35 SLES 10 - Release Notes

19. Once you have reviewed the release notes, click **Next**.

At this point the system will attempt to detect printers connected to the system (Figure 6-36):



Figure 6-36 SLES 10 - Printer Detection

20. Click **Continue** on the Printer detection dialog box.

The Hardware Configuration window is displayed for configuration of printers (as shown in Figure 6-37).

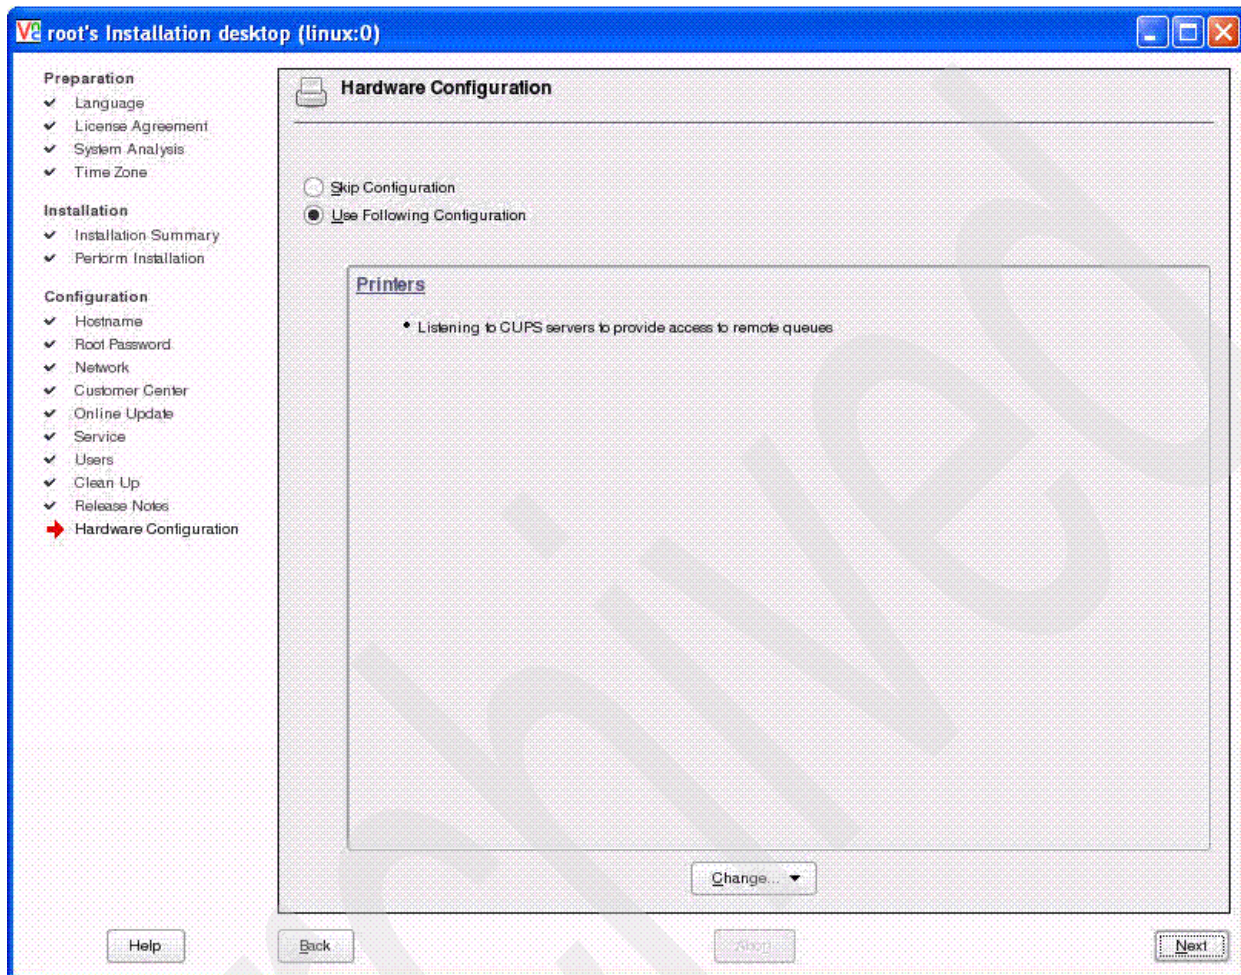


Figure 6-37 SLES 10 - Hardware Configuration

21. If you want the system to “listen” for network attached printers, then leave the configuration as-is and click **Next**. Otherwise, select **Skip Configuration** and then click **Next**.

The final screen that is displayed is the Installation Completed screen as shown in Figure 6-38.

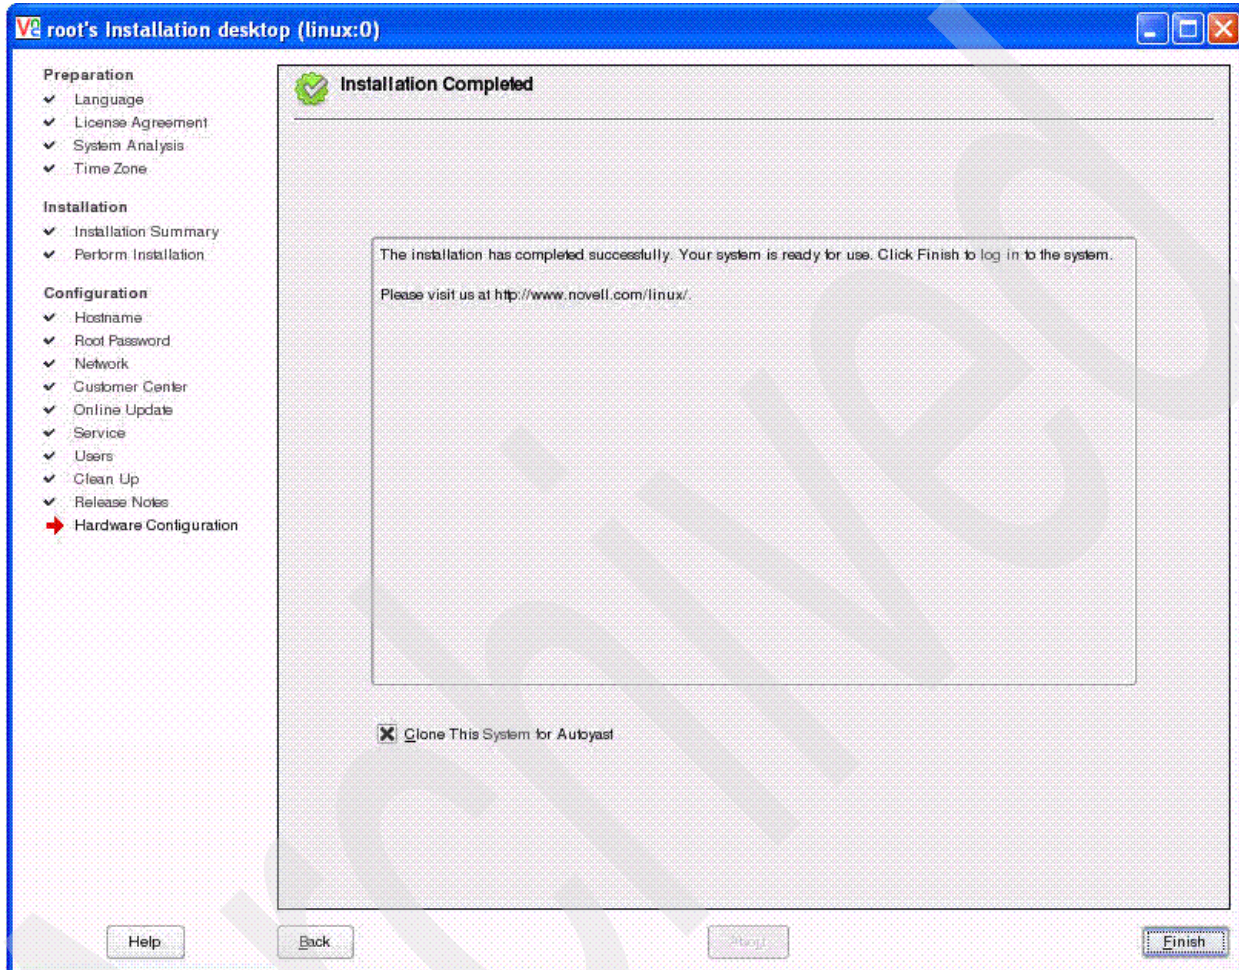


Figure 6-38 SLES 10 - Installation Completed

The Clone This System to Autoyast will create a save of the installed configuration that can be used for auto-installation of additional Linux instances. You can either leave the setting as is or disable it based on your environment needs.

22. Once the setting has been made on the Installation Completed screen, click **Finish**.

At this point, the installation is complete and the installed Linux will finish its initial boot.

## 6.2.2 Red Hat RHEL 4

At the time of writing of this book, RHEL 4 Update 3 was the latest Linux distribution available from Red Hat.

### Network server setting of interest

The IPL Parms field of the Network Server provides the capability to define parameters that will be passed directly to the Linux kernel. The distributor provided documentation will provide

a detailed list of possible parameters. One parameter of interest is *vnc*. This parameter instructs the installation program to perform the installation in a graphical mode making use of the Virtual Network Computing server/client technology.

## Installation steps

**Note:** The following example assumes that the Network Server has been configured to boot off of the installation kernel (typically /qopt/RED\_HAT/IMAGES/PSERIES/NETBOOT.IMG), the IPL parameters line includes 'vnc', the virtual console has been accessed, and you have access to the VNC client which you can download from the Web at:

<http://www.realvnc.com>

You can also perform a text-based installation by omitting the vnc value from the IPL parameters field of the Network Server.

1. On the i5/OS hosting partition, issue the following command to vary on the Network Server as shown in Example 6-4.

*Example 6-4 Vary on of Network Server*

```
VRYCFG CFGOBJ(RHEL4) CFGTYPE(*NWS) STATUS(*ON)
```

The value for CFGTYPE is the name of the Network Server.

At this point the logical partition is started and the installation program begins. After a few seconds, the first screen of the installer is displayed as shown in Figure 6-39.

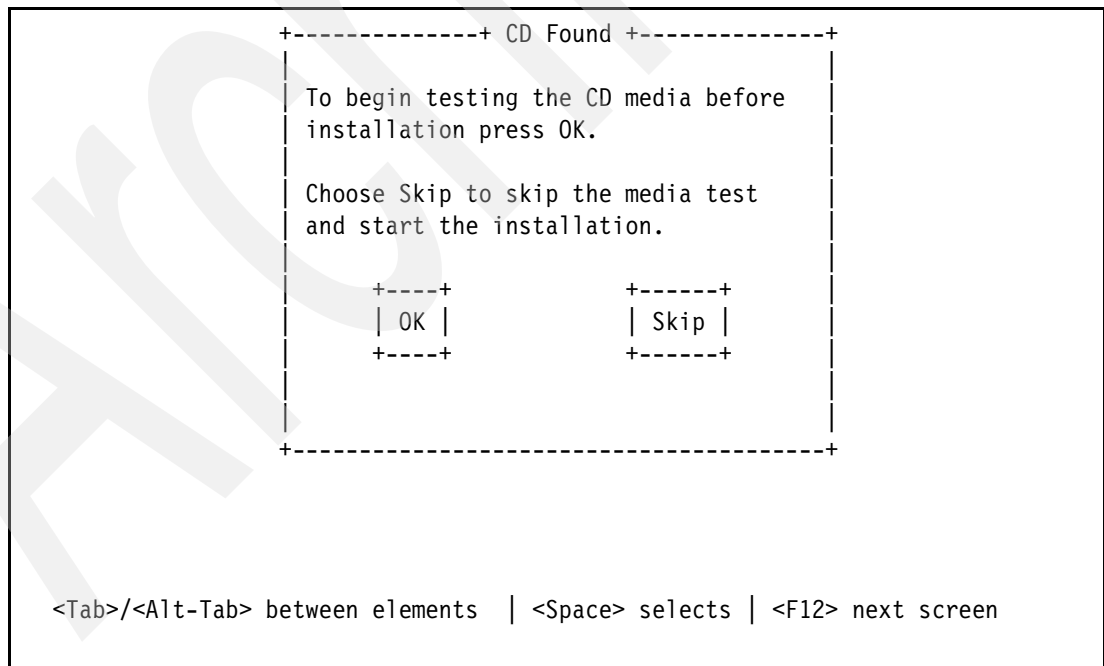


Figure 6-39 RHEL 4 - Media test



2. Use the interface to test the validity of the installation CDs.

Once the CDs have been validated, the Language selection screen is displayed as shown in Figure 6-40.

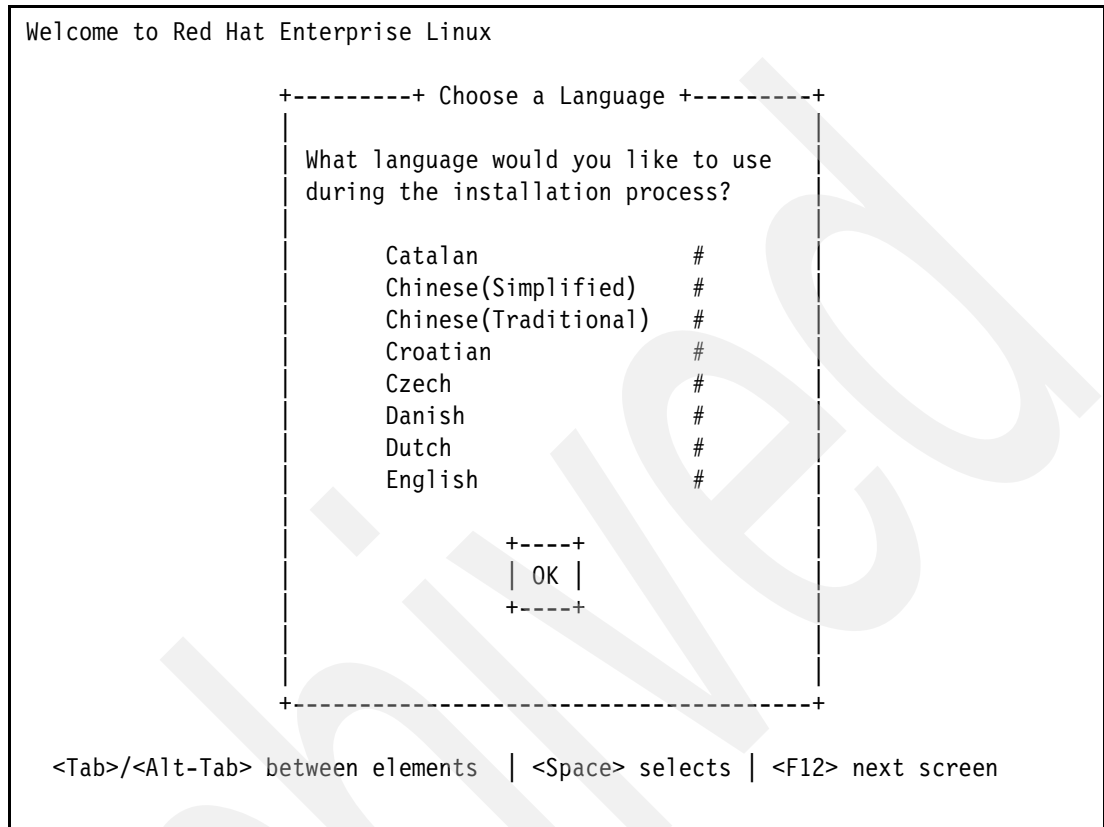


Figure 6-40 RHEL 4 - Language selection

3. The language selected on the Language selection screen (Figure 6-40) is the language that will be used for the remainder of the installation process. Use the up and down arrow keys to position the cursor on the desired language and then press **Tab** to highlight [OK] and press **Enter**.

The Configure TCP/IP screen (Figure 6-41) provides the ability to configure the network parameters for the first network interface identified by the installer. This screen is displayed when the IPL parameters indicate that a graphical-based installation is to be

performed (that is, vnc) or when the installation method was identified as a network-based installation.

```
Welcome to Red Hat Enterprise Linux

+-----+ Configure TCP/IP +-----+
|
| Please enter the IP configuration for this machine. Each
| item should be entered as an IP address in dotted-decimal
| notation (for example, 1.2.3.4).
|
| [] Use dynamic IP configuration (BOOTP/DHCP)
|
| IP address: 9.5.110.85__
| Netmask: 255.255.255.248_
| Default gateway (IP): 9.5.110.81__
| Primary nameserver: _____
|
| +----+ +-----+
| | OK | | Back |
| +----+ +-----+
|
+-----+

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen
```

Figure 6-41 RHEL 4 - Configure TCP/IP

4. Provide the requested values for the network parameters. Once the values are complete, press **Tab** until [OK] is highlighted and then press **Enter**.

At this point the installer configures and starts the network interface and then starts the VNC server to enable the remainder of the installer to be completed via a Graphical User Interface. Once the VNC server has been started, connection information similar to Figure 6-42 is displayed.

```
Starting VNC...

WARNING!!! VNC server running with NO PASSWORD!
You can use the vncpassword=<password> boot option
if you would like to secure the server.

The VNC server is now running.
Please connect to 9.5.110.85:1 to begin the install...
Starting graphical installation...
XKB extension not present on :1
```

Figure 6-42 RHEL 4 - VNC startup messages

- To complete the installation, start the VNC viewer client and provide the server connection details provided by the installer, as shown in Figure 6-43.




Figure 6-43 RHEL 4 - VNC viewer connection details

- Once the connection details have been provided, click **OK**.

At this point the graphical portion of the installer is started and the welcome screen is displayed as shown in Figure 6-44.



Figure 6-44 RHEL 4 - Welcome screen

**Note:** Make sure not to close the virtual console connection as it will be required for a later portion of the installation.

- Click **Next** to continue the installation.

8. On the 'Disk Partitioning Setup' screen (shown in Figure 6-45), leave the setting at Automatically partition and then click **Next**.

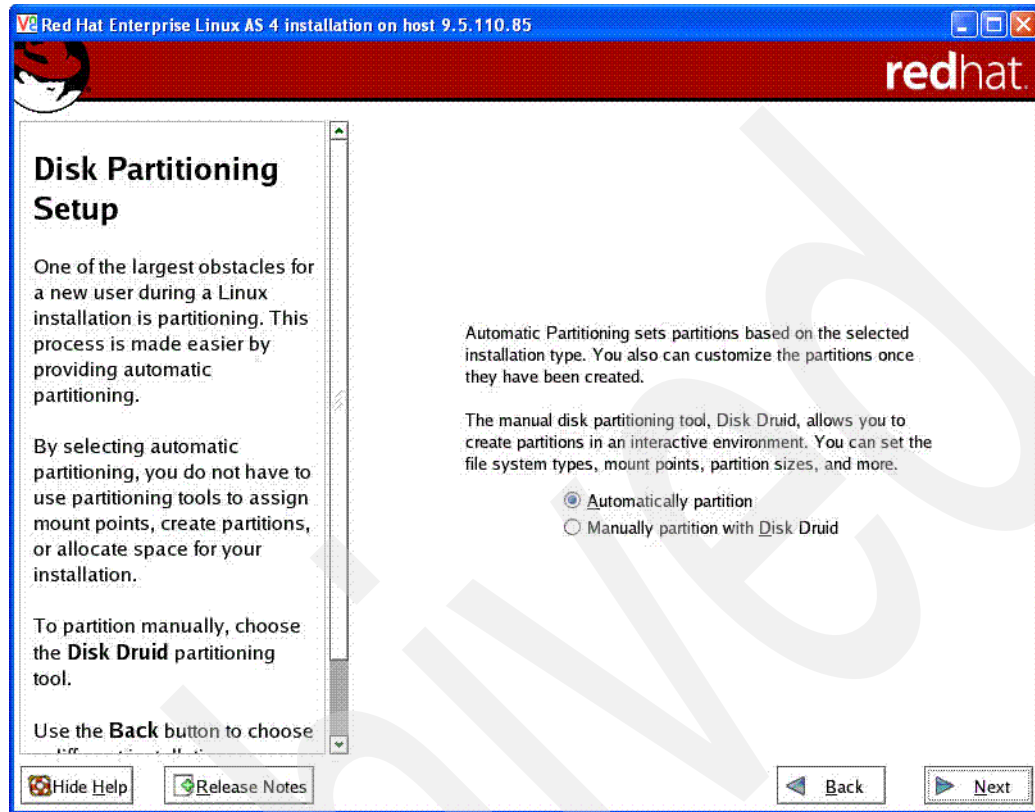


Figure 6-45 RHEL 4 - Disk partitioning setup

9. The Automatic Partitioning screen (Figure 6-46) allows you to specify how any existing partitions on the disk should be handled. Typically, the setting will be left at Remove all partitions on this system. Additionally, make sure to leave the Review (and modify if needed) the partitions created option set and then click **Next**.

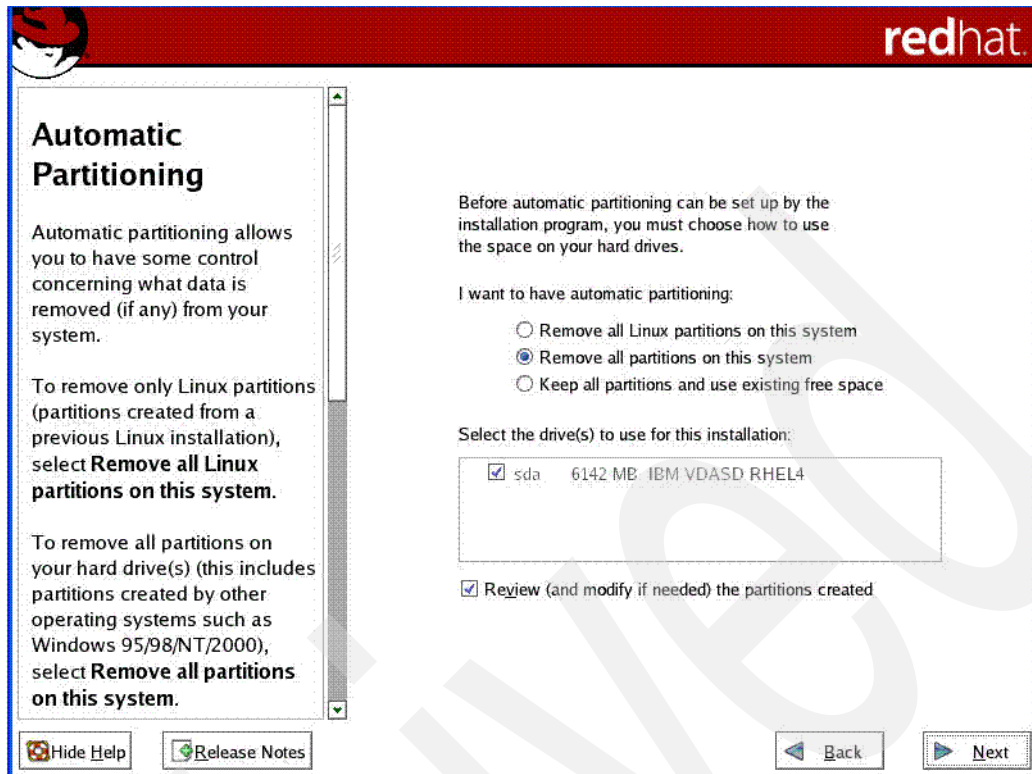


Figure 6-46 RHEL 4 - Automatic partitioning

Since you selected to remove the disk partitions, a dialog box prompting for confirmation is displayed as shown in Figure 6-47.



Figure 6-47 RHEL 4 - Partition deletion warning

10. Click **Yes** on the dialog box as shown in Figure 6-47 to verify that all of the existing partitions should be deleted.

The next screen that is displayed is the Disk Setup screen as shown in Figure 6-48.

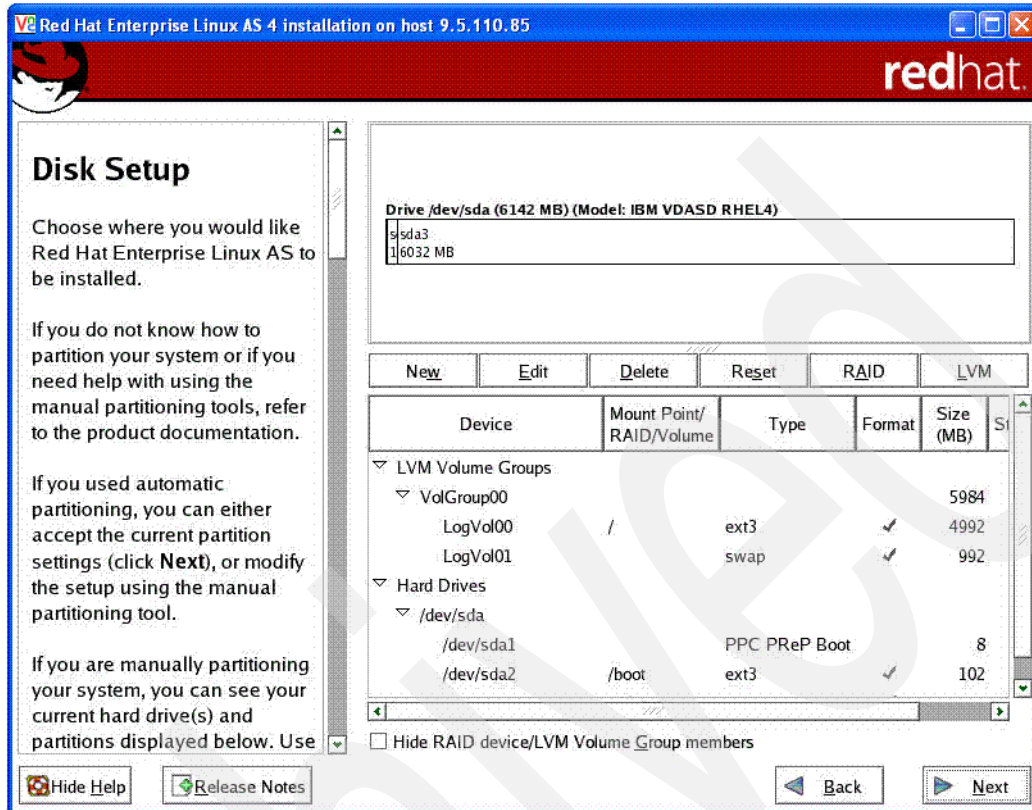


Figure 6-48 RHEL 4 - Disk Setup

The partitioning used by the installer is sufficient for a successful installation and it uses Logical Volume Manager (LVM) for the root (/) partition so that resizing of the disk partition is possible. You can either leave the partitioning as is or change it to better configure the disk for resizing of the virtual storage space.

Perform the following steps to repartition the disk:

**Note:** If you want to leave the partitioning of the disk as provided by the installer, then click **Next** and skip to step 11 on page 220.

- a. Delete the existing partitions. This actually should be accomplished from the bottom to the top. Select and delete the partitions in the following order:
  - i. /dev/sda2            /boot                    ext3
  - ii. /dev/sda1                               PPC PReP Boot
  - iii. LogVol01                               swap
  - iv. LogVol00                               ext3
  - v. VolGroup00
  - vi. /dev/sda1                               LVM PV

Now that all of the partitions have been deleted, the four partitions required by the installer should be created.

- b. Create the PReP Boot partition. Click the **New** button and complete the fields of the Add Partition dialog as shown in Figure 6-49.

The screenshot shows the 'Add Partition' dialog box with the following configuration:

- Mount Point:** <Not Applicable>
- File System Type:** PPC PReP Boot
- Allowable Drives:**  sda, 6142 MB, IBM VDASD RHEL4
- Size (MB):** 8
- Additional Size Options:**
  - Fixed size
  - Fill all space up to (MB): 8
  - Fill to maximum allowable size
- Force to be a primary partition

Buttons:

Figure 6-49 RHEL 4 - PReP boot partition creation

**Note:** In order to ensure that the disk partitions are placed on the disk in the correct order to support possible later resizing, it is imperative that on this and the remaining disk creation steps that the 'Force to be a primary partition' item is selected.

- c. Once the fields for the PReP Boot partition have been completed, click **OK**.

- d. Click **New** on the Disk Setup screen to create the /boot partition. Complete the fields as shown in Figure 6-50.

The screenshot shows the 'Add Partition' dialog box with the following settings:

- Mount Point: /boot
- File System Type: ext3
- Allowable Drives:  sda 6142 MB IBM VDASD RHEL4
- Size (MB): 102
- Additional Size Options:
  - Fixed size
  - Fill all space up to (MB): 102
  - Fill to maximum allowable size
- Force to be a primary partition

Buttons: Cancel, OK

Figure 6-50 RHEL 4 - /boot partition creation

- e. Once the fields for the /boot partition have been completed, click **OK**.
- f. Click **New** on the Disk Setup screen to create the swap partition. Complete the fields as shown in Figure 6-51.

The screenshot shows the 'Add Partition' dialog box with the following settings:

- Mount Point: <Not Applicable>
- File System Type: swap
- Allowable Drives:  sda 6142 MB IBM VDASD RHEL4
- Size (MB): 512
- Additional Size Options:
  - Fixed size
  - Fill all space up to (MB): 512
  - Fill to maximum allowable size
- Force to be a primary partition

Buttons: Cancel, OK

Figure 6-51 RHEL 4 - swap partition creation



- g. Once the fields for the swap partition have been completed, click **OK**.
- h. Click **New** on the Disk Setup screen to create the root (/) partition. Complete the fields as shown in Figure 6-52.

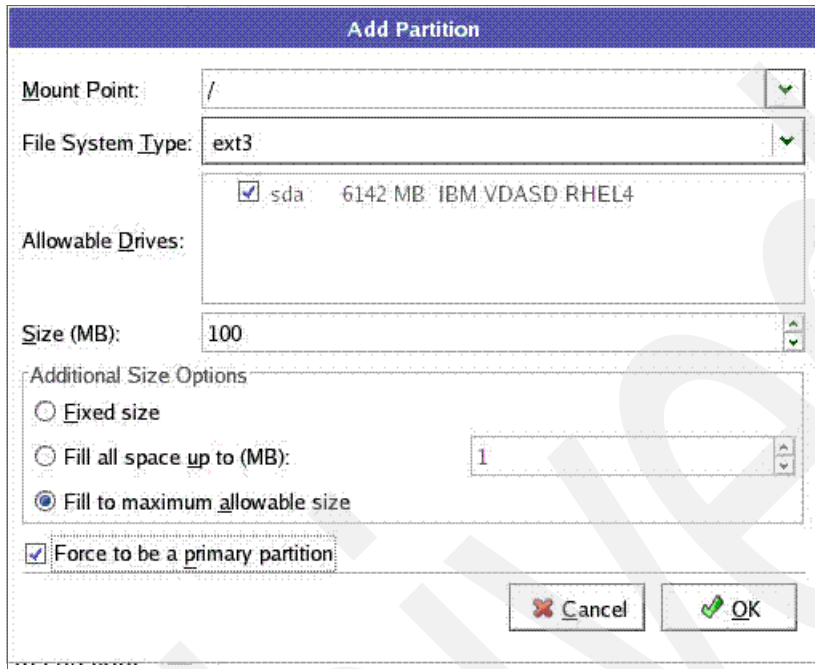


Figure 6-52 RHEL 4 - root partition creation

- i. Once the fields for the root partition have been completed, click **OK**.

Once all of the partitions have been created, the disk layout should resemble Figure 6-53.

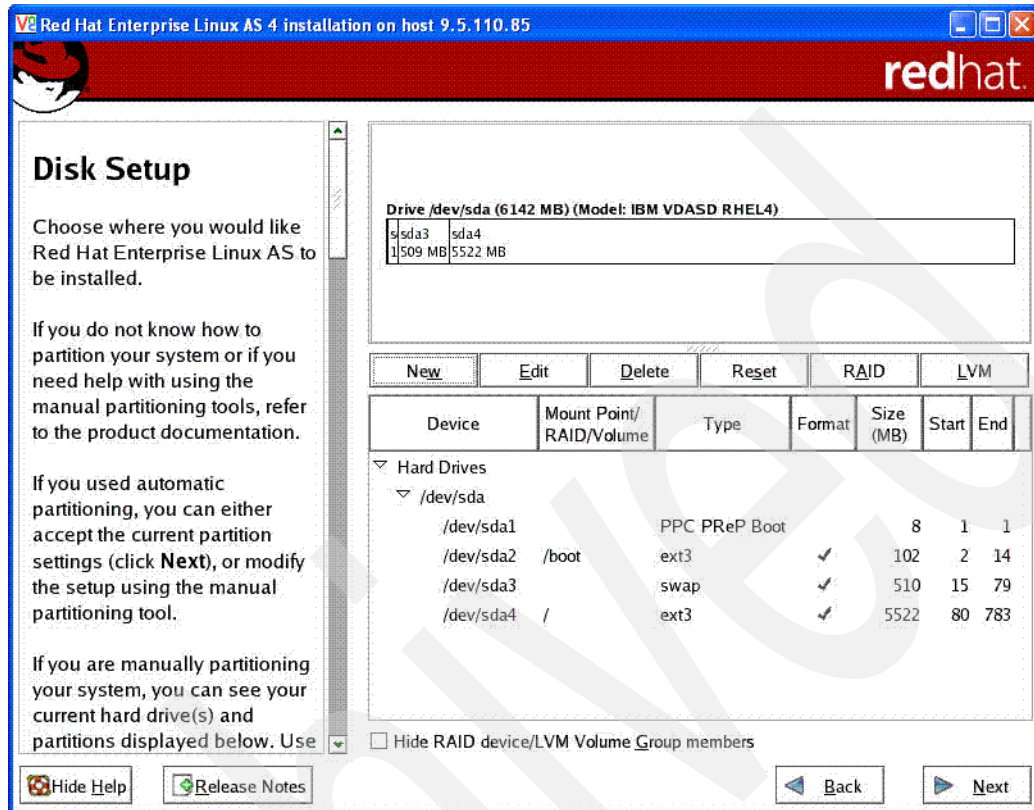


Figure 6-53 RHEL 4 - Disk setup (complete)

- j. Click **Next** to accept the disk partitioning.
11. The Network Configuration screen shown in Figure 6-54 provides an opportunity to configure the network parameters for the system. Since the network parameters were provided for the VNC based install those settings have been reflected in the network

parameters on this screen. Make any necessary changes for your environment (if any) and then click **Next**.

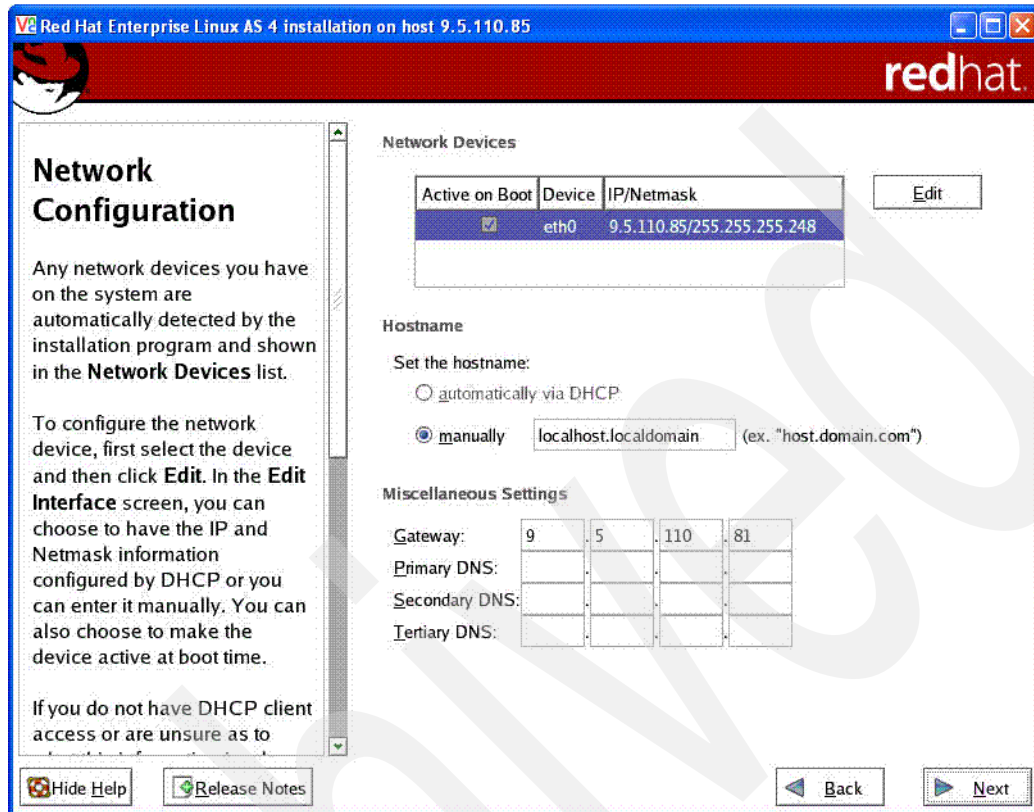


Figure 6-54 RHEL 4 - Network Configuration

The Firewall Configuration screen shown in Figure 6-55 provides an opportunity to configure security settings for the system.

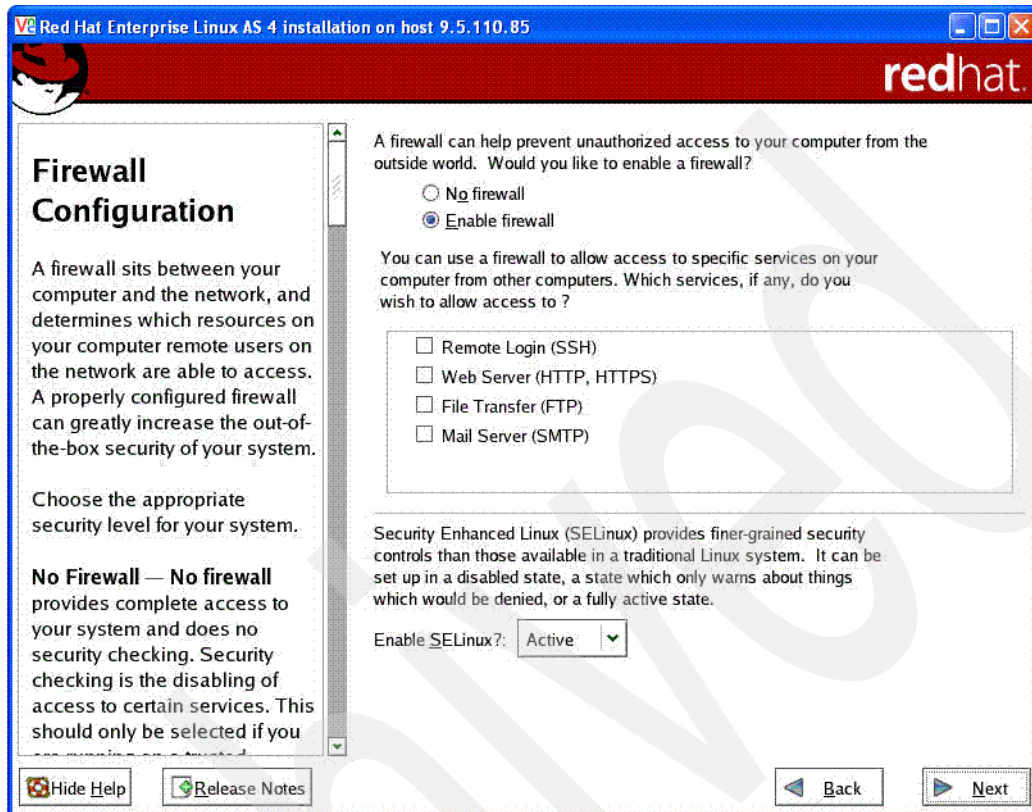


Figure 6-55 RHEL 4 - Firewall configuration

The settings you choose on the Firewall Configuration screen are dependent on the environment that you are installing Linux into, the function that Linux is going to perform, as well as the security mechanisms you wish to employ.

The firewall function provides for the ability to secure network access to services provided by Linux based on characteristics of the network traffic (such as source IP address, destination ports, and so on). The SELinux (Secure Enterprise Linux) function is a discretionary access control (DAC) security mechanism implemented in Linux that adds a layer of security based on functions, capabilities, and security containers.

**Note:** SELinux represents additions to the Linux kernel that have either been developed or heavily influenced by the National Security Agency / Central Security Service. The implementation of the SELinux security model is usually used by companies bidding or working on United States Government - Department of Defense contracts.

For our example, we will disable the firewall (No firewall) as well as SELinux (Disabled).

12. Once you have set the Firewall and SELinux items to the desired values for your environment, click **Next**.

13. The Additional Language Support screen shown in Figure 6-56 is used to define additional language support for the installation. Select any additional language support that you want for the installed Linux instance and then click **Next**.

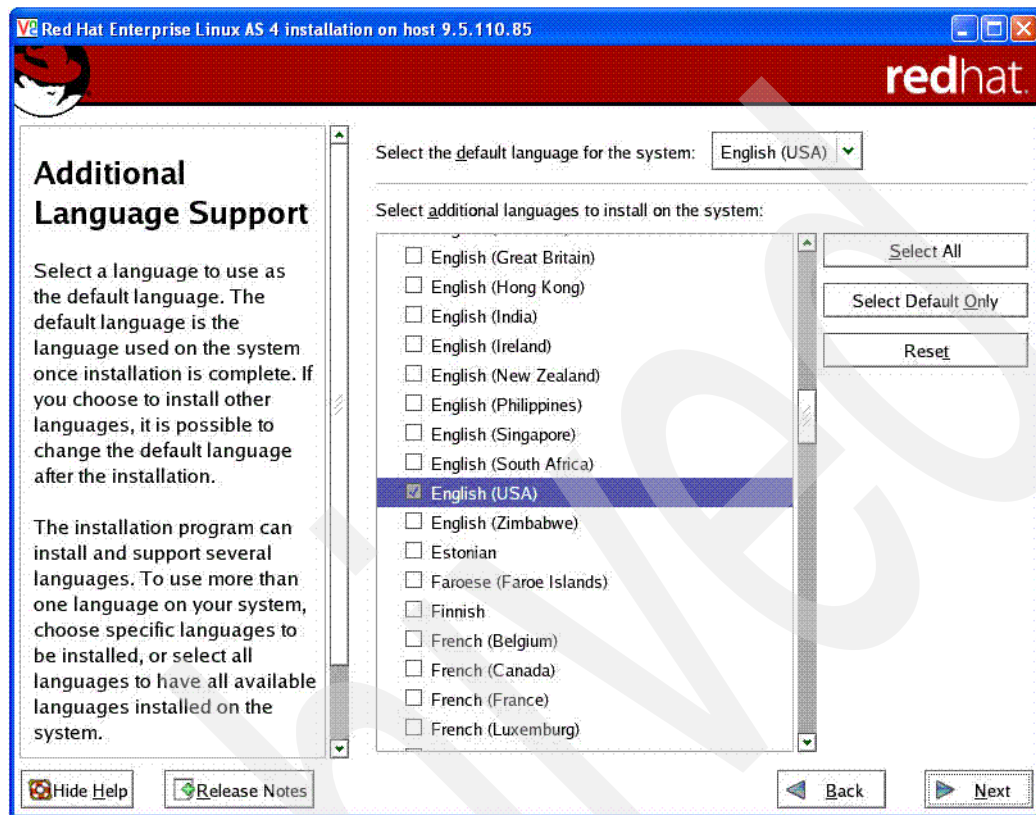


Figure 6-56 RHEL 4 - Additional language support

14. The Time Zone Selection screen shown in Figure 6-57 is used to define the time zone for the system. Choose your time zone and then click **Next**.

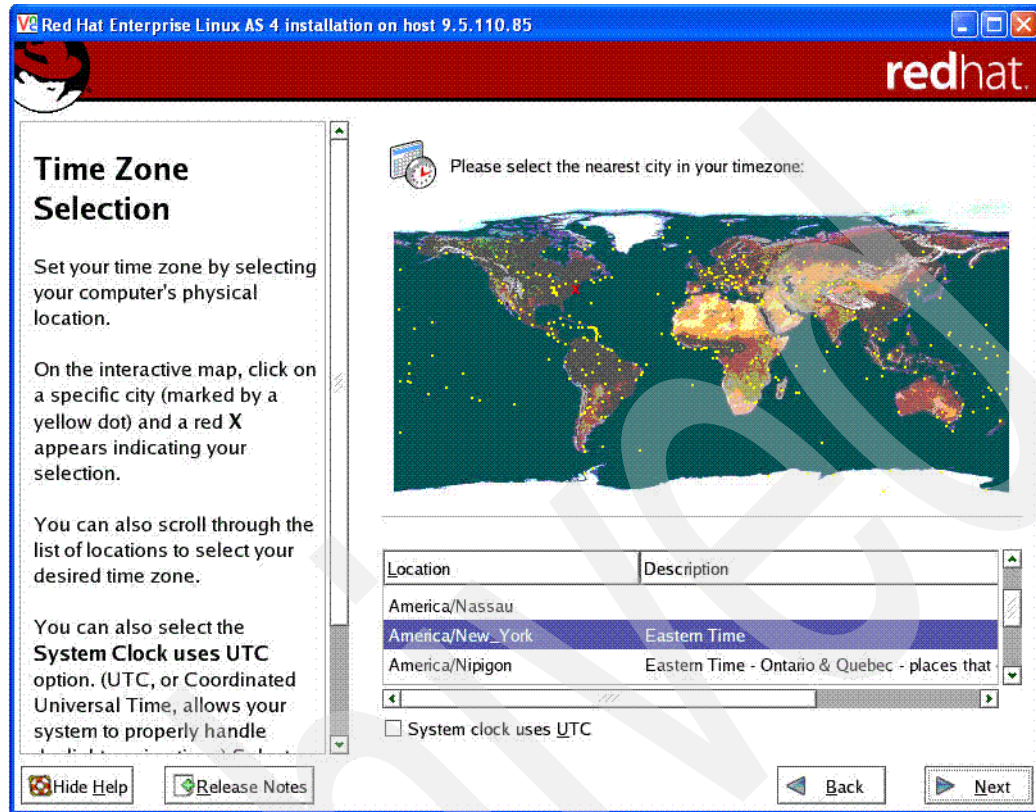


Figure 6-57 RHEL 4 - Time zone selection

15. The Set Root Password screen (Figure 6-58) is used to set the root password for the system. The root user is the user ID that is used for system administration functions. Set and confirm the password and then click **Next**.



Figure 6-58 RHEL 4 - Set root password

The Package Installation Defaults screen (Figure 6-59) shows the default package scheme selected by the installer based on the size of the disk allocated to Linux. You can

either accept the defaults or modify the package selection by selecting **Customize software packages to be installed**.

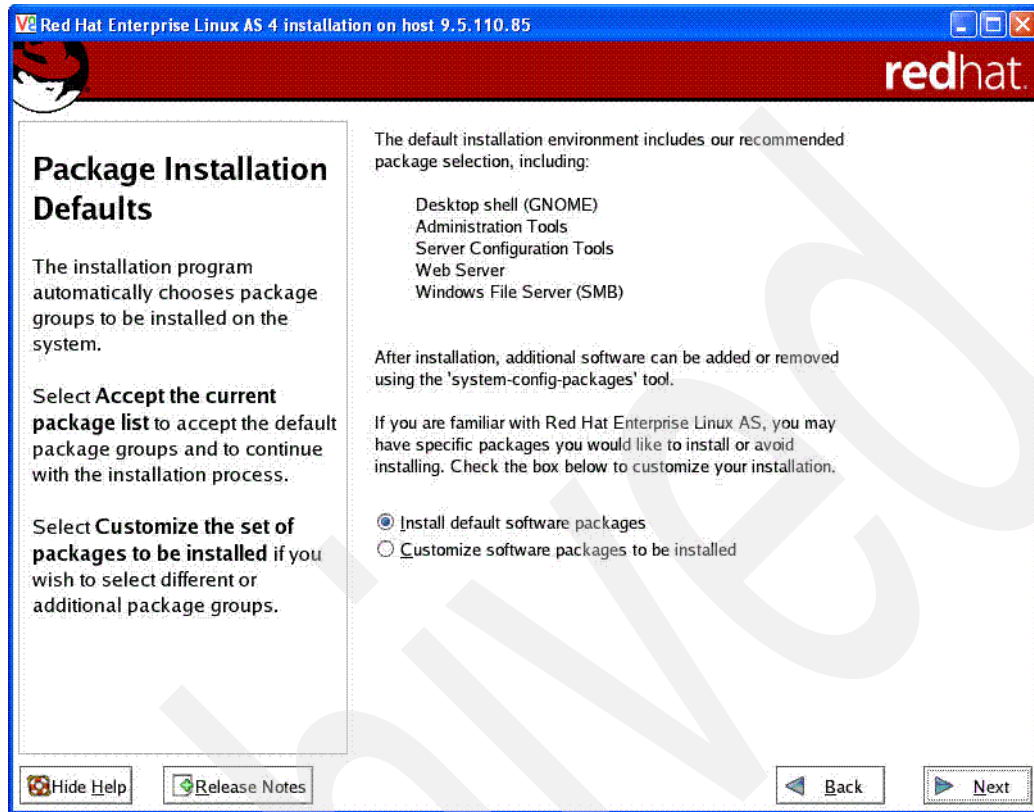


Figure 6-59 RHEL 4 - Package Installation Defaults



16. Once the package selection is complete, click **Next**.

At this point, all of the information needed to perform the installation has been collected and a final screen is displayed as shown in Figure 6-60.



Figure 6-60 RHEL 4 - About to Install

17. Click **Next** on the “About to Install” screen to start the installation of packages.

A dialog box will be displayed listing the required installation media as shown in Figure 6-61.

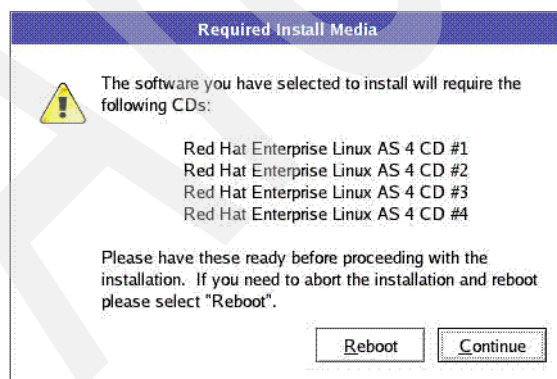


Figure 6-61 RHEL 4 - Required install media

18. Make sure that you have access to all of the required media and then click **Continue**. Progress of the installation of packages will be displayed as shown in Figure 6-62.

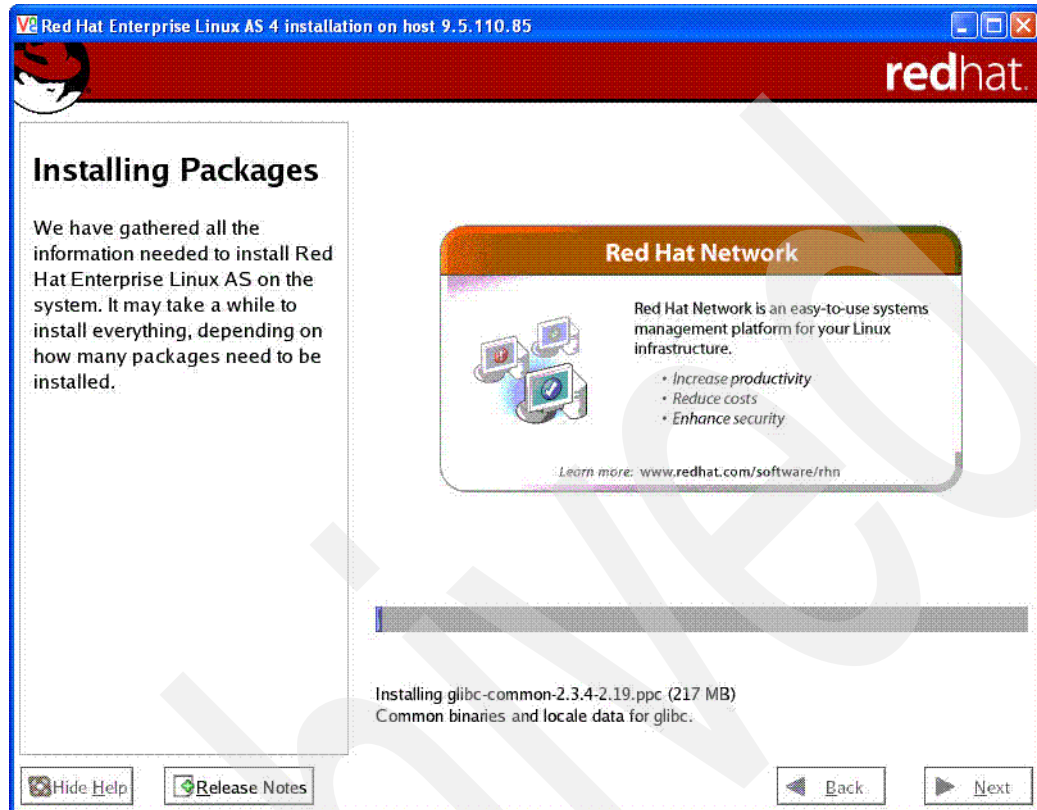


Figure 6-62 RHEL 4 - Installing packages

Prompts for additional CDs will be displayed as required as shown in Figure 6-63.



Figure 6-63 RHEL 4 - Change CDROM

Once all of the packages have been installed a completion message will be displayed as shown in Figure 6-64.

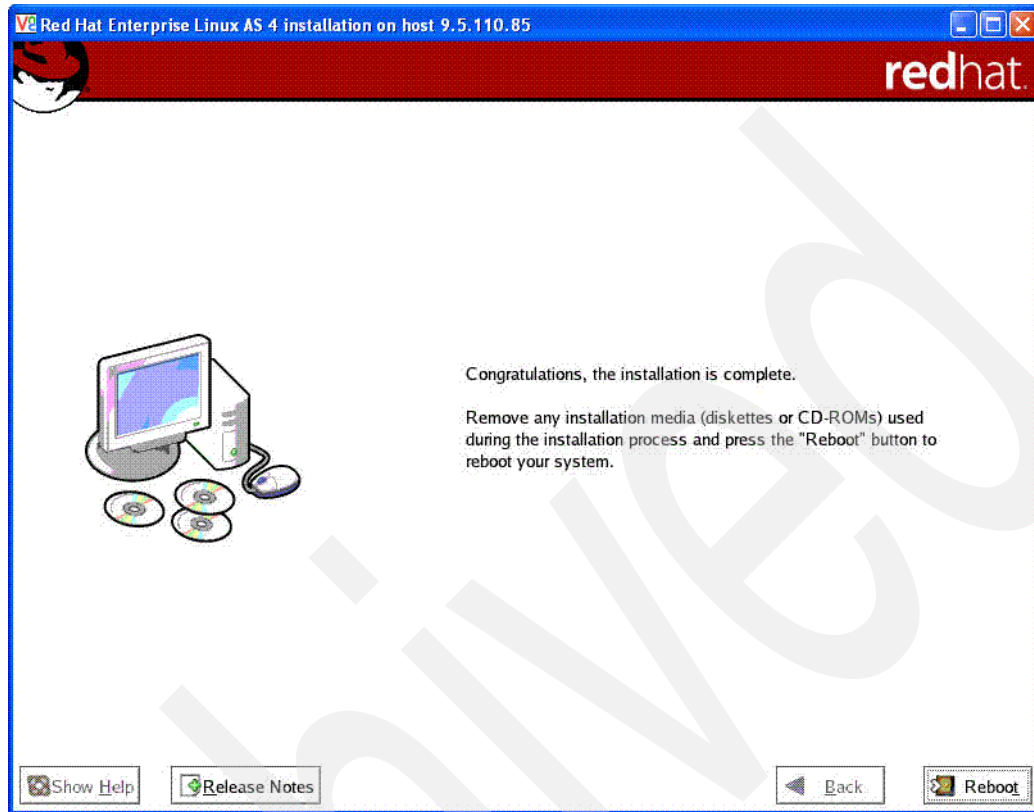


Figure 6-64 RHEL 4 - Installation Complete

19. Click **Reboot** on the Installation Complete window.

**Note:** Ensure that you have access to the virtual console for the Linux partition that has just been installed.

At this point the server is restarted. Since the Network Server Description is still configured to boot off of the installation CD the installation program will be restarted. At this point, you should vary off the Network Server and change it to boot off of the kernel installed into the virtual storage space.

20. From the hosting i5/OS partition, issue the following command shown in Example 6-5 to vary off the Network Server.

*Example 6-5 Vary off Network Server*

```
VRYCFG CFGOBJ(RHEL4) CFGTYPE(*NWS) STATUS(*OFF)
```

The value for CFGOBJ is the name of your Network Server.

The packages required to support shutdown of the operating system are not yet installed so the vary off of the network server will not successfully shutdown the operating system and the logical partition.

21. From the HMC, right-click the Linux partition and select **Shut Down Partition**.

22. From the Shut Down Partitions dialog box shown in Figure 6-65, select **Immediate** and then click **OK**.

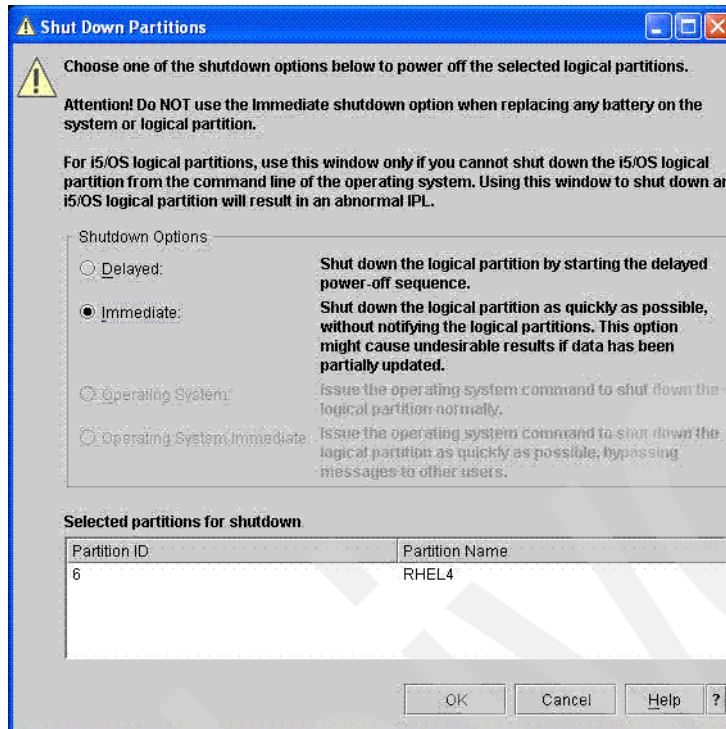


Figure 6-65 RHEL 4 - Shut Down Partitions

23. On the Shutdown confirmation dialog box shown in Figure 6-66, click **Yes** to confirm the shutdown request.

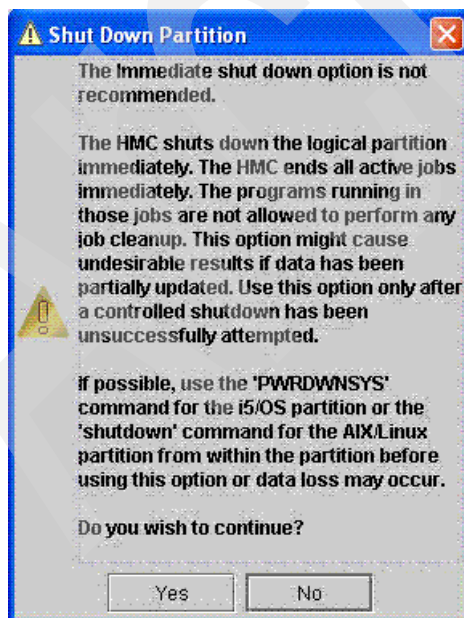


Figure 6-66 RHEL 4 - Shut Down Partition confirmation

24. Depending on how recent your version of the HMC is, an additional dialog box concerning replacement of a cache battery may be displayed (Figure 6-67). If this dialog box is displayed, click **No** to indicate that a cache battery is not being replaced.



Figure 6-67 RHEL 4 - Cache battery replacement

At this point, the settings of the Network Server Description can be changed and the Network Server can be varied back on.

25. In the i5/OS hosting partition, issue the following command shown in Example 6-6 to modify the Network Server Description to boot off the installed kernel.

*Example 6-6 Modify Network Server Description -- boot off of Network Storage*

```
CHGNWSD NWS(D(RHEL4) IPLSRC(*NWSSTG) IPLSTMF(*NONE) IPLPARM(*NONE)
```

The value for the NWS parameter indicates the name of the Network Server Description.

26. In the hosting i5/OS partition issue the following command shown in Example 6-7 to vary the Network Server back on.

*Example 6-7 Vary on Network Server*

```
VRYCFG CFGOBJ(RHEL4) CFGTYPE(*NWS) STATUS(*ON)
```

The value for the CFGOBJ parameter indicates the name of the Network Server.

At this point, the initial boot of the Linux partition is performed and a setup agent is displayed as shown in Figure 6-68.

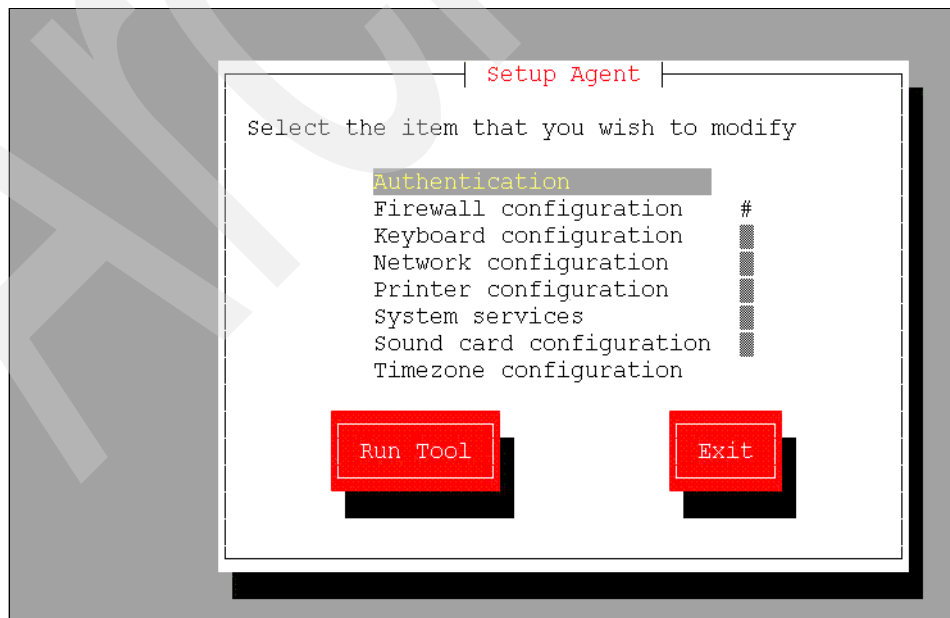


Figure 6-68 RHEL 4 - Setup agent

You can use this dialog box to perform additional configuration tasks at this time. If you do nothing, the dialog box will be dismissed after approximately 20 seconds and Linux will complete booting.

### 6.2.3 Installation of Linux on POWER Utilities

There are a number of additional Linux on POWER utilities that should be installed to support specific functions of running Linux on a POWER based system including support of Power Off function via vary off of the Network Server as well as support for the Dynamic LPAR function.

**Note:** These additional packages needs to be installed regardless of whether you installed the Novell / SuSE distribution or the RedHat distribution of Linux.

The utilities can be downloaded as RPM packages from the “Service and productivity tools” site on the Web at:

<https://www14.software.ibm.com/webapp/set2/sas/f/1opdiags/home.html>

You can either download the packages to a local client and then ftp them to the Linux system or you can download them directly to the Linux system.

The initial Web page lists the Linux distributions as well as the possible hardware platforms they could be installed on as shown in Figure 6-69.

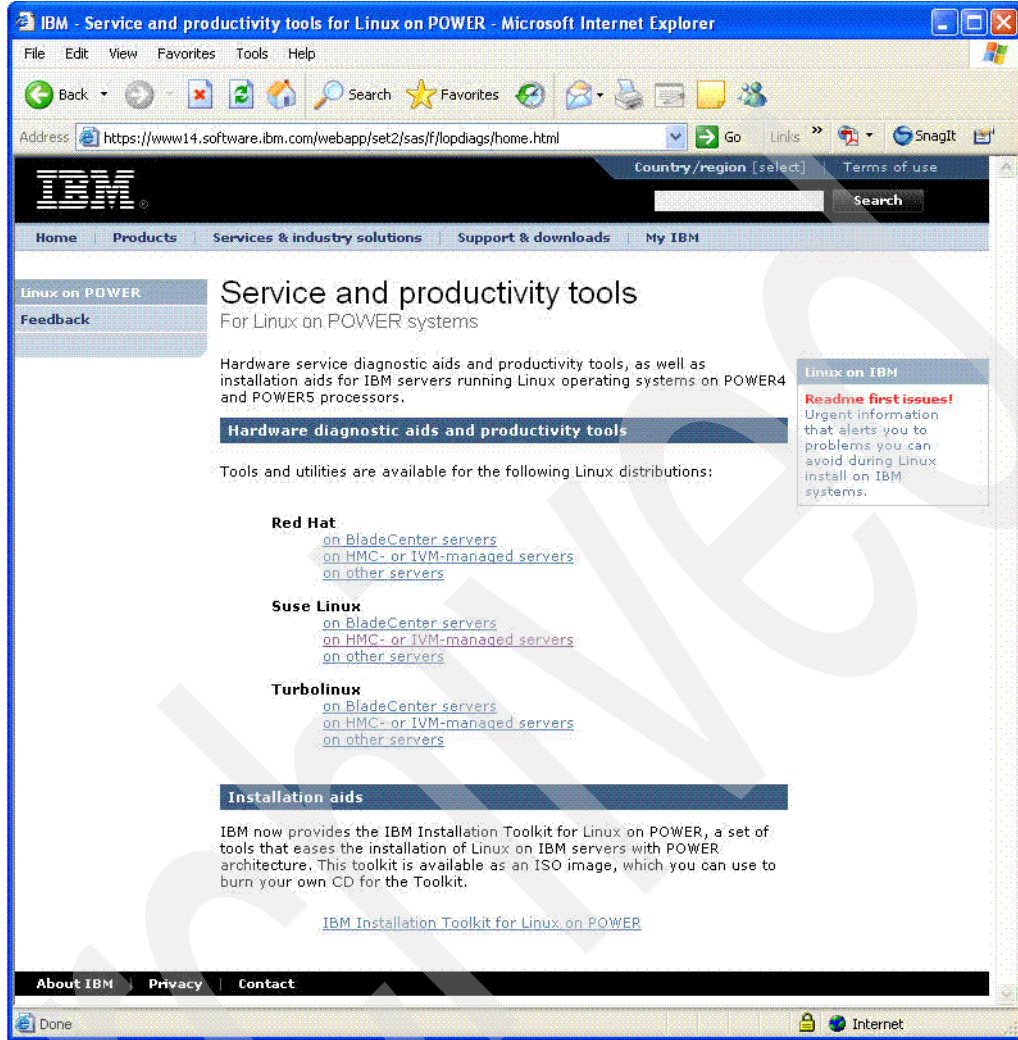


Figure 6-69 Service and productivity tools Web page

Select the link for the appropriate distribution and platform. In this case, you would select on **HMC- or IVM-managed servers** under either the “Red Hat” or “Suse Linux” heading

(depending on which distribution you have installed). This list of available packages is displayed as shown in Figure 6-70.

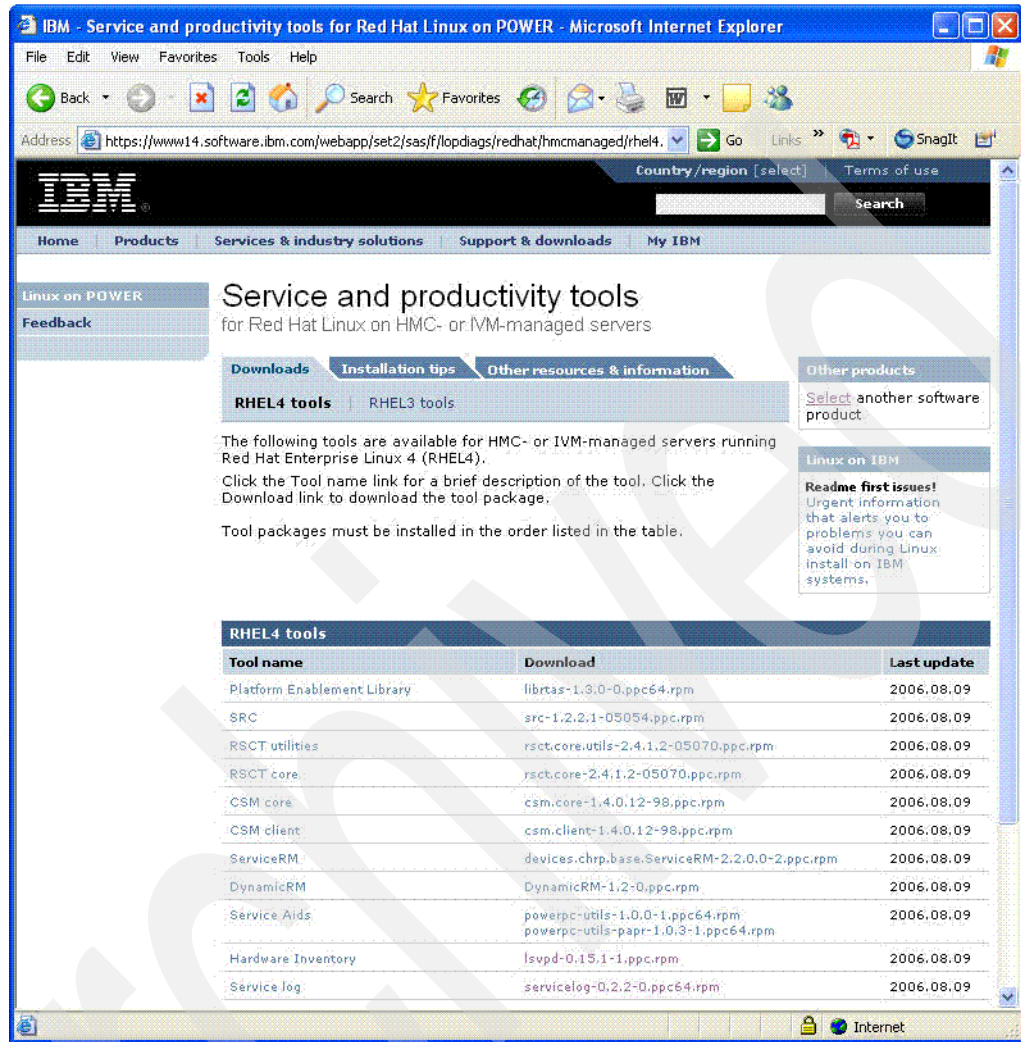


Figure 6-70 Service and productivity tools - package list

For support of the power-off function, download and install the following packages in the order listed:

- ▶ Platform Enablement Library                      librtas
- ▶ Service Log                                              servicelog
- ▶ Hardware Inventory                                  lsvpd
- ▶ Error Log Analysis                                    diagela

Once the packages have been downloaded to the Linux partition, you can install them with the following **rpm** command:

```
rpm -ivh <package name>
```

You can verify that the functions are providing support for the power-off function by a vary off of the Network Server. You should observe shutdown messages in the Linux virtual console and the Linux partition should go into a "Not Activated" state.



## 6.3 Native disk install - installation notes

Typically, most installations that are using native disk for Linux will leverage virtual I/O for access to the i5/OS CD device for the installation source. In this environment, a virtual SCSI client adapter will be created in the Linux partition that connects to a virtual SCSI server adapter in the i5/OS partition. For the installation process, the Linux partition will be configured with a power-controlling partition of the i5/OS partition and the Network Server will be configured with Power Control set to \*YES.

To start the installation, the Network Server will be varied on (in i5/OS).

Once the installation is completed, perform the following actions:

- ▶ Vary off the Network Server.
- ▶ Shutdown the Linux partition.
- ▶ Modify the Linux partition to remove the Power Controlling Partition.
- ▶ Modify the Network Server to set Power Control to \*NO.
- ▶ Start the Linux partition from the HMC, override the boot mode and set it to SMS.
- ▶ Use the SMS menus to set the boot order to boot from the native disk.
- ▶ Shutdown the partition.
- ▶ Restart the partition from the HMC and verify that it boots from the physical disks.

For the process, perform the following steps:

Depending on the Linux distribution being installed, the installer will attempt to reboot the partition in the middle of the installation process. When this occurs the partition will need to be shutdown and the boot modes and network server will need to be modified.

**Note:** As an example, for the Novell / SUSE SLES 10 installation, after all of the packages have been installed from the first CD, the partition will be restarted. This is the point where the partition will need to be shutdown and the boot modes modified.

1. When the installation is restarted (after packages from the first CD have been read), the Language screen will be displayed (Figure 6-71) press **Tab** to highlight Next and press **Enter**.

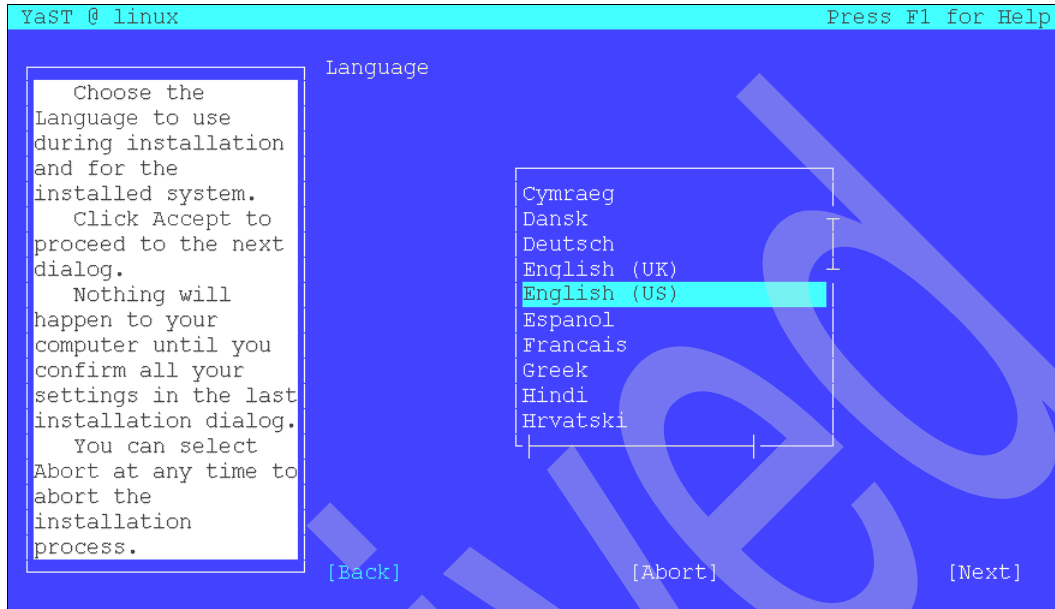


Figure 6-71 SLES 10 installer - language selection screen

2. When the License Agreement screen is displayed (Figure 6-72), select **No** and then press **Tab** to highlight Next and press **Enter**.

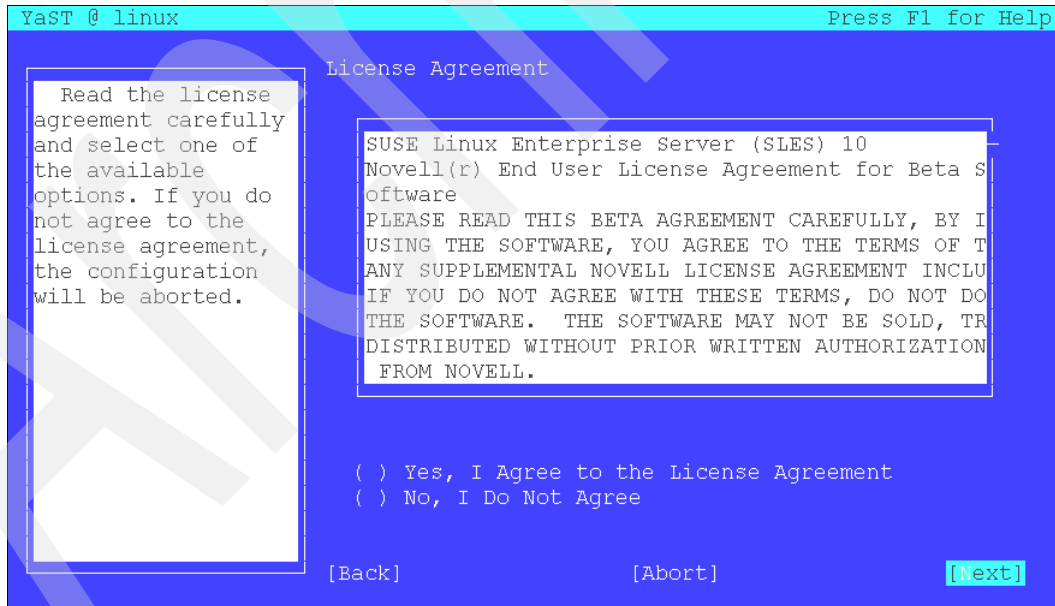


Figure 6-72 SLES 10 installer - license agreement

3. Because No was selected, a prompt to confirm the selection (Figure 6-73) is displayed. Confirm the declining of the License agreement and press **Enter**.

```
Reread or reconsider the license agreement?
[yes] [No]
```

Figure 6-73 SLES 10 installer - license agreement confirmation screen

4. When the Main Menu is displayed (Figure 6-74), type 8 (Power off) and press **Enter**.

```
Main Menu

1) Settings
2) System Information
3) Kernel Modules (Hardware Drivers)
4) Start Installation or System
5) Verify Installation CD-ROM/DVD
6) Eject CD
7) Exit or Reboot
8) Power off

>
```

Figure 6-74 SLES 10 installer - main menu

5. At the prompt to confirm the power off request (Figure 6-75), type 1 (Yes) and press **Enter**.

```
Do you want to halt the system now?

1) Yes
2) No

>
```

Figure 6-75 SLES 10 installer - power off request confirmation

6. Now that the operating system has been shutdown, the network server can be varied off. In i5/OS, enter the WRKCFGSTS \*NWS command.

- The Work with Configuration Status screen will be displayed (Figure 6-76). Type 2 (Vary off) in the opt field for the network server and press **Enter**.

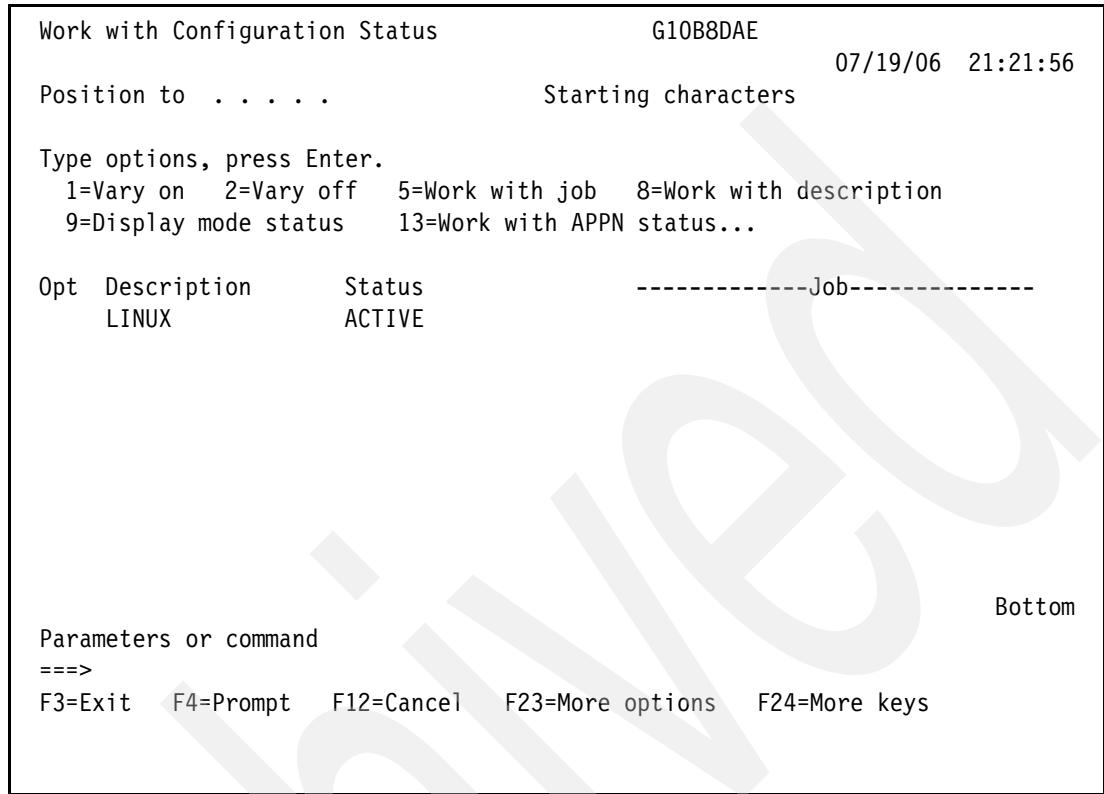


Figure 6-76 Work with configuration status

- Still on the Work with Configuration Status screen, type 8 (Work with description) in the Opt field of the Network Server and press **Enter**.

9. On the Work with Network Server Descriptions screen (Figure 6-77), type 2 (Change) in the Opt field of the Network Server and press **Enter**.

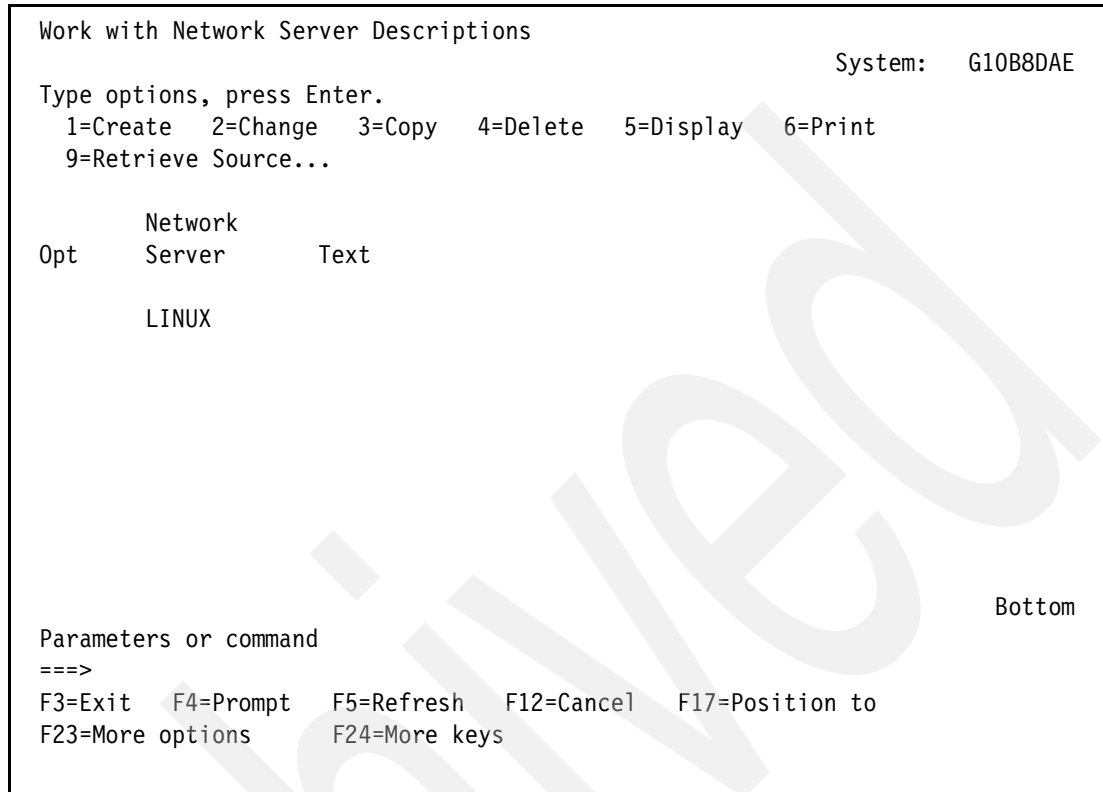


Figure 6-77 Work with network server descriptions

10. When the Network Server parameters are displayed, press **PgDn** *thrice* to display the last page of parameters. Change the value of the Power control field to \*NO and press **Enter**.
11. Press **F3** to exit the Work with Network Server Descriptions screen.
12. Back on the Work with Configuration Status screen, type 1 (Vary on) in the Opt field of the Network Server and press **Enter**. The Network Server is varied back on to make the i5/OS CD drive available for the remainder of the installation process.

## 6.4 Updating SLES 9 to SLES 10

The SLES 10 version of the Novell / SUSE Linux distribution includes support for upgrading from the previous version (SLES 9). This section describes the process of updating an instance of SLES 9 to SLES 10.

**Note:** Similar to any other update, a full backup of the SLES 9 installation should be made prior to attempting the update process.

1. The first step is to vary off the Network Server. From the i5/OS hosting partition issue the following command as shown in Example 6-8.

*Example 6-8 Vary off network server*

---

```
VRYCFG CFGOBJ(SLES9) CFGTYPE(*NWS) STATUS(*OFF)
```

---

The CFGOBJ value is the name of the Network Server.

The update process is started by IPLing the Linux partition off of the SLES10 distribution media. The network server will need to be modified to boot off of the SLES 10 installation kernel.

2. Place CD1 of the distribution media into the CD drive of the hosting i5/OS partition and issue the command shown in Example 6-9.

*Example 6-9 Modify network server description*

---

```
CHGNWSD NWSD(SLES9) IPLSRC(*STMF) IPLSTMF('/QOPT/SLES100.001/SUSEBOOT/INST64')
```

---

The NWSD value is the name of the Network Server Description.

**Note:** At the time of writing of this book, the path for the installation kernel is /qopt/SLES100.001/SUSEBOOT/INST64. The path should be validated through the documentation provided with the Linux distribution.

3. Once the Network Server Description has been modified, vary on the Network Server with the following command shown in Example 6-10.

*Example 6-10 Vary on network server*

---

```
VRYCFG CFGOBJ(SLES9) CFGTYPE(*NWS) STATUS(*ON)
```

---

**Note:** Make sure to have access to the virtual console for the Linux partition to be able to respond to the installation prompts.

The first screen to be displayed is the language selection screen shown in Figure 6-78.

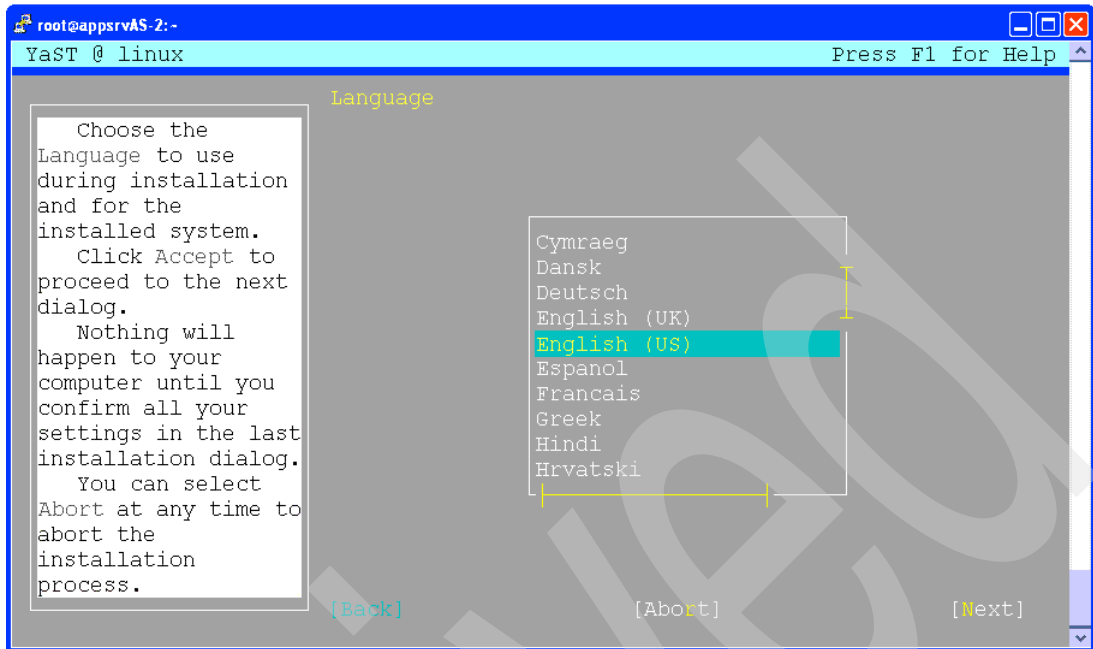


Figure 6-78 SLES 10 Update - Language Selection

4. Choose the desired language then press **Tab** until [Next] is highlighted and press **Enter**.

The next screen to be displayed is the License Agreement screen as shown in Figure 6-79.

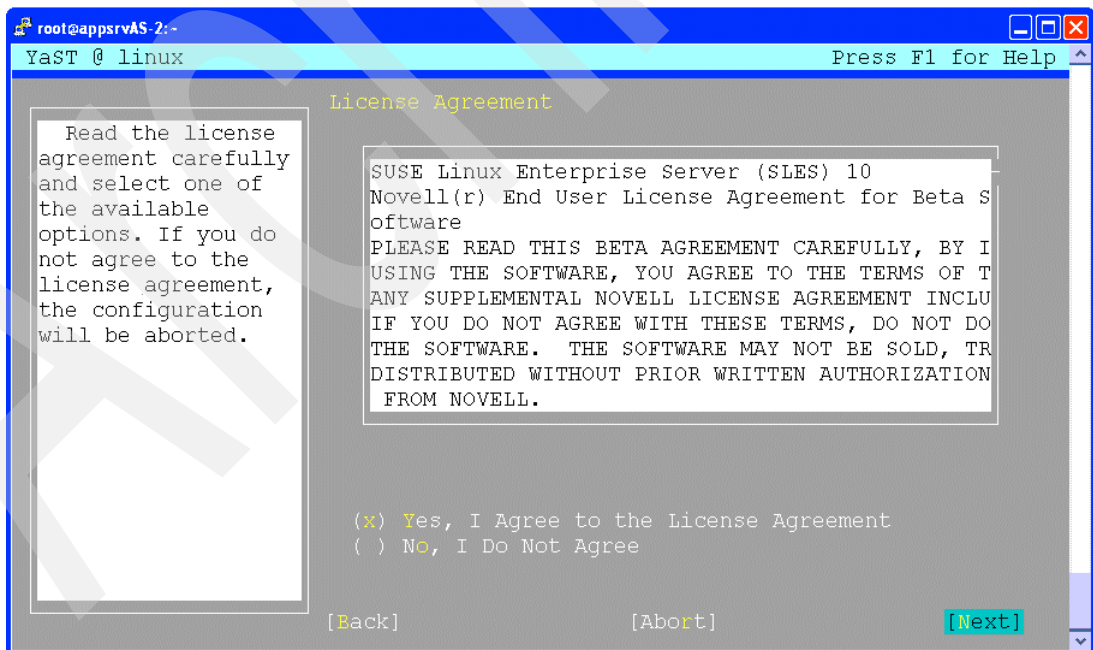


Figure 6-79 SLES 10 Update - License Agreement

5. Press **Tab** until [Next] is highlighted and then press **Enter**.

At this point the installer analyzes the attached disk and determines that an installation is already on the disk. The installation Mode screen (Figure 6-80) provides the ability to select a **New Installation** or an **Update**.

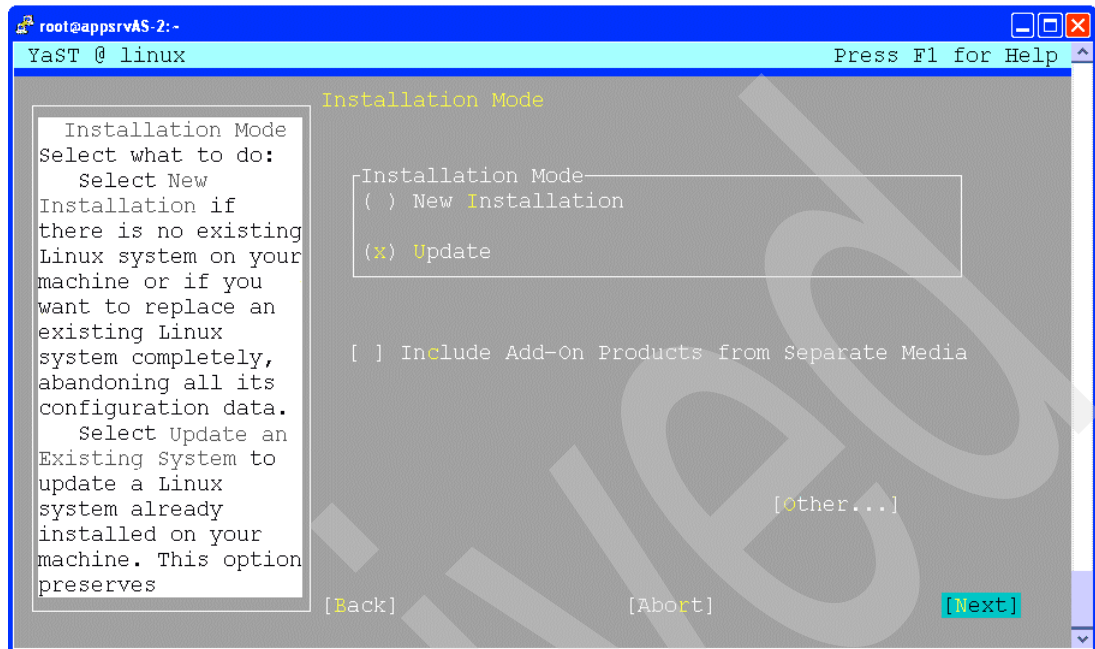


Figure 6-80 SLES 10 Update - Installation Mode

6. Press **Tab** until Update is highlighted and then press **Space Bar** to select. Then press **Tab** until [Next] is highlighted and then press **Enter**.
7. From the list of available installations (Figure 6-81) ensure that the installation to be updated is at the top of the list (it should in fact be the only distribution listed). Press **Tab** until [Next] is highlighted and then press **Enter**.

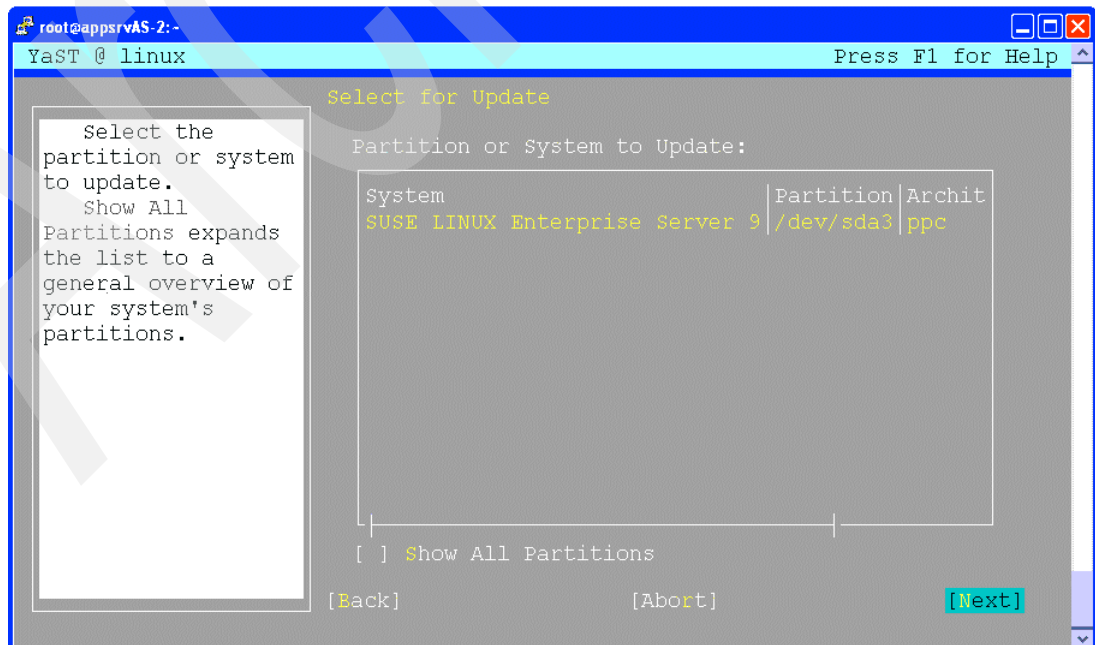


Figure 6-81 SLES 10 Update - Select installation to update



At this point the installer analyzes the installed packages and determines what packages need to be updated. Once that is complete the “Installation Settings” screen shown in Figure 6-82 is displayed.

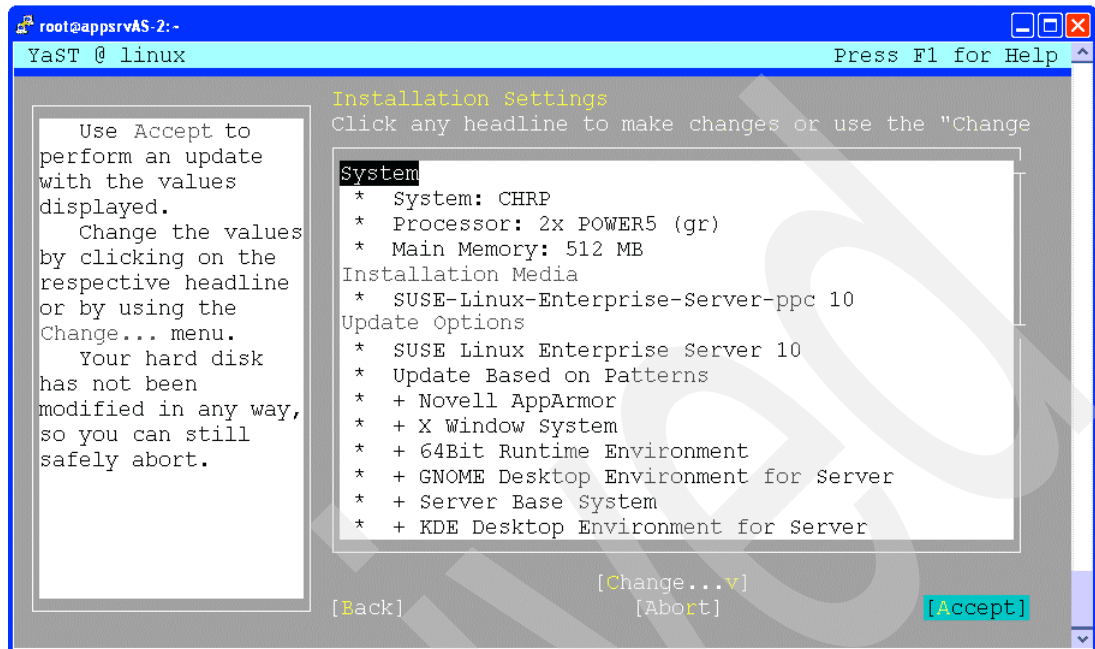


Figure 6-82 SLES 10 Update - Installation Settings

8. Press **Tab** on the Installation Settings screen until **[Accept]** is highlighted and then press **Enter**.
9. On the “Confirm Update” screen (Figure 6-83), press **Tab** until **[Start Update]** is highlighted and then press **Enter**.



Figure 6-83 SLES 10 Update - Confirm Update

At this point the installer copies and installs the updated packages onto the installed system. A constantly updated status of the update process is displayed as shown in Figure 6-84.

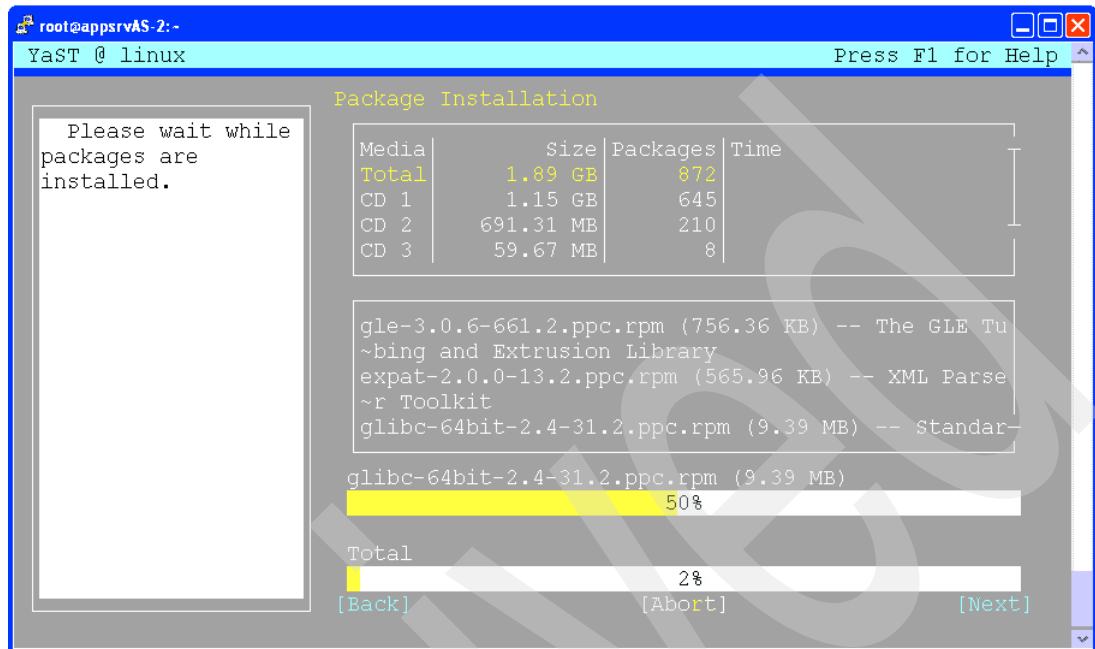


Figure 6-84 SLES 10 Update - Package Installation progress

After the updated packages have been installed, the installation program automatically reboots the system. Since the Network Server is still configured to start the installation kernel, the installation program attempts to restart. At this point, you should shutdown the Network Server so that the Network Server Description can be updated to boot off of the Linux kernel installed/updated in the virtual storage space.

10. When the Language selection screen (Figure 6-85) is displayed, press **Tab** until [Abort] is highlighted and then press **Enter**.

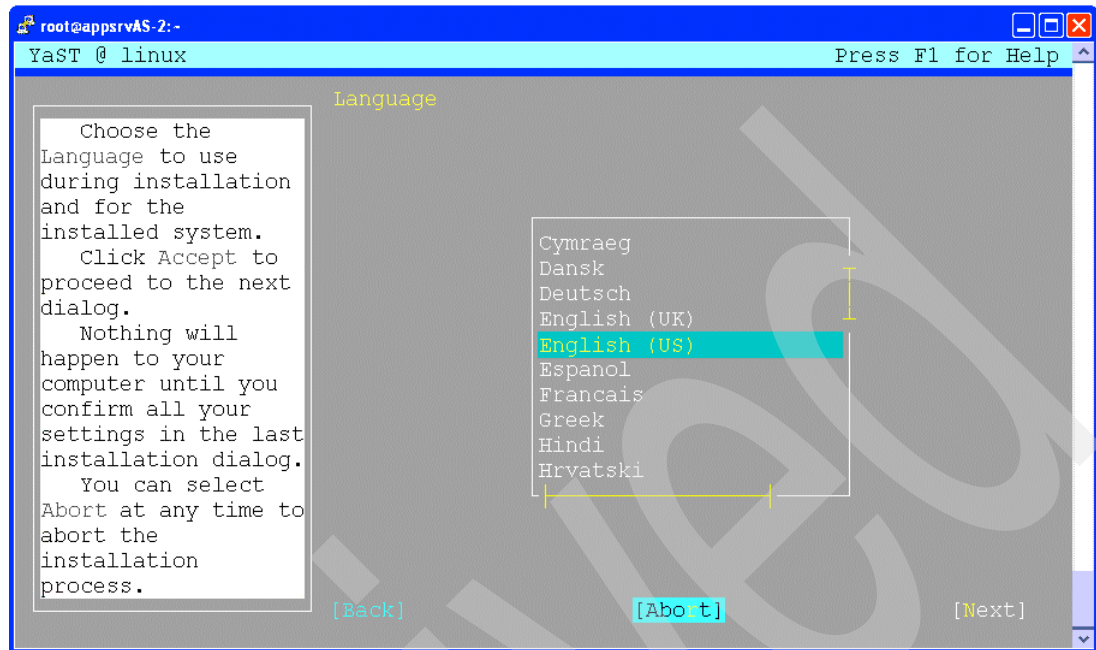


Figure 6-85 SLES 10 Update - Language selection - Abort Installation

11. When the confirmation dialog box (Figure 6-86) is displayed, press **Tab** until [Abort Installation] is highlighted and then press **Enter**.




Figure 6-86 SLES 10 Update - Confirm abort request

At this point the installation program terminates and the main menu is displayed as shown in Figure 6-87.

```
Main Menu

1) Settings
2) System Information
3) Kernel Modules (Hardware Drivers)
4) Start Installation or System
5) Verify Installation CD-ROM/DVD
6) Eject CD
7) Exit or Reboot
8) Power off

>
```

Figure 6-87 SLES 10 Update - Main Menu

12.Type 8 (Power off) and then press **Enter**.

13.When the confirm power off message (Figure 6-88) is displayed, type 1 (Yes) and press **Enter**.

```
Do you want to halt the system now?

1) Yes
2) No

>
```

Figure 6-88 SLES 10 Update - Confirm power off request

At this point the operating system will be shut down.

14.In the hosting i5/OS partition, issue the following command to vary off the Network Server as shown in Example 6-11.

*Example 6-11 Vary off Network Server*

---

```
VRYCFG CFGOBJ(SLES9) CFGTYPE(*NWS) STATUS(*OFF)
```

---

The value for CFGOBJ is the name of the Network Server.

15.In the hosting i5/OS partition, issue the following command to update the Network Server Description to boot off the installed kernel as shown in Example 6-12.

*Example 6-12 Update Network Server Description*

---

```
CHGNWSD NWSD(SLES9) IPLSRC(*NWSSTG) IPLSTMF(*NONE)
```

---

The value for NWSD is the name of the Network Server Description.

16. In the hosting i5/OS partition, issue the following command to vary the Network Server back on as shown in Example 6-13.

*Example 6-13 Vary on network server*

```
VRYCFG CFGOBJ(SLES9) CFGTYPE(*NWS) STATUS(*ON)
```

The value for CFGOBJ is the name of the Network Server.

**Note:** You should have access to the virtual console for the Linux partition. If the virtual console window has been terminated, restart it either through the HMC or the virtual console support of the hosting i5/OS partition.

Once the partition is restarted, the installer will prompt for any additional CDs required as shown in Figure 6-89.



Figure 6-89 SLES 10 Update - Prompt for additional CDs

17. Insert the requested CD into the CD drive of the hosting i5/OS partition and then press **Enter**.

Once all of the updated packages have been installed a prompt is displayed for downloading additional updates from the Internet as shown in Figure 6-90.

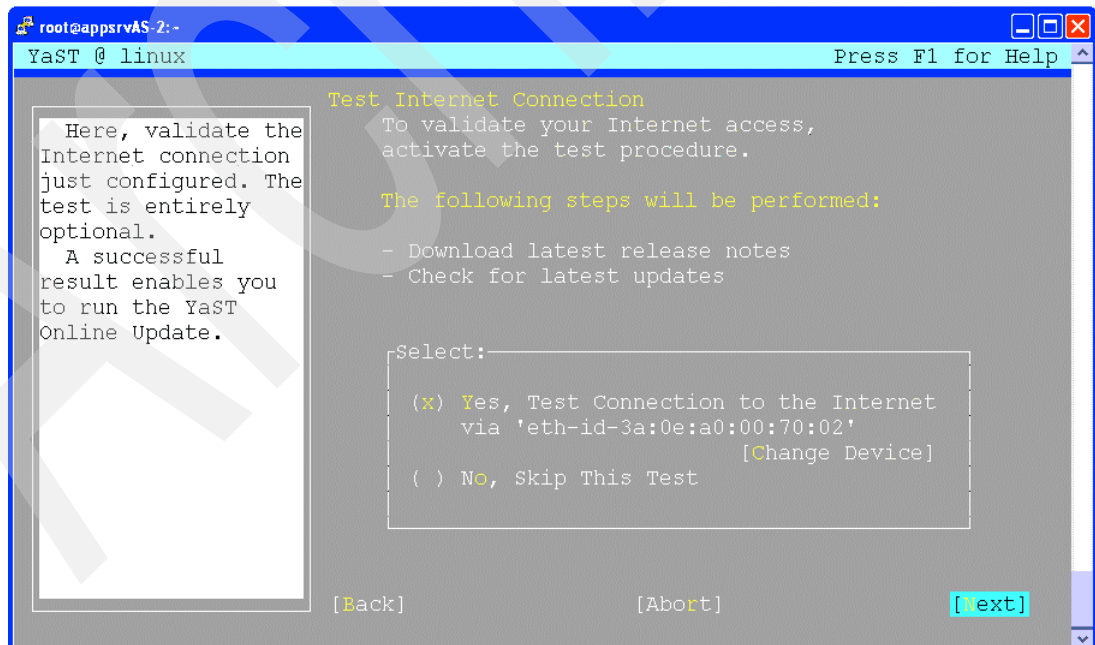


Figure 6-90 SLES 10 Update - Download additional updates

The selection you make will be dependent on whether your environment supports a connection to the Internet for the Linux partition.

18. For this example, we will assume that a connection to the Internet is not available, therefore press **Tab** until **No. Skip This Test** is highlighted then press **Space Bar** to make the selection. Press **Tab** until **[Next]** is highlighted and press **Enter**.

At this point, the system configuration files are updated by the installer. Once the configuration files have been updated, the release notes are displayed.

19. After review of the release notes, press **Tab** until **[Next]** is highlighted and then press **Enter**.

An Installation Completed dialog box is displayed as shown in Figure 6-91.

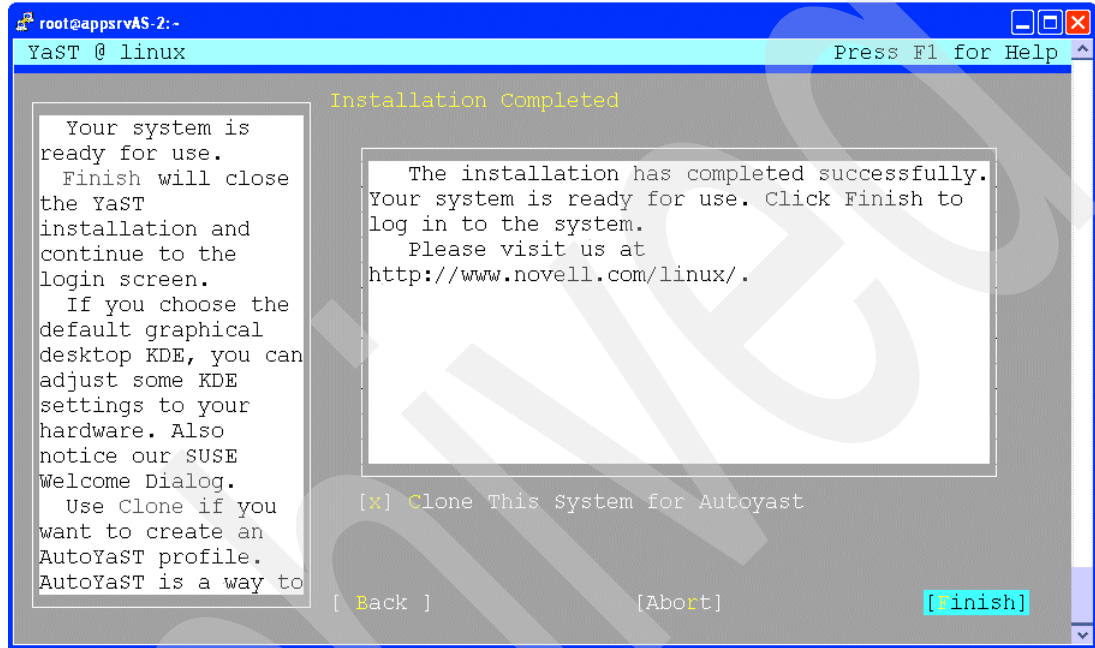


Figure 6-91 SLES 10 Update - Installation Completed

20. Press **Enter** to select the **[Finish]** option.

At this point, the updated Linux instance will complete starting. One way to confirm the update is to notice that the Logon herald now indicates that the distribution is SLES 10 as shown in Figure 6-92.




Figure 6-92 SLES 10 Update - Updated Logon Herald

# Administering the Linux environment

This chapter discusses the administration and operation topics of the Linux server including:

- ▶ Checking the Linux partition resources
  - From the Hardware Management Console (HMC)
  - From Linux
- ▶ Webmin discussion
  - Common administration task examples using WebMin
- ▶ Backup and restore of Linux partition
  - Backup and restore strategy
  - i5/OS side backup and restore
  - Linux side backup and restore
  - HMC configuration backup and restore

## 7.1 Checking Linux partition resources

We need to know the current resource utilization of the Linux partitions for proper management of the resources. Those resources can be processors, memory utilization, process jobs. This section explains how to check the Linux resources from HMC and from Linux OS.

### 7.1.1 From HMC

On runtime you can monitor your Linux partitions from HMC. Here you can see the processors and memory usage and the I/O devices that are attached to the partition. Perform the following steps:

1. On HMC right-click over partition and select **Properties** as illustrated in Figure 7-1.

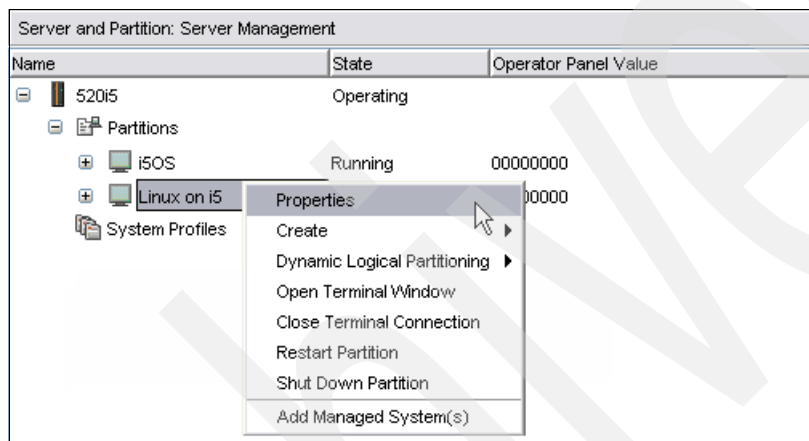


Figure 7-1 HMC - partition properties

2. Over the partition properties window, select the **Hardware** tab. In this window, select the **I/O** tab as illustrated in Figure 7-2. Under the I/O tab, you can see the resources that are used for the logical partition. In our case, the Linux partition has attached an ethernet adapter as Physical I/O.



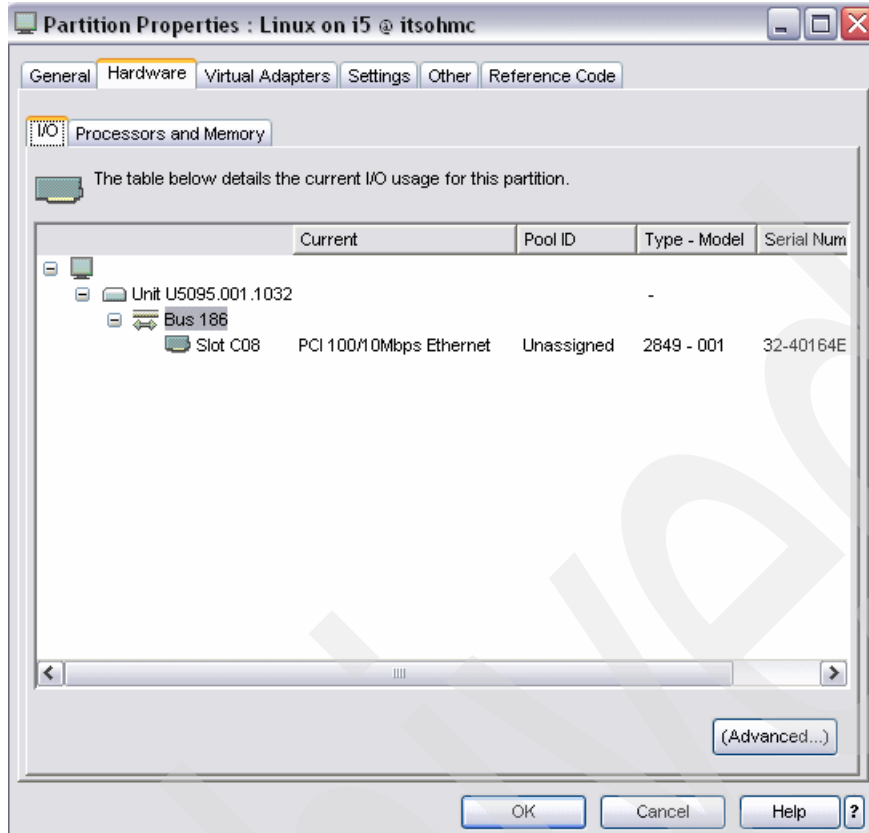


Figure 7-2 Partition properties - Hardware I/O

3. Select the **Processors and Memory** tab as illustrated in Figure 7-4. This tab shows the resource definition as well as the current resource allocation.

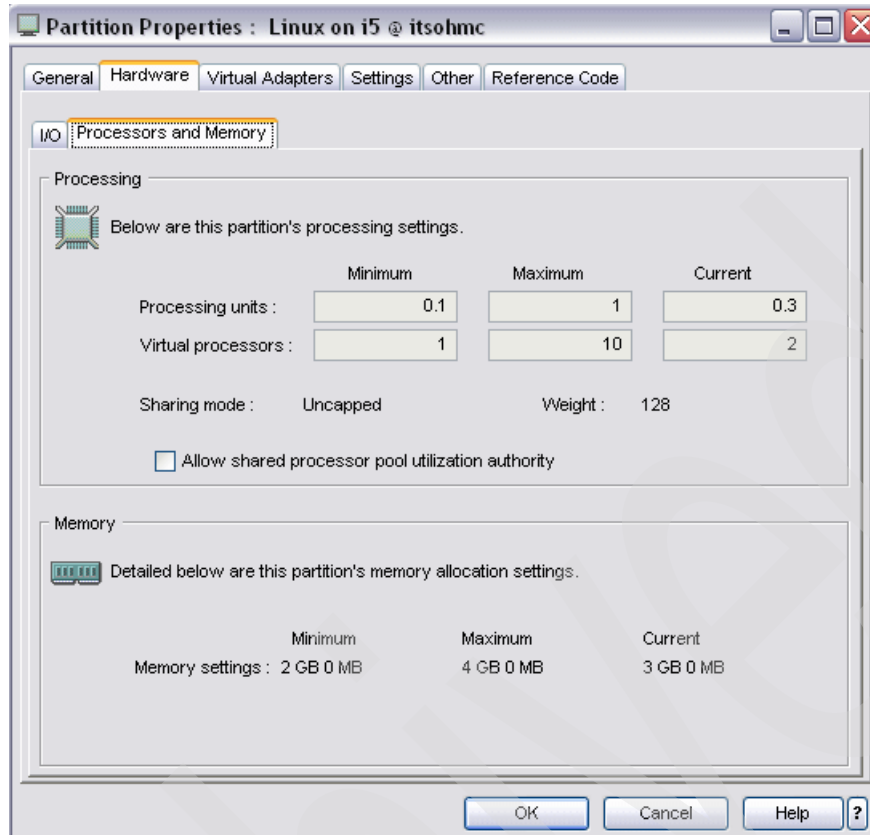


Figure 7-3 Partition properties - Hardware processors and memory

## 7.1.2 From Linux OS

There are a number of tools that exist in Linux to determine the current resource allocation or utilization. This section briefly covers the following commands:

- ▶ top
- ▶ vmstat
- ▶ free
- ▶ iostat

### The top command

The **top** command provides a dynamic real-time view of a running system. It can display system summary information as well as a list of tasks currently being managed by the Linux kernel. The types of system summary information shown and the types, order and size of information displayed for tasks are all user configurable and that configuration can be made persistent across restarts. The program provides a limited interactive interface for process manipulation as well as a much more extensive interface for personal configuration.

When operating the **top** command, press **h** or **?** for help and **q** to quit. Alternatively, you could use the traditional interrupt key, **Ctrl+C** when you are done.

To start **top**, run the following command in any Linux emulation:

```
linux:~ # top
```

As illustrated in Figure 7-4, the **top** command is composed for the following screen areas:

1. Summary area: The summary area gives a brief status of the system showing the quantity of tasks, the CPU, memory and swap utilization.
2. Message/prompt line: This is for the command entry. For more information about the commands, press the **h** or **?** keys for help.
3. Columns header: This file defines the column header. In our example, you see processes by PID (Process ID), user, PR (priority), NI (nice value, that is another type of priority), VIRT (Virtual memory used in the task), RES (non swapped physical memory used), SHR (Shared Memory size), S (Process status), %CPU (CPU usage), %MEM (Memory usage), TIME+ (CPU time), COMMAND (Command line or program name).
4. Task area: In this area, every line or file is one independent process.

```

top - 08:51:36 up 13:02, 2 users, load average: 0.00, 0.00, 0.00
Tasks: 95 total, 1 running, 94 sleeping, 0 stopped, 0 zombie
Cpu(s): 0.0% us, 0.0% sy, 0.0% ni, 100.0% id, 0.0% wa, 0.0% hi, 0.0% si
1
Mem: 3076376k total, 259024k used, 2817352k free, 23408k buffers
Swap: 1541984k total, 0k used, 1541984k free, 111260k cached

-
2
 PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND
3
4015 root 16 0 2472 1508 2228 R 0.3 0.0 0:02.39 top
 1 root 16 0 628 284 480 S 0.0 0.0 0:00.20 init
 2 root RT 0 0 0 0 S 0.0 0.0 0:00.00 migration/0
 3 root 34 19 0 0 0 S 0.0 0.0 0:00.00 ksoftirqd/0
 4 root RT 0 0 0 0 S 0.0 0.0 0:00.00 migration/1
 5 root 34 19 0 0 0 S 0.0 0.0 0:00.00 ksoftirqd/1
 6 root RT 0 0 0 0 S 0.0 0.0 0:00.00 migration/2
 7 root 34 19 0 0 0 S 0.0 0.0 0:00.00 ksoftirqd/2
4
 8 root RT 0 0 0 0 S 0.0 0.0 0:00.00 migration/3
 9 root 34 19 0 0 0 S 0.0 0.0 0:00.00 ksoftirqd/3
 42 root 5 -10 0 0 0 S 0.0 0.0 0:00.00 events/0
 43 root 5 -10 0 0 0 S 0.0 0.0 0:00.00 events/1
 44 root 5 -10 0 0 0 S 0.0 0.0 0:00.00 events/2
 45 root 5 -10 0 0 0 S 0.0 0.0 0:00.00 events/3
 46 root 5 -10 0 0 0 S 0.0 0.0 0:00.00 kblockd/0
 47 root 5 -10 0 0 0 S 0.0 0.0 0:00.00 kblockd/1
 48 root 5 -10 0 0 0 S 0.0 0.0 0:00.00 kblockd/2

```

Figure 7-4 top command

### The vmstat command

The **vmstat** command reports information about processes, memory, paging, block IO, traps, and CPU activity. The first report produced gives averages since the last reboot. Additional reports give information about a sampling period of length delay. The process and memory reports are instantaneous in either case.

To start **vmstat**, run the Linux command as illustrated in Figure 7-5.

```
linux:~ # vmstat
procs -----memory----- --swap-- -----io----- --system-- -----cpu----
 r b swpd free buff cache si so bi bo in cs us sy id wa
 0 0 0 2800192 26920 129336 0 0 1 1 2 6 0 0 100
 0
linux:~ #
```

Figure 7-5 *vmstat* command

Execute **vmstat -h** or **man vmstat** for further information.

### The free command

The **free** command displays the total amount of free and used physical and swap memory in the system, as well as the buffers used by the kernel. The shared memory column should be ignored as it is obsolete.

To run the **free** command, execute it on Linux as illustrated in Figure 7-6.

```
linux:~ # free
 total used free shared buffers cached
Mem: 3076376 282088 2794288 0 27184 130100
-/+ buffers/cache: 124804 2951572
Swap: 1541984 0 1541984
linux:~ #
```

Figure 7-6 *free* command

Execute **free -h** or **man free** for further information.

### The iostat command

The **iostat** command is used for monitoring system input/output device loading by observing the time the devices are active in relation to their average transfer rates. The **iostat** command generates reports that can be used to change system configuration to better balance the input/output load between physical disks. The first report generated by the **iostat** command provides statistics about the time since the system was booted.

**Note:** To run this command, it is necessary to install the **sysstat** package. Refer to your distribution software installer to obtain this package.

In Figure 7-7, we illustrate how to execute the `iostat` on a Linux terminal. This example shows two system reports at intervals of two seconds. Each new report contains statistics collected during the interval since the previous report.

```
linux:~ # iostat 2 2
Linux 2.6.5-7.97-pseries64 (linux) 09/08/04

avg-cpu: %user %nice %sys %iowait %idle
 0.06 0.01 0.02 0.02 99.90

Device: tps Blk_read/s Blk_wrtn/s Blk_read Blk_wrtn
sda 0.58 10.03 5.76 515518 295872
sdb 0.00 0.03 0.00 1288 0
sr0 0.01 0.15 0.00 7804 0

avg-cpu: %user %nice %sys %iowait %idle
 0.00 0.00 0.12 0.00 99.88

Device: tps Blk_read/s Blk_wrtn/s Blk_read Blk_wrtn
sda 0.00 0.00 0.00 0 0
sdb 0.00 0.00 0.00 0 0
sr0 0.00 0.00 0.00 0 0

linux:~ #
```

Figure 7-7 `iostat` command

## 7.2 Linux utilities

The following software packages are the most commonly used to connect to and manage your Linux servers.

## 7.2.1 YaST for SUSE

YaST stands for *Yet Another System Tool*. It is used to configure and manage your Linux server. You can start YaST by typing `~> yast` as system administrator (root) on the Linux command line. YaST Control Center window appears as illustrated in Figure 7-8.

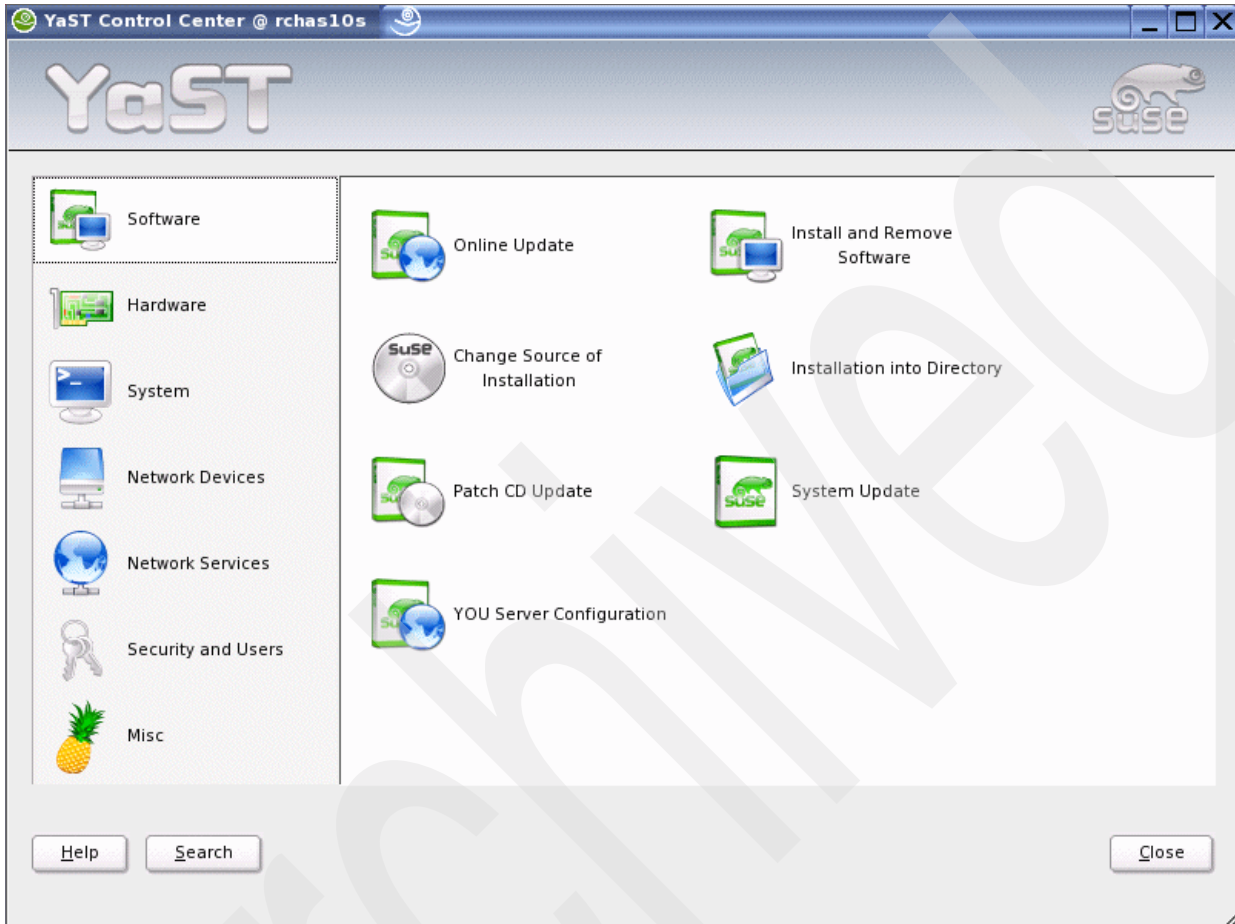


Figure 7-8 YaST Control Center

This tool is available on the SUSE Linux distributions.

## 7.2.2 Red Hat Setup

Setup, shown in Figure 7-9, is the Red Hat implementation of the Linux configuration tool. It is used to configure and manage your Linux server. As system administrator (root) on the Linux command line, you can start Setup by typing `~> setup`.

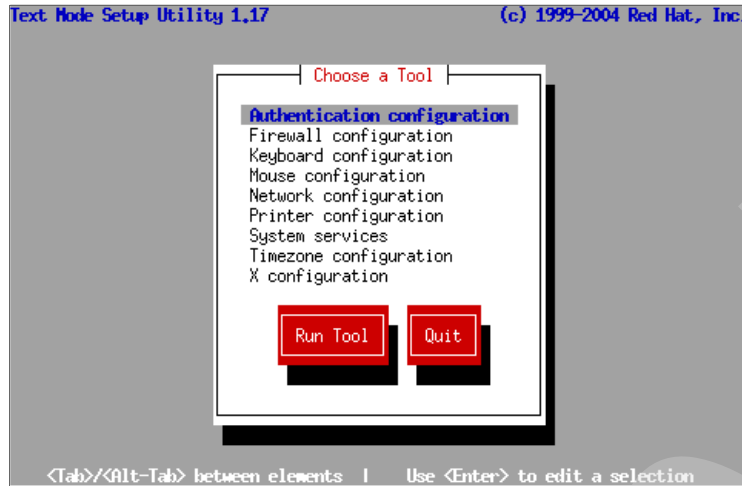


Figure 7-9 Red Hat Setup

This tool is available on the Red Hat Linux distributions.

## 7.2.3 PuTTY

PuTTY, shown in Figure 7-10, is a free Telnet and SSH client for Win32® and UNIX platforms. The PuTTY version used in this book is version 0.57.

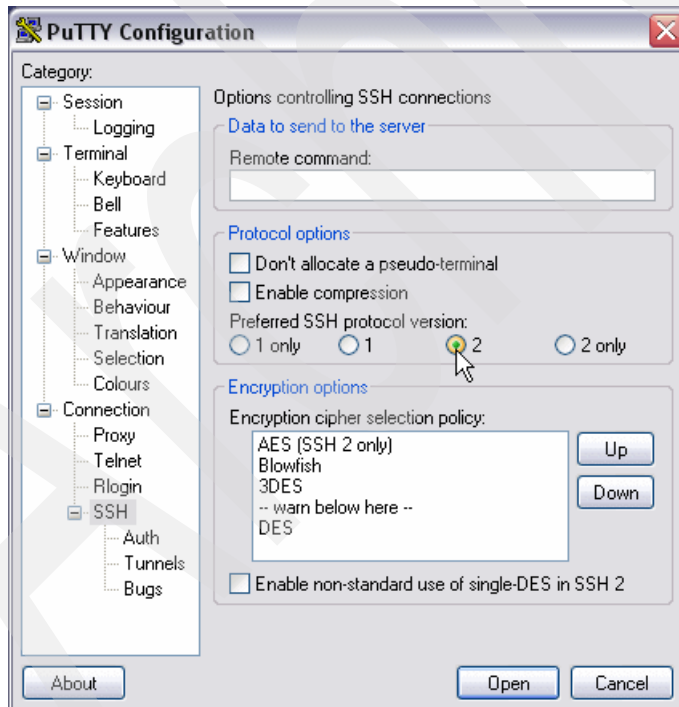


Figure 7-10 PuTTY

You can download PuTTY from the Web at:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/>

## 7.2.4 VNC

Virtual Network Computing (VNC), shown in Figure 7-11, is free software which allows you to view and interact with another computer which may or may not have the same operating system.




Figure 7-11 VNC

You can download VNC from the Web at:

<http://www.realvnc.com/>

## 7.2.5 Webmin

Webmin, shown in Figure 7-12, is a Web-based interface for system administration of UNIX systems. This is an open-source tool that runs on a variety of Linux distributions including SUSE and Red Hat and a variety of hardware platforms. With Webmin, you can access your Linux server using a Web browser with the following address:

<https://linuxhostname:10000>

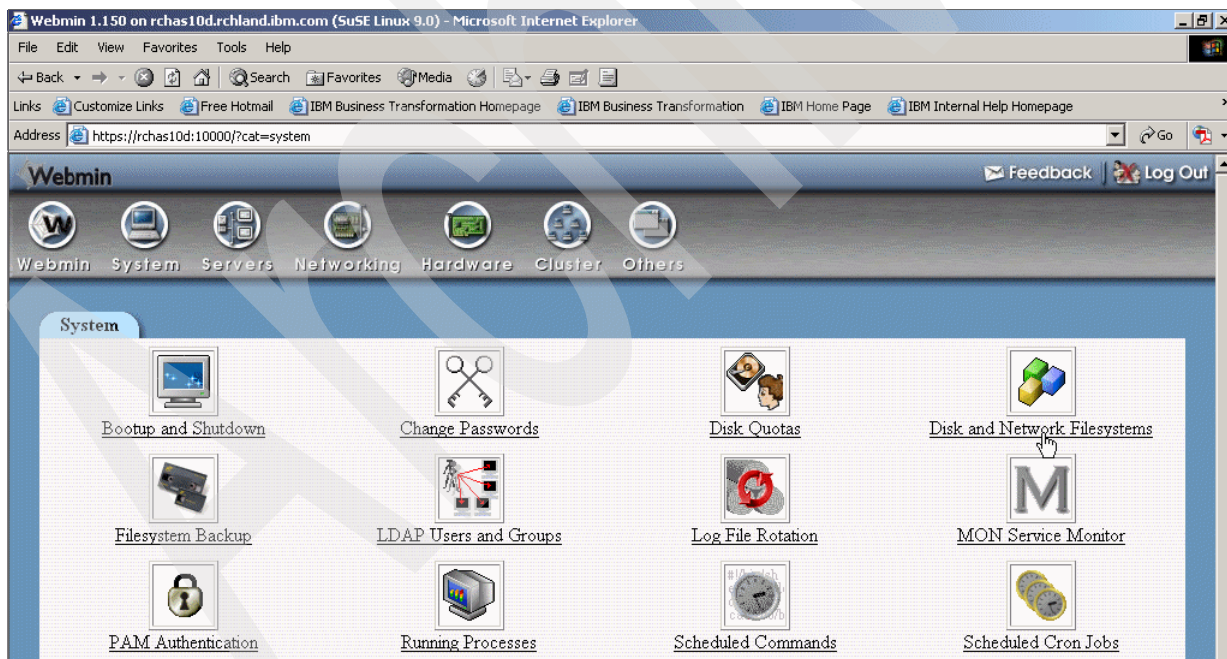


Figure 7-12 Webmin

You can download Webmin from the Web at:

<http://www.webmin.com/>



## Using Webmin for Linux administration tasks

Before you start any administration task, it is good idea to set up Webmin environment to be most convenient for you. For example, if English is not your native language, select **Webmin** → **Change Language and Theme** and change Webmin UI language.

### Viewing and editing files

This section explains how to use Webmin to find, view, and edit the files.

1. From the main Webmin menu, select **Others** → **File Manager** as illustrated in Figure 7-13.

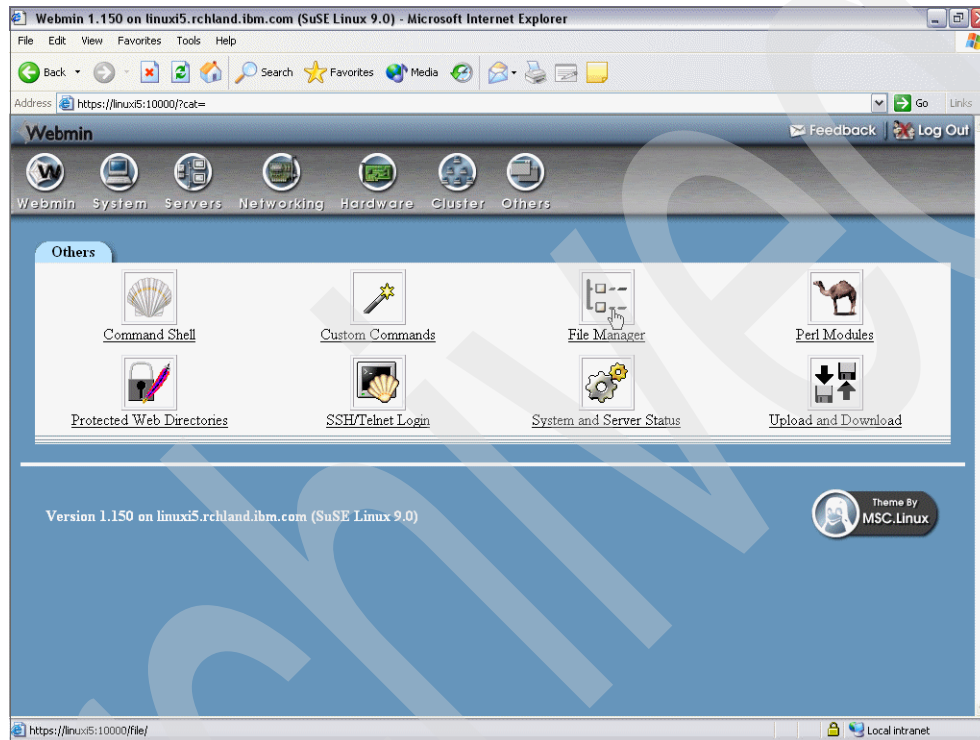


Figure 7-13 Webmin - File manager option

- You can click the directory or file of your interest or you can search for a file by clicking the Find icon in the navigation bar. Once you find the file you want, click the Edit icon for editing as illustrated in Figure 7-14.

You can double-click the file to view the contents, however, you cannot edit it.

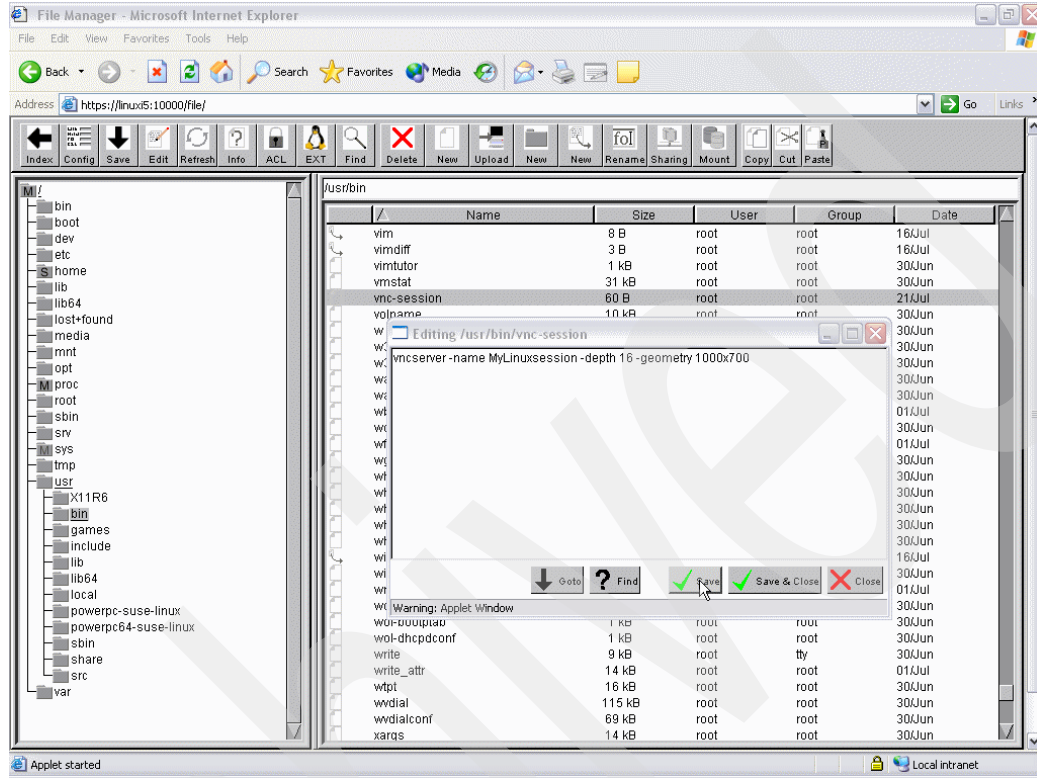


Figure 7-14 Webmin - File Manager

### Security settings

This section explains how to use Webmin's navigation bar icons as illustrated in Figure 7-15 for a specific file's security settings:

- ▶ **Info** icon: You can view and edit the permission and ownership of the file
- ▶ **ACL** icon: You can view and edit the Access Control List over the file
- ▶ **EXT** icon: You can view and edit the extra attributes of the file where some of the attributes are security-related.



Figure 7-15 Webmin - File Manager - navigation bar

Figure 7-16 illustrates the examples of these icons.

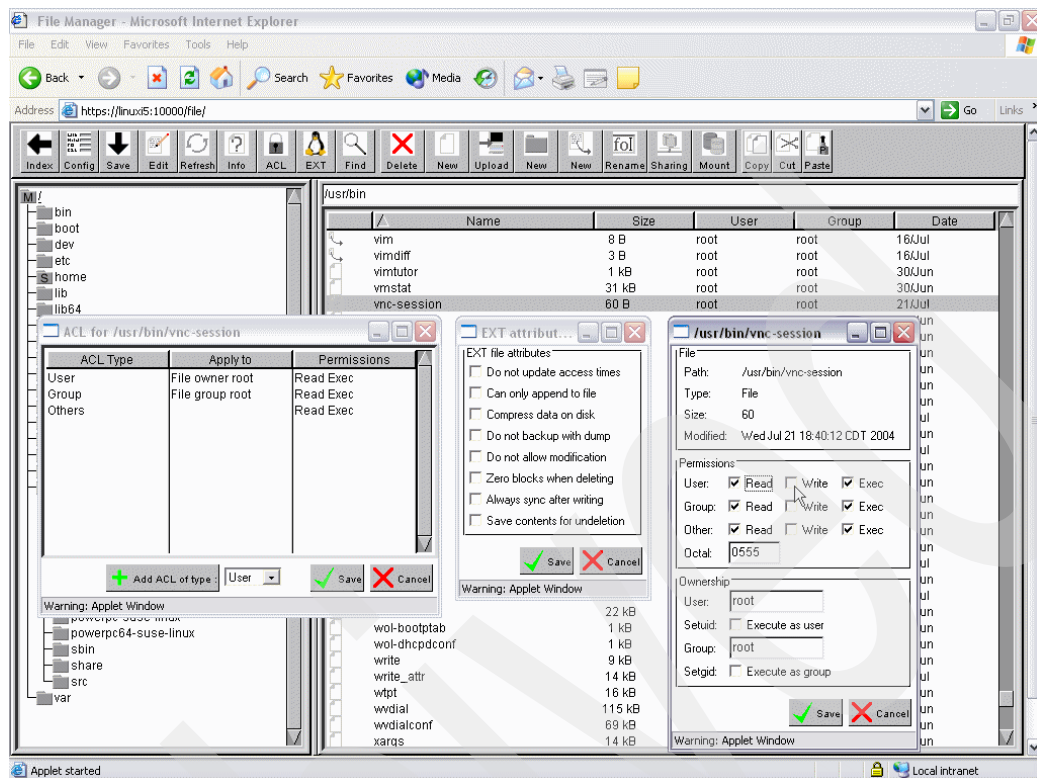


Figure 7-16 Webmin - File Manager - file privileges

## Managing users and groups

Click the Users and Groups icon in the System pane to administrate Linux users (Figure 7-17).

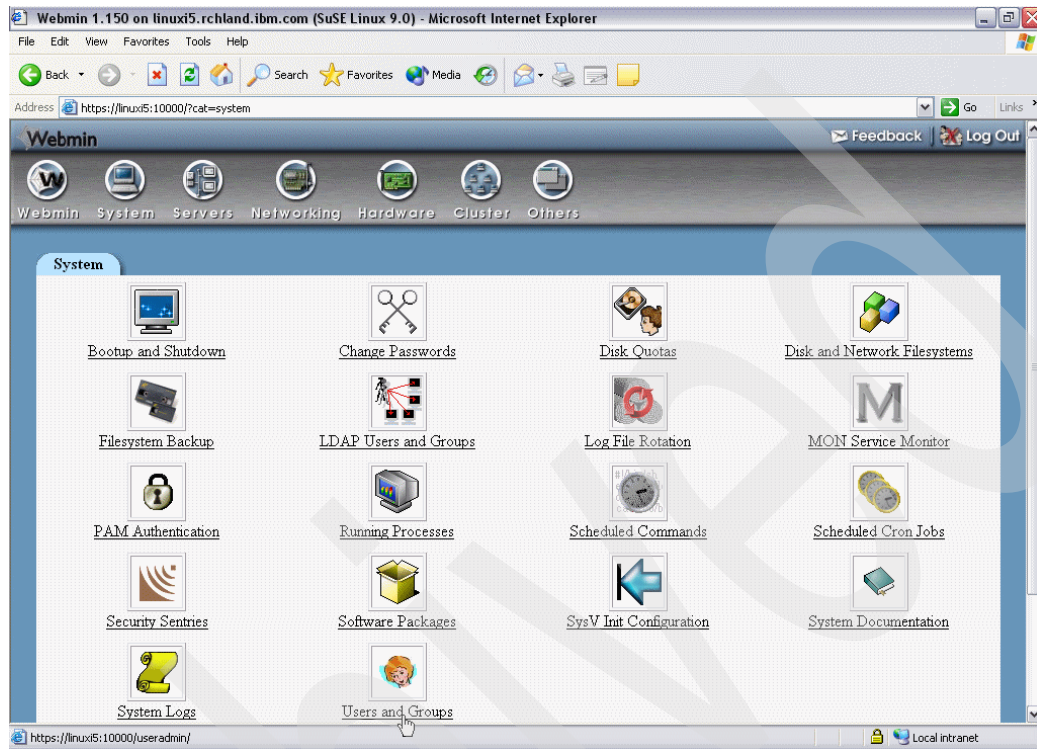


Figure 7-17 Webmin - Users and Groups

From the opened windows you can create, modify, or delete users or groups only by filling the appropriate fields or pressing buttons.

You can find details about using Webmin on the Web at:

<http://www.swelltech.com/support/pdfs/webminguide.pdf>

There are many other Webmin modules that are very useful which you can find on the Web at:

<http://webadminmodules.sourceforge.net/>

Alternatively, from the main Webmin page, you can click the Third-Party Modules icon. Install these modules by selecting **Webmin** → **Webmin Configuration** → **Webmin Modules** → **Install Module**.

## 7.2.6 WinSCP

WinSCP is an open source SFTP and SCP client for Windows using SSH to connect to UNIX or Linux machines.

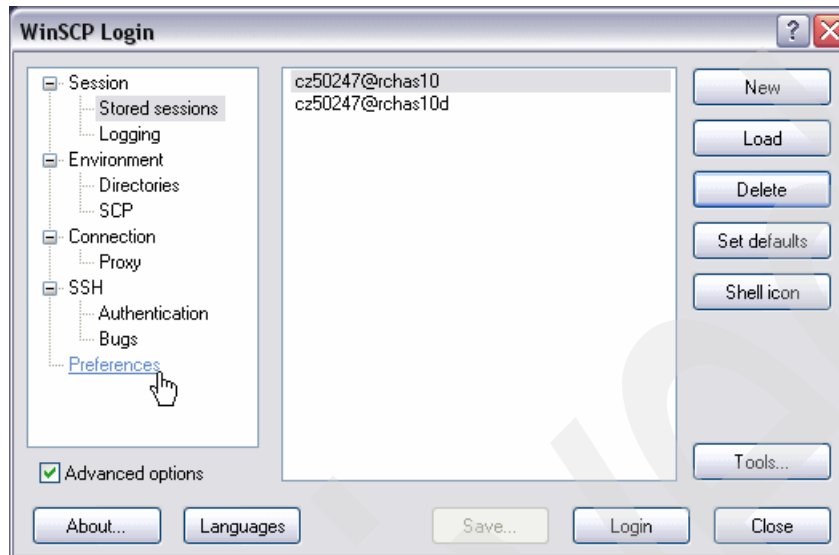


Figure 7-18 WinSCP preferences

You can download WinSCP from the Web at:

<http://winscp.sourceforge.net/eng/>

## 7.3 Backup and recovery

In this section, we discuss backup strategy for Linux partitions. This section also introduces some tools for backup.

### 7.3.1 Backup strategy for Linux partition

We assume you have virtual I/O for your Linux partition where disk storage resources are served by the i5/OS partition. Therefore, there are four components which should be part of your save/restore strategy:

- ▶ NWS Storage Space (NWSSTG): All Linux disks are parts of objects called Network Server Storage Space in i5/OS terminology.
- ▶ NWS Description (NWSD): All information about image of Linux server are described inside of objects called Network Server Description (NWSD) in i5/OS terminology.
- ▶ NWS Storage Space Link: Information about links among NWSD and NWS Storage Spaces are part both of NWS Storage Space and NWSD.
- ▶ Partition: All information about setting Linux partition are stored on HMC configuration.

## 7.3.2 i5/OS oriented backup

i5/OS gives you the following possibilities to save data for Linux partition:

- ▶ Save entire system (option 21 in i5/OS Save menu)
  - This option save all data from i5/OS including NWSD or NWS Storage Spaces or links among NWSD and NWS Storage Spaces.
  - Data saved this way can be used for recovery of an entire system and for recovery of a part of system like NWS Storage Space or NWSD.
  - It is a common practice that this is done once in a month.
- ▶ Save system data only (option 22 in i5/OS Save menu)
  - This option saves NWSSTG and NWSD.
  - It is a common practice that this is done once in a week.
- ▶ Save all user data (option 23 in i5/OS Save menu)
  - This option saves NWSSTG.

These three options are base backup options in an i5/OS strategy. They give us a weekly copy of two parts of Linux environment. This is probably enough for NWSD because it does not change frequently. But for the data in NWSSTG we need to implement the daily backup strategy.

There are two possibilities. First is to save it on NWSSTG level and second is to save on files or directories level.

We also assume that a base i5/OS strategy includes a daily backup of all new and changed document (option 31 in i5/OS Save menu).

### NWSSTG level backup

This is backing up the whole NWSSTG space, that is, whole disk in Linux term. This section explains the steps for this task.

All NWSSTG spaces are in /QFPNWSSTG directory in IFS of i5/OS. One of the i5/OS object types is Save file (\*SAVF). This section describes the method of backing up NWSSTG using this object type. It does compress data file while it saves it and is similar to the way how tar file in Linux works.

1. If the Linux server is still up and running, stop it first. On Linux terminal session, type:

```
shutdown -h now
```

2. Then you need to vary off NWSD. On i5/OS command entry, type:

```
VRYCFG CFGOBJ(LINUX) CFGTYPE(*NWS) STATUS(*OFF)
```

Give the name of corresponding NWSD to your Linux partition to CFGOBJ parameter. In our example, it is "LINUX".

3. Create a library where you will store NWS Storage Spaces on a daily basis by typing the following command:

```
CRTLIB LNXBACKUP
```

4. Create SAVF (with name NWSSSx where NWSSS is a name of NWS Storage Space and x is day in week) to which you save your Linux NWS Storage Space with the following command:

```
CRTSAVF LNXBACKUP/NWSDSSx
```

5. Repeat CRTSAVF step for all other NWS Storage Spaces you want to save.

6. Save your NWS Storage Spaces to this SAVF with data compression, assuming that LINUX is the name of NWSD, with the following command:  

```
SAV DEV('/QSYS.LIB/LNXBACKUP.LIB/NWSSSx.FILE') OBJ('/QFPNWSSTG/LINUX')
DTACPR(*YES)
```
7. Repeat SAV step for all other NWS Storage Spaces you want to save.
8. Vary on your NWSD.
9. The SAVF will be saved during daily save of new or changed documents.
10. The LNXBACKUP library can be cleaned after weekly save of all user data.

**Tip:** You can save NWSD in LNXBACKUP library as well as part of saved configuration data. Create SAVF and then use option 10 from Save menu on i5/OS.

### Files and directories level backup

This section describes the method of files and directories level backup.

**Note:** This type of backup has the advantage of not requiring the Linux partition to be varied off, that is, to be down.

This task consists of a few sub tasks:

- ▶ Creating a new directory in i5/OS IFS and export it to NFS
- ▶ Mounting this NFS directory to Linux file system
- ▶ Creating script in Linux for backup files and directories to this NFS drive
- ▶ Scheduling this operation

This IFS directory exported to NFS will be then saved in i5/OS on base daily backup of new and changed documents.

**Note:** We assume that the virtual ethernet line between the i5/OS and Linux partition already exists.

## Creating a new directory in i5/OS IFS and exporting it to NFS

1. Create a new directory in i5/OS IFS. On iSeries Navigator, select **File Systems** → **Integrated File System** → **Root**. Right-click the root and select **New Folder** as illustrated in Figure 7-19.

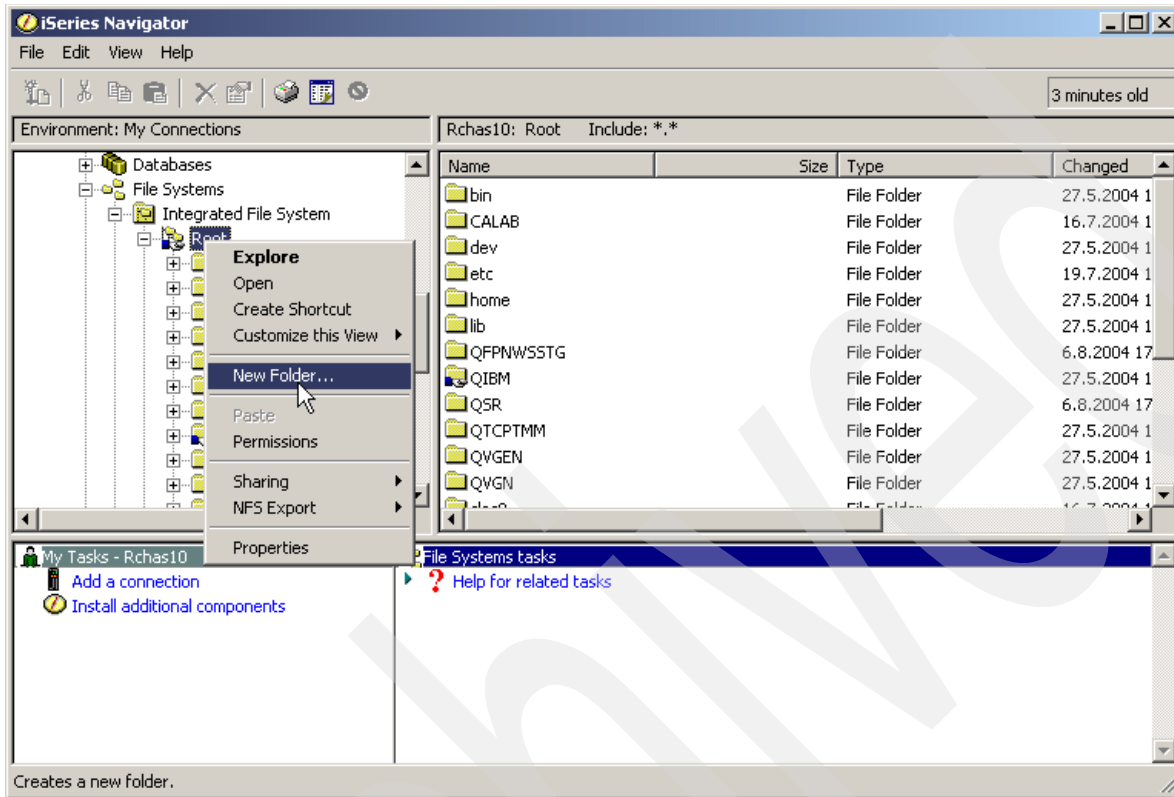


Figure 7-19 Create new folder

2. In the New Folder window, as illustrated in Figure 7-20, give a folder name and select **No** on the Scan objects created in folder value. Then click **OK**.

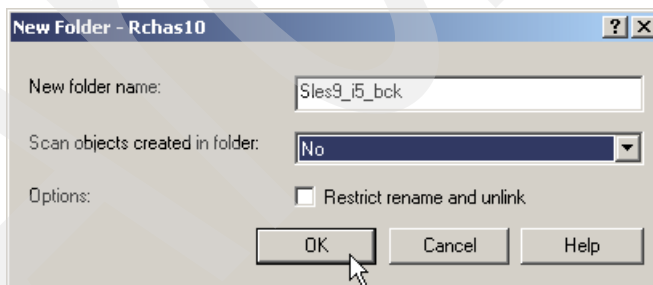


Figure 7-20 New folder parameters



3. On the iSeries, start the NFS server. Select **Network** → **Servers** → **TCP/IP**. Right-click **NFS** and then select **Start All**.

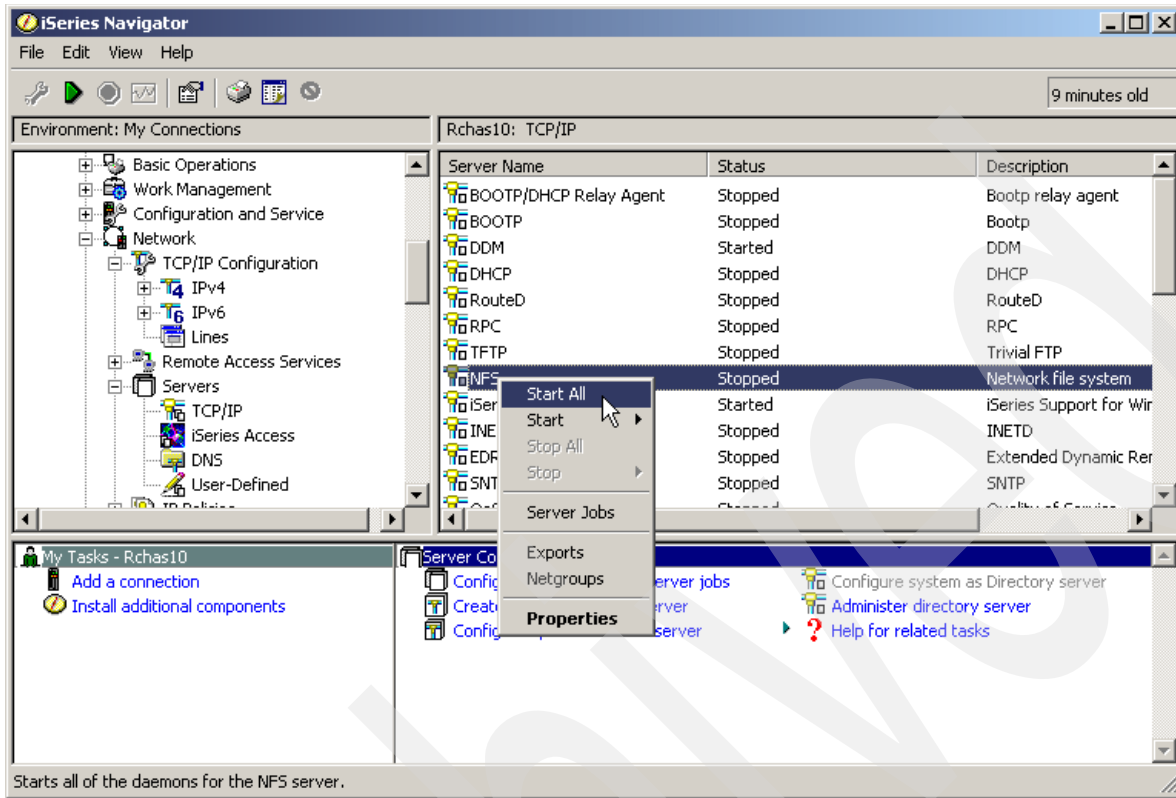


Figure 7-21 NFS - select all

- Right-click the **NFS** again and select **Properties**. Check the check box Start when TCP/IP is started and click **OK** as shown in Figure 7-22.

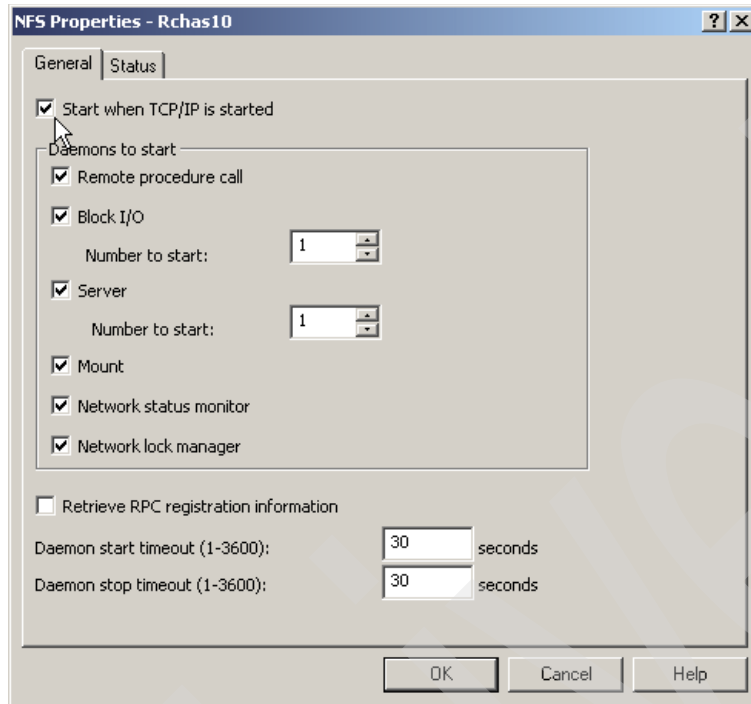


Figure 7-22 Start NFS when TCP/IP is started

- From the IFS, select the desired directory that will hold the Linux backup. Right-click and select **NFS Export** and then select **Properties** as shown in Figure 7-23.

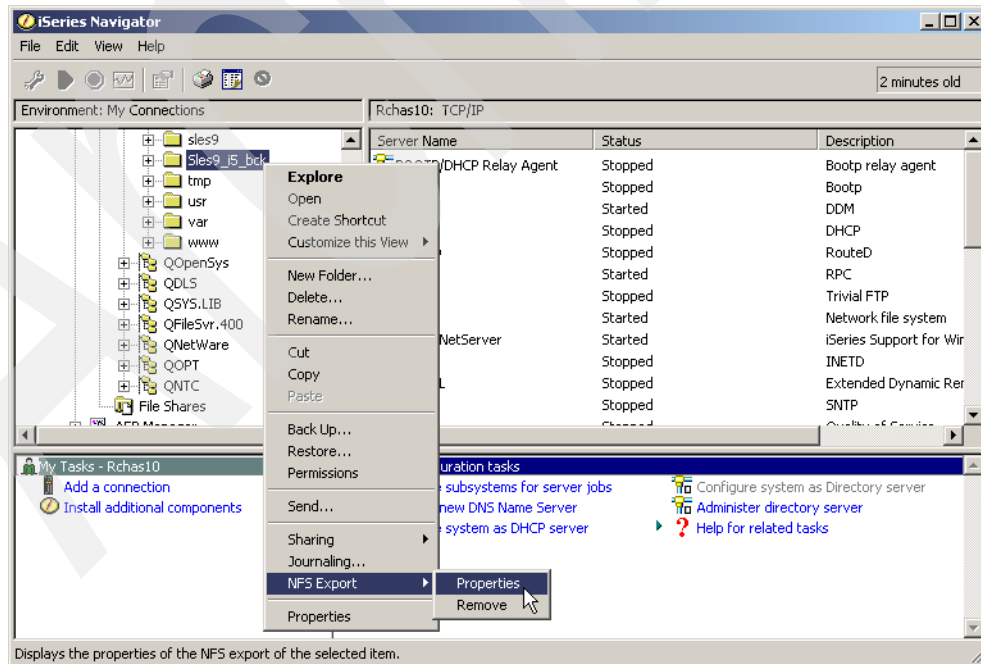


Figure 7-23 Select NFS Export - Properties

6. Check the **Add to list of permanent defined export** check box as illustrated in Figure 7-24. An entry is added to IFS /etc/EXPORTS file.

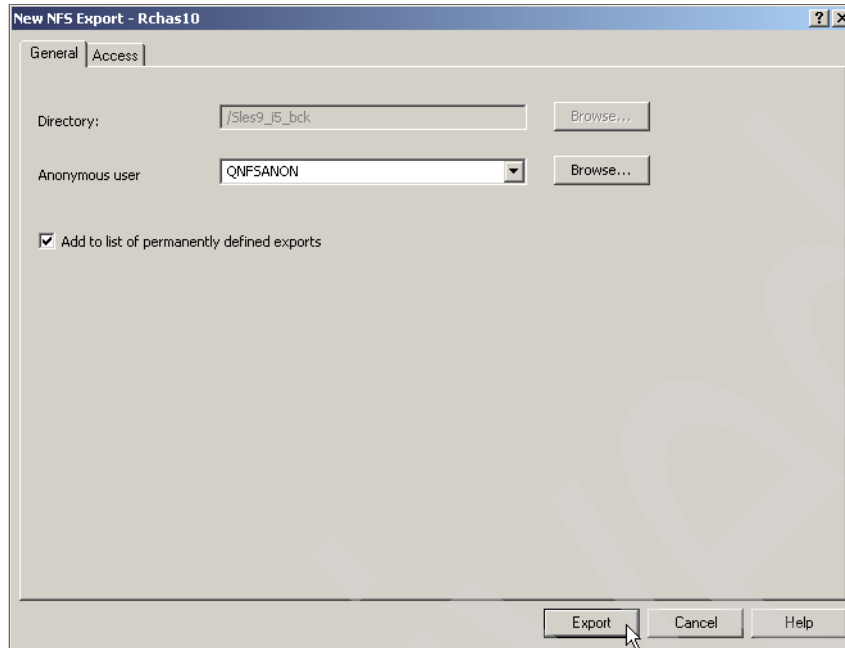


Figure 7-24 Export an IFS directory

7. Mount this NFS directory to Linux file system.

### **Mounting the NFS directory to Linux file system**

1. Mount NFS directory to Linux. Use **Webmin** → **System** → **Disk and Network Filesystems** as shown in Figure 7-25.

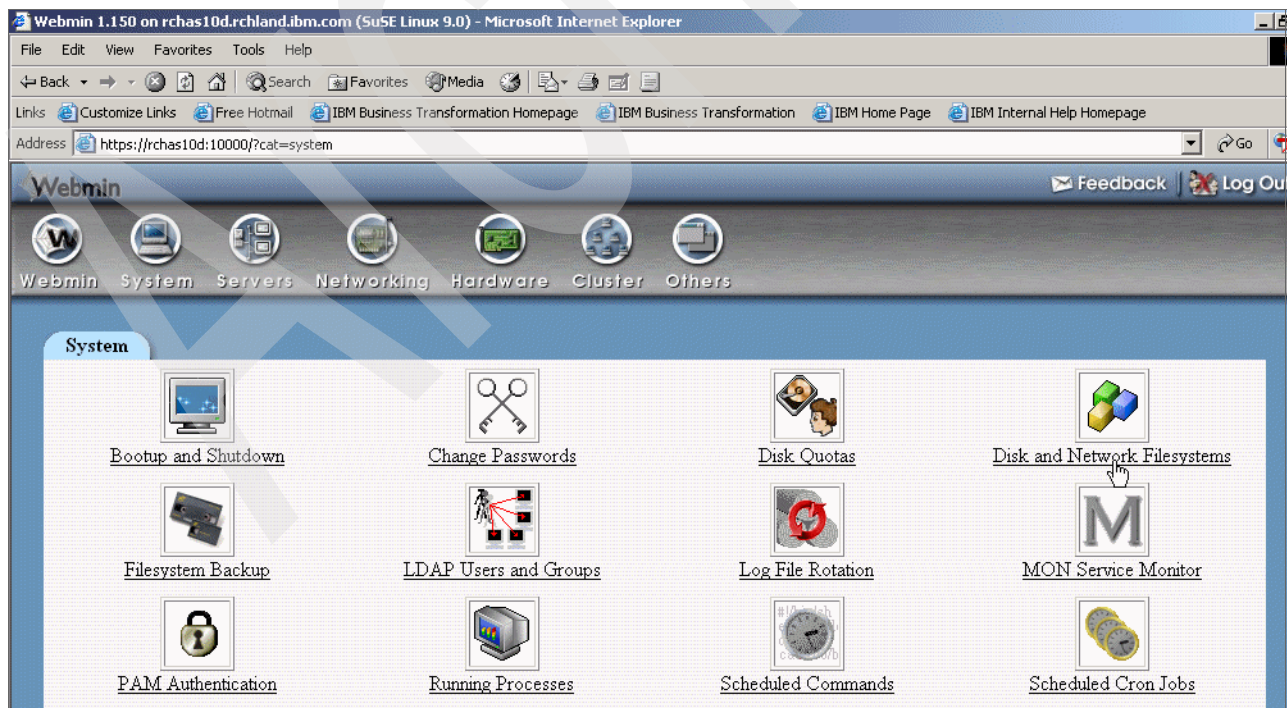


Figure 7-25 Select Disk and Network Filesystems

- From the pull-down menu, first select **Network Filesystem (nfs)**, then click **Add Mount** as illustrated in Figure 7-26.

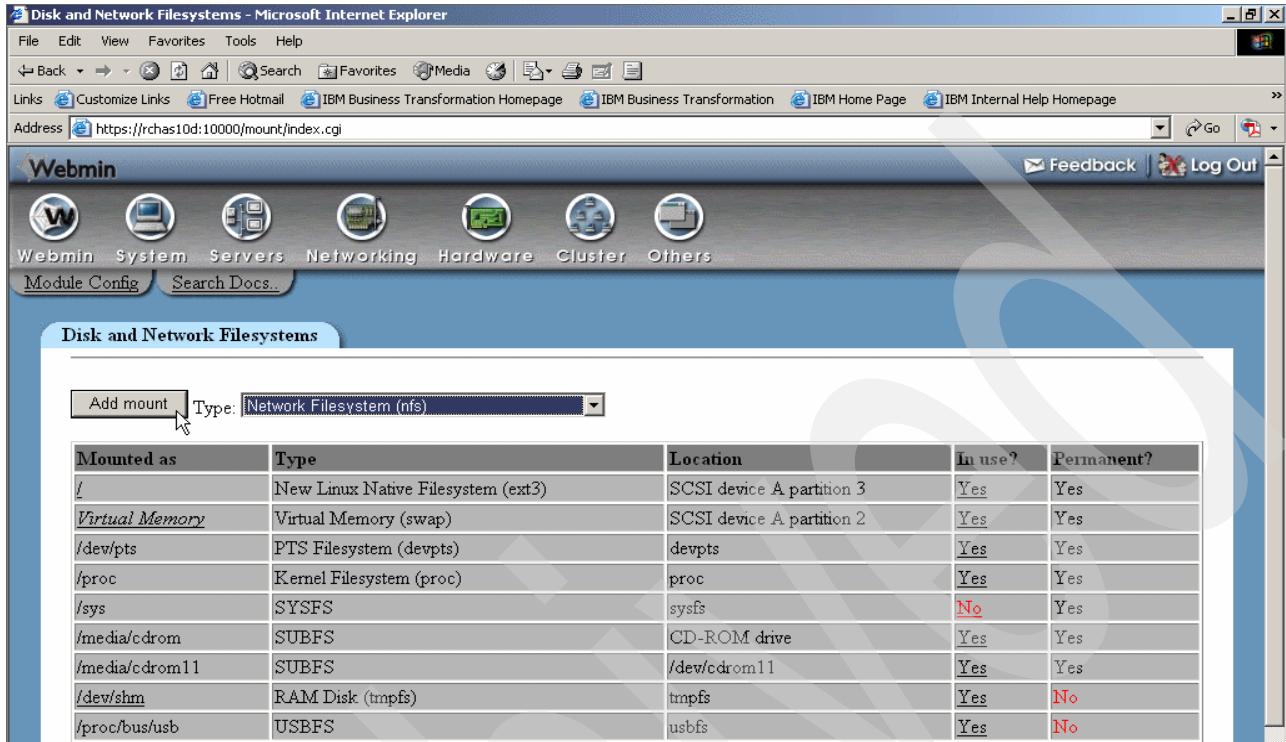


Figure 7-26 Add Mount nfs

3. Select **NFS Hostname** (192.168.1.1 is IP address of i5/OS partition on virtual ethernet line) as shown in Figure 7-27.

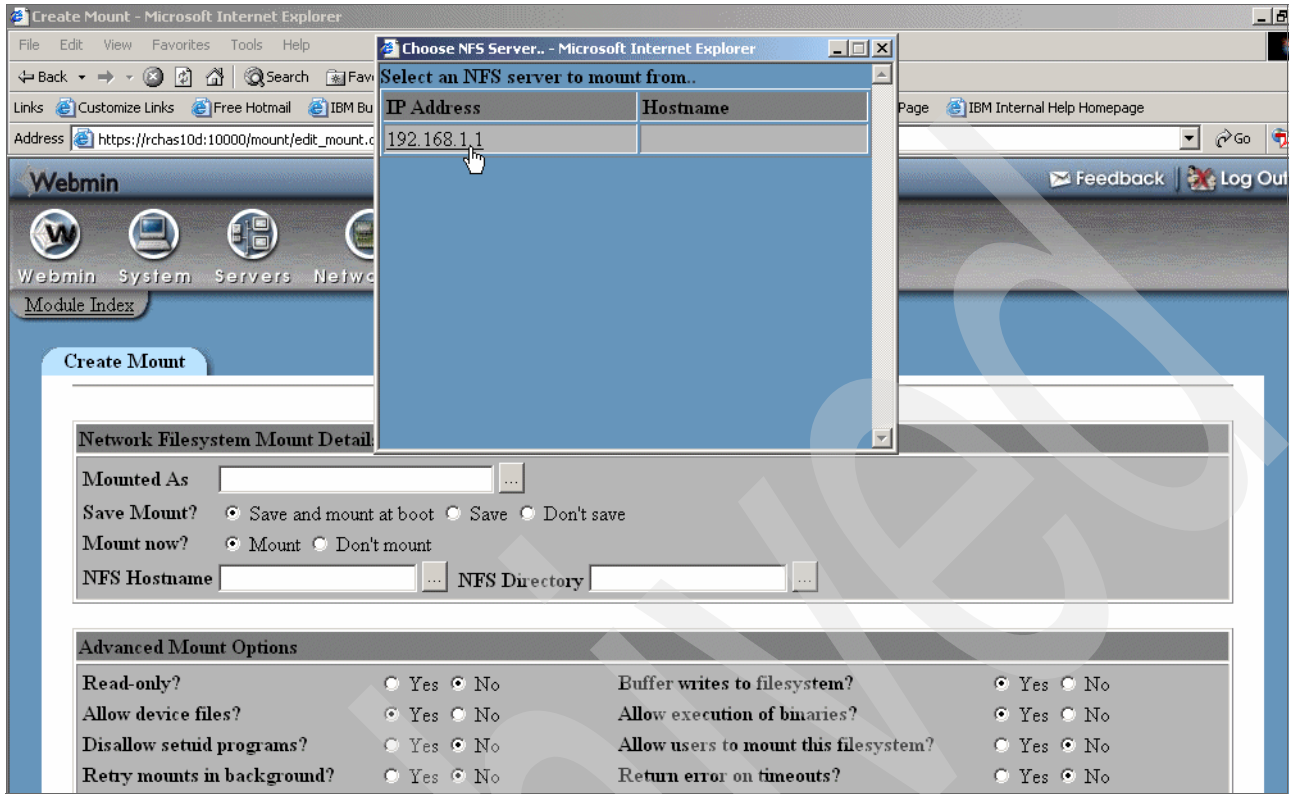


Figure 7-27 Select NFS Hostname

4. Select NFS directory which we exported using iSeries Navigator as shown in Figure 7-28.

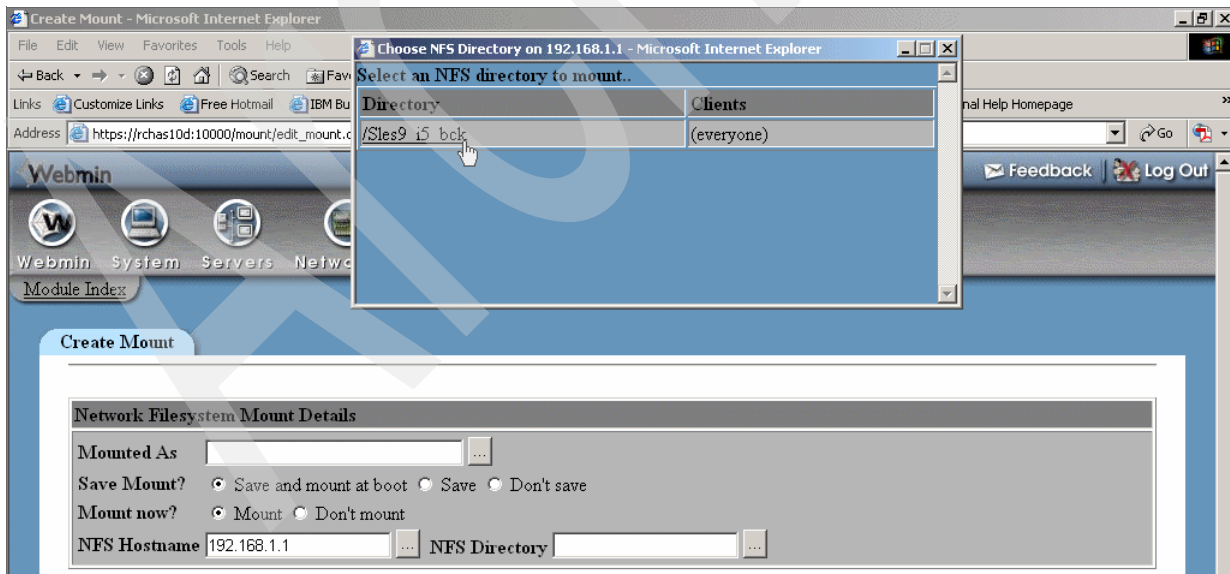


Figure 7-28 Select NFS directory

5. Type a name of directory which you want to create for this operation to **Mounted As** field. Check the radio buttons **Save and restart at boot** and **Mount**. Then press the **Create** button as shown in Figure 7-29.

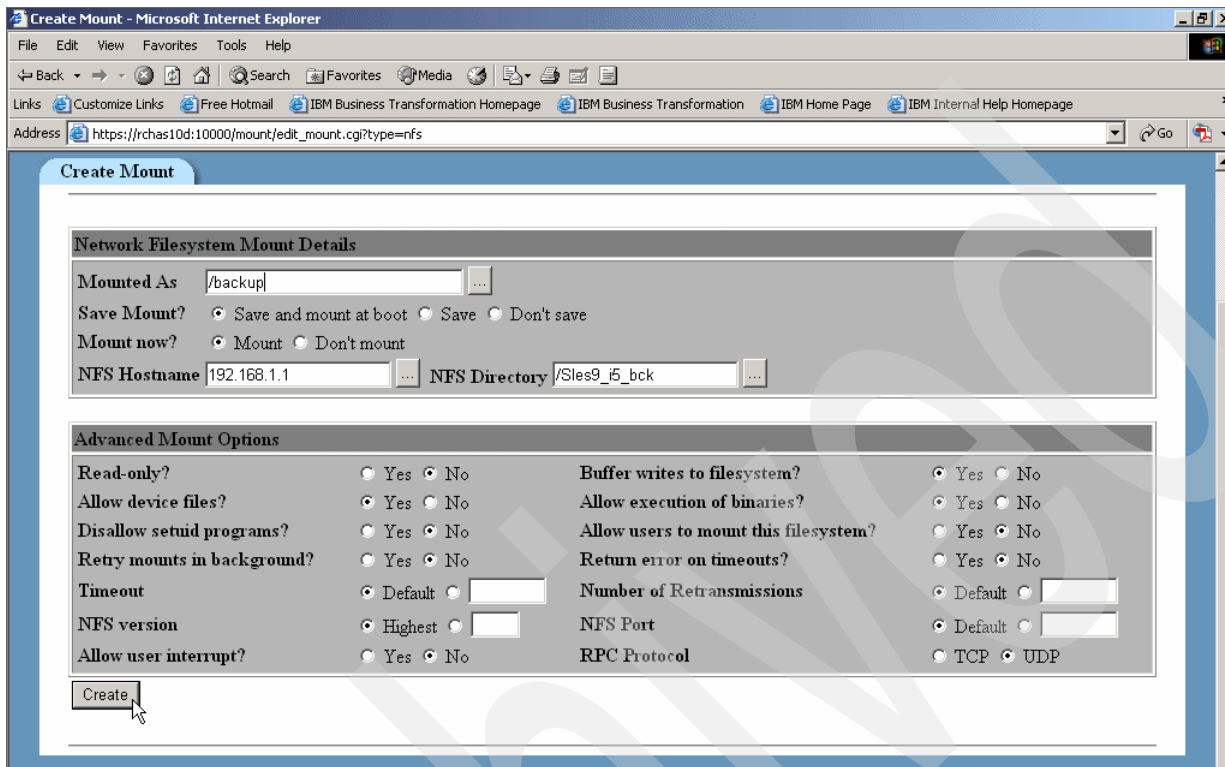


Figure 7-29 Create mount in /backup directory

6. Check if mount was created as shown in Figure 7-30.

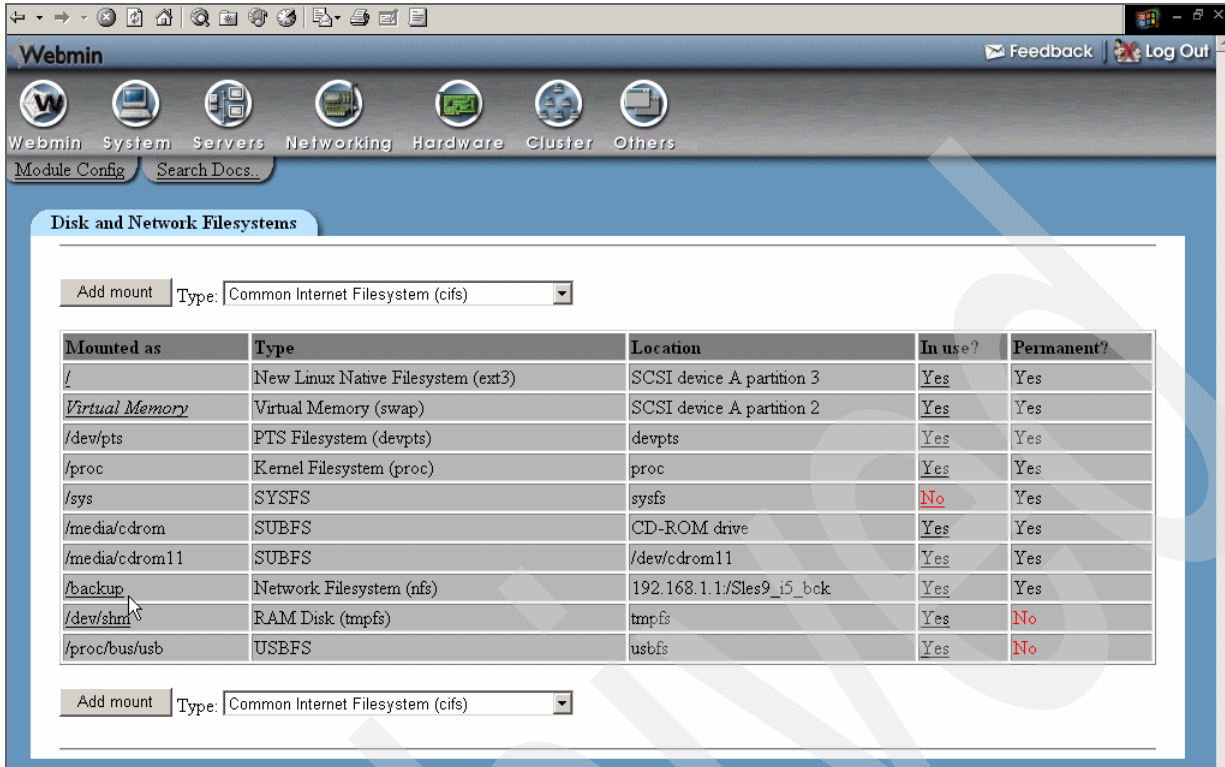


Figure 7-30 Check created mount

### Creating a backup script

There can be a number of backup methods in Linux depending on the commands you use or depending on the type of applications or services you are backing up. In this section, we provide a method in case of a file server application.

Assume that our Linux system is defined as File server and the actual directory we are saving is “/home” directory.

1. From Webmin main menu, select **Others** → **File Manager**. Then create bckhome.sh script file under root directory as illustrated in Figure 7-31.

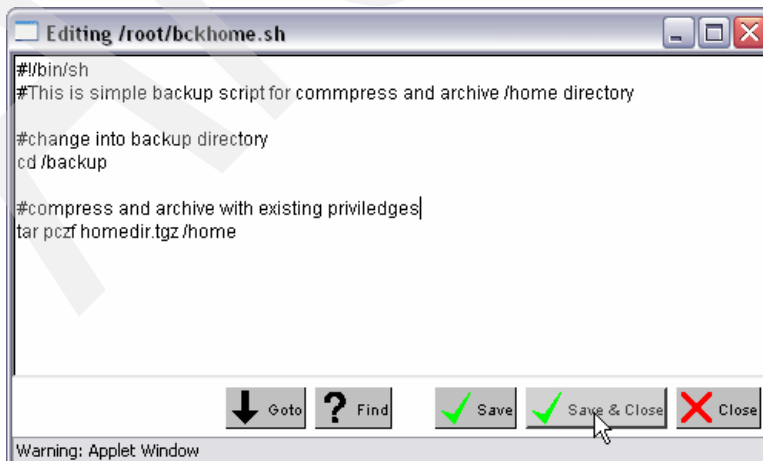


Figure 7-31 Backup script

2. Add it the privileges to be executed like on Figure 7-32.

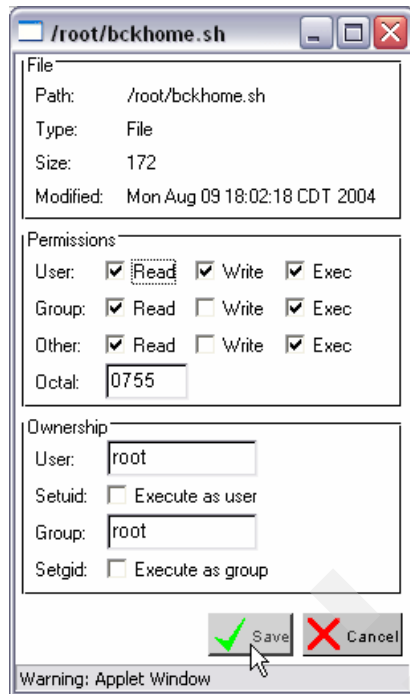


Figure 7-32 File privileges for backup script

Now you can execute this script any time you want to backup /home directory to the i5/OS IFS file, "homedir". Next we will schedule it for a periodic backup.

### **Scheduling backup**

Let's assume that your backup strategy dictates a daily backup for all new and changed documents should take place at 23:00 and all the backup copies should be placed on i5/OS side. That means we need the backup of Linux files and directories on i5/OS side should take



place right before 23:00, for example at 22:30. This section shows you how to accomplish this task.

1. Select **Webmin** → **System** → **Scheduled Cron Jobs**. Choose **Create a new scheduled cron job** link. Fill Job Details as shown in Figure 7-33.

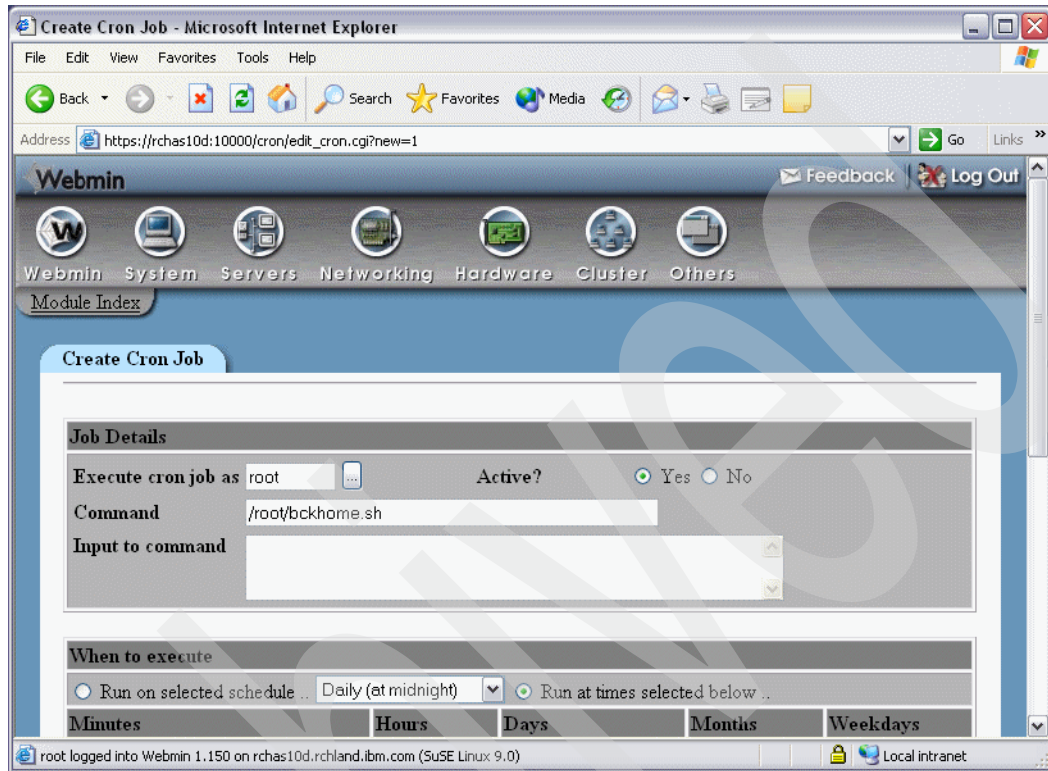


Figure 7-33 Create Cron Job - Job Details

2. Then choose the time when to execute and press **Create** as shown in Figure 7-34.

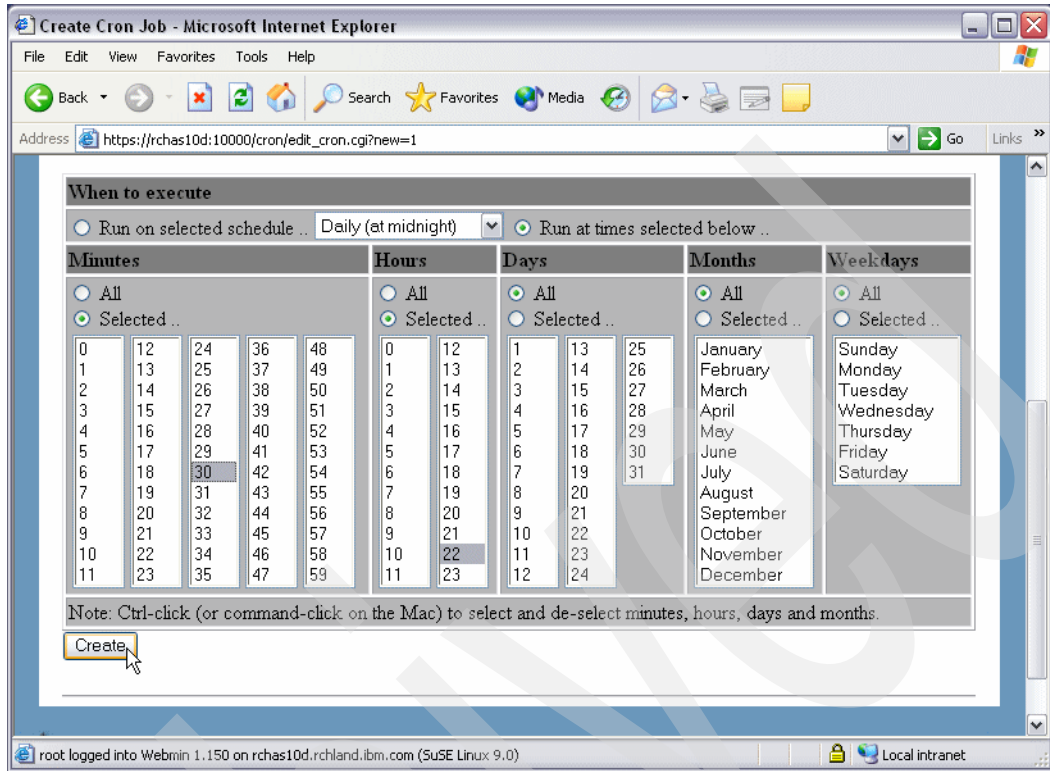


Figure 7-34 Create Cron Jobs - When to execute

3. You will now be able to see this job in Scheduled CronJobs as shown in Figure 7-35.

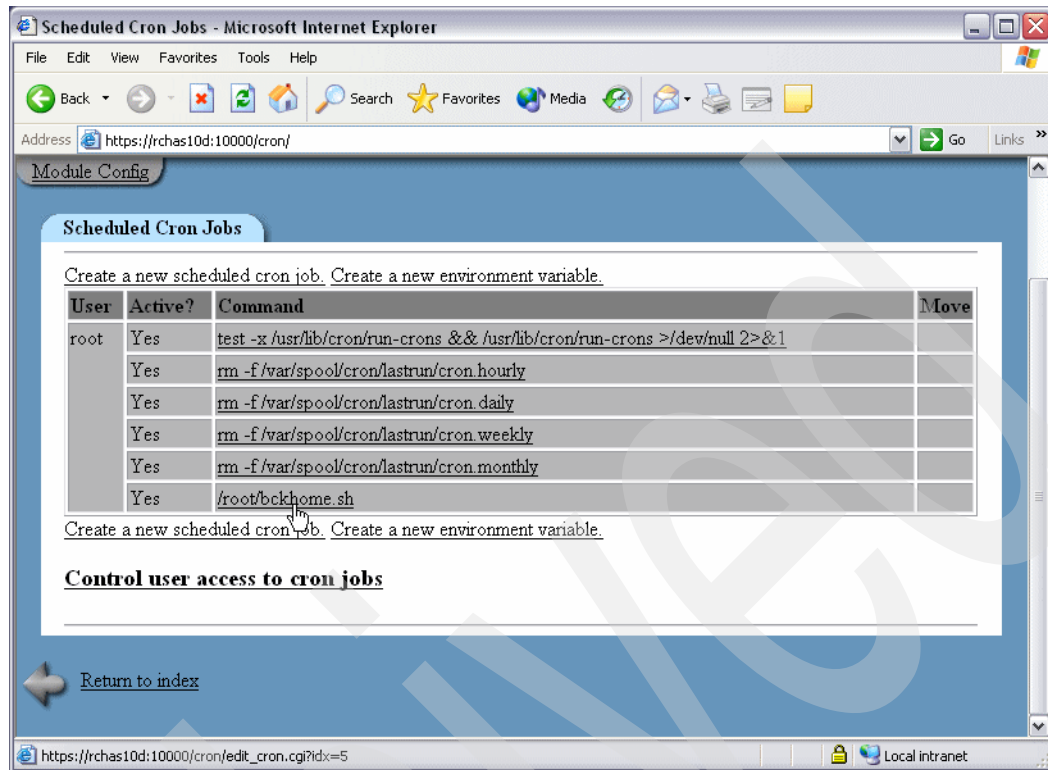


Figure 7-35 Schedule Cron Jobs

4. Synchronize time between i5/OS and Linux partitions.

- The next day, you will see the homedir.tgz file in the /Sles9\_i5\_bck directory in IFS, your i5/OS with time of creation about 22:30 as shown in Figure 7-36.




Figure 7-36 iSeries Navigator - Linux backup file

### 7.3.3 Linux-oriented backup

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 NFS directory but directly on tape. Tape is directly connected to i5/OS environment but can be used in Linux as virtual tape.

There are three steps which you should follow to successfully save Linux files and directories on tape connected to i5/OS partition:

- Initialize tape for Linux. On i5/OS, the standard code for coding letters and numbers used is EBCDIC, while on Linux it is ASCII. This means in the scenario of Linux-oriented backup, you need to strictly sort out tapes 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)
```
- Find the name convention for i5 tape in Linux. Use the `dmesg | grep st` command in ssh Putty to find the name of tape you want to use to Linux backup. In our scenario, we have two tape devices, st0 and st1, and the messages during boot regarding tapes are shown in Figure 7-37.

```
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 7-37 dmesg - boot messages regarding tapes

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 following command:

```
find / -name st1
```

We receive the following output as illustrated in Figure 7-38.

```
rchas10d:~ # find / -name st1
/sys/class/scsi_tape/st1
/dev/st1
```

Figure 7-38 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 the following command:  
VRYCFG CFGOBJ(TAP02) CFGTYPE(\*DEV) STATUS(\*OFF) RANGE(\*OBJ)
5. Use ssh Putty screen to check status of tape. Use the following command:

```
mt -f /dev/st1 status
```

We receive the output as illustrated in Figure 7-39.

```
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 7-39 Tape status

**Tip:** When you receive error message such as “mt: /dev/st1: Input/output error” try following:

- ▶ Firstly, check that the tape cartridge is not ejected.
- ▶ Use the `sg_start /dev/sgx 1` command.

Here, you find `x` by completing information from `sg_map` command and `dmesg | grep sg` and `dmesg | grep st` commands.

6. Use **Webmin** → **System** → **Filesystem Backup**. Choose the directory you want to backup from the pull-down menu first, then click the **Add a new backup of directory** button as illustrated in Figure 7-40.

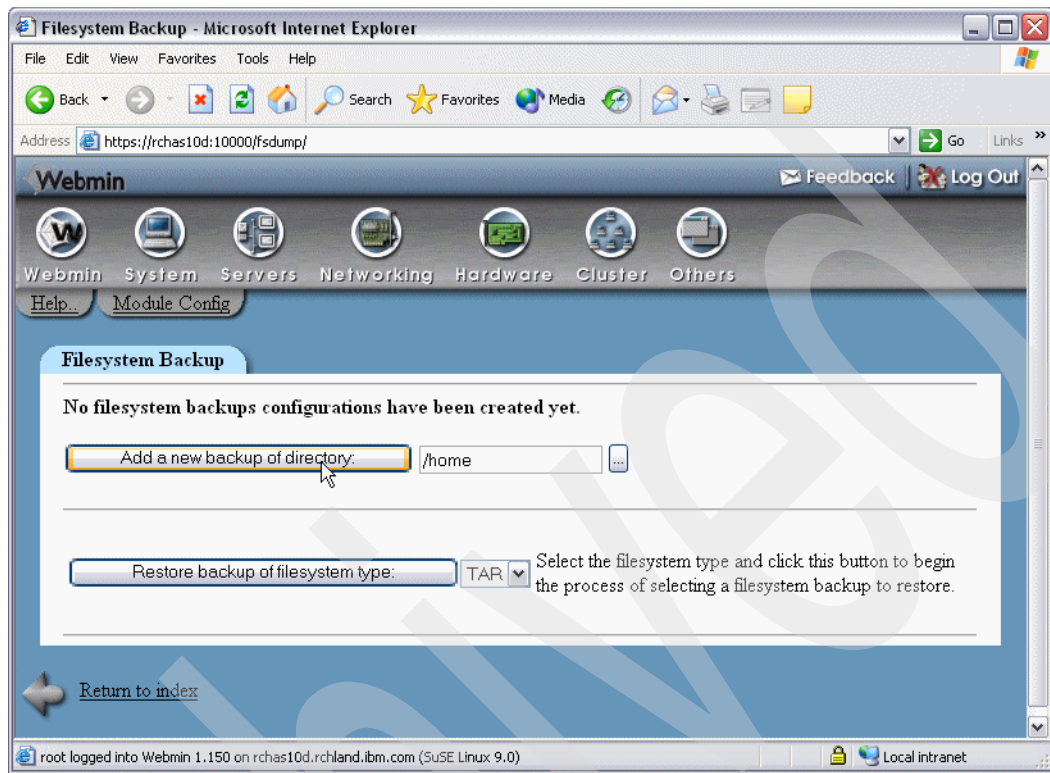


Figure 7-40 Webmin - Filesystem Backup - Add a new backup directory

7. Fill the appropriate data to Backup to and Backup label fields as illustrated in Figure 7-41.

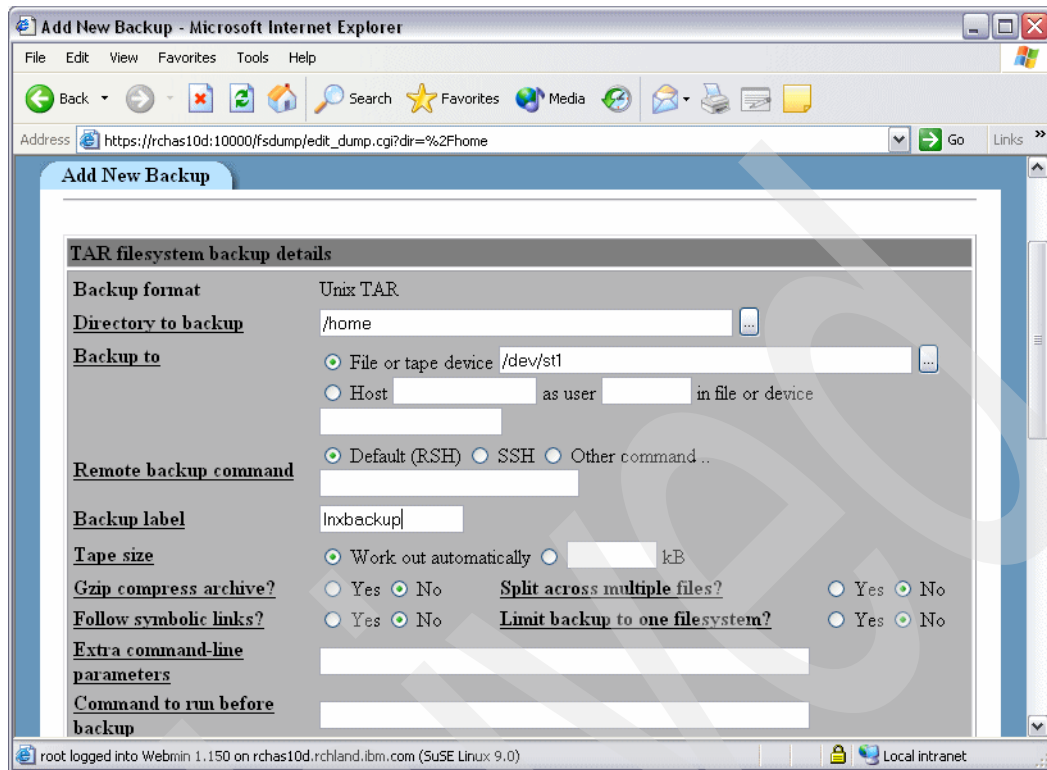


Figure 7-41 Webmin - Add New Backup

8. If you want to schedule this backup choose this option and set the appropriate time. Then press one of the buttons below depending if you choose scheduling or not as illustrated in Figure 7-42.

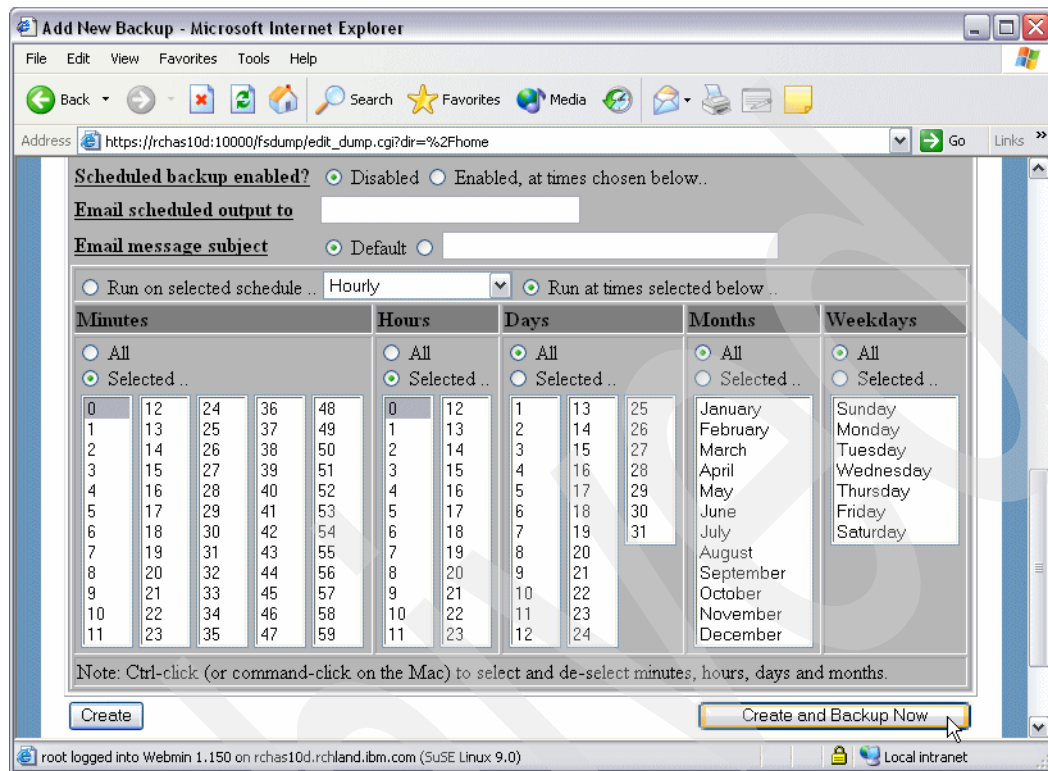


Figure 7-42 Webmin - Add New Backup - continue



9. Press the **Create and Backup Now** button and you will receive following message as illustrated in Figure 7-43.

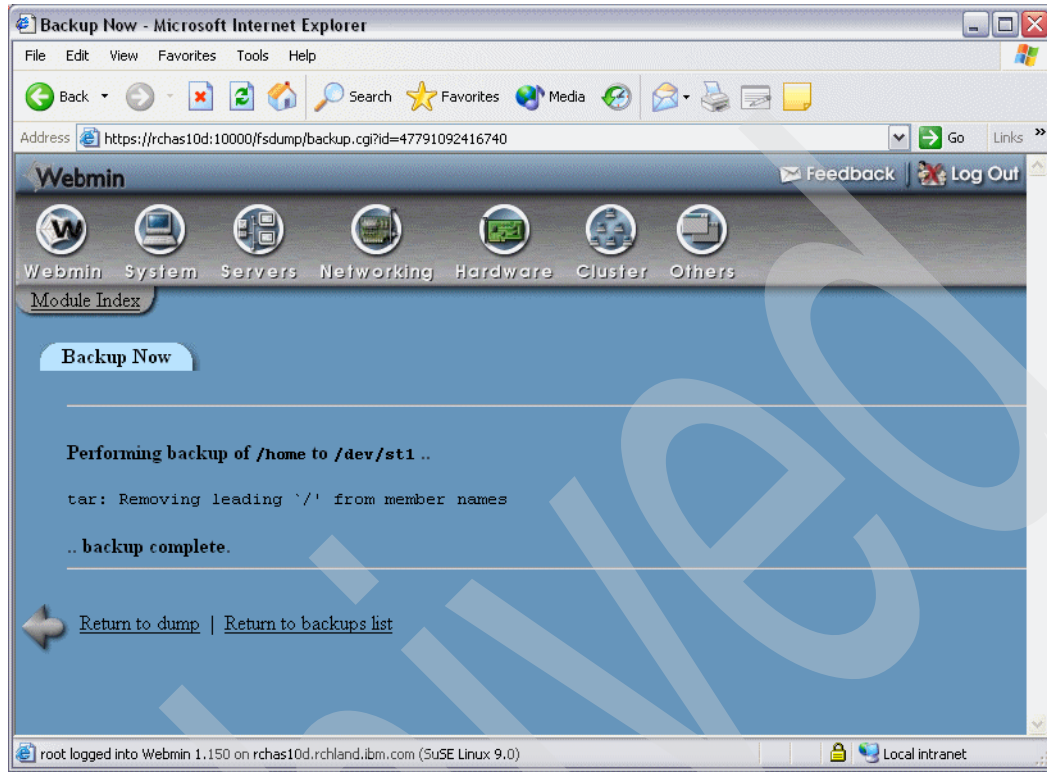


Figure 7-43 Webmin - Backup now

10. Eject the cartridge and deactivate tape device on Linux to be able use it on i5/OS. Use following Linux commands:

```
mt -f /dev/st1 rewind
```

To rewind the tape, use the following command:

```
mt -f /dev/st1 eject
```

To eject the tape, use the following command:

```
sg_map
```

To recognize mapping among default and generic devices, we receive output which is shown in Example 7-1.

Example 7-1 Mapping among default and generic devices

```
rchas10d:~ # sg_map
Error opening /dev/st0 : Input/output error
/dev/sg0 /dev/sda
/dev/sg1 /dev/sdb
/dev/sg2 /dev/sr0
/dev/sg3 /dev/sr1
/dev/sg4
/dev/sg5 /dev/st1
```

To release tape device from Linux, use the following command:

```
sg_reset -d /dev/sg5
```

11. Label this tape correctly to be recognized among i5/OS backup tapes. Now the Linux side backup operation is completed.

12. Vary on tape in i5/OS partition so that the device now can be used by i5/OS folks. Use the following command:

```
VRYCFG CFGOBJ(TAP02) CFGTYPE(*DEV) STATUS(*ON)
```

### 7.3.4 i5/OS oriented restore

This section describes how to restore data saved in 7.3.2, “i5/OS oriented backup” on page 264.

#### Restoring NWSD

You can restore NWSD called LINUXD saved by option 10 from Save menu, as mentioned in 7.3.2, “i5/OS oriented backup” on page 264 by using the following command:

```
RSTCFG OBJ(LINUXD) DEV(*SAVF) OBJTYPE(*NWS) SAVF(LNXBACKUP/I5CFG)
```

When saved to tape by option 21 or 22 from Save menu, you can restore it with the following command:

```
RSTCFG OBJ(LINUXD) DEV(TAPOX) OBJTYPE(*NWS)
```

#### Recovery for NWS Storage Spaces level

You can restore the NWS Storage Space saved by SAV command to SAVF from 7.3.2, “i5/OS oriented backup” on page 264 with the following command:

```
RST DEV('/QSYS.LIB/LNXBACKUB.LIB/NWSS1.FILE') OBJ('/QFPNWSSTG/LINUX')
```

To display objects saved in <savf>, if you are not sure which NWS Storage Space is saved in the SAVF, you can use the following command:

```
DSPSAVF <library>/<savf>
```

When saved to tape by option 21 or 23 from Save menu, you can restore it with the following command:

```
RST DEV('/QSYS.LIB/TAPOX.DEVD') OBJ('/QFPNWSSTG/LINUX')
```

**Important:** The link between NWSD and NWS Storage space is saved in both.

#### Recovery for file and directory level

Restore files from tape with the following command:

```
RST DEV('/QSYS.LIB/TAPOX.DEVD') OBJ('/S1es9_i5_bck/*')
```

You can use following script used for uncompress and extract files from archive file to /tmp directory as illustrated in Figure 7-44.

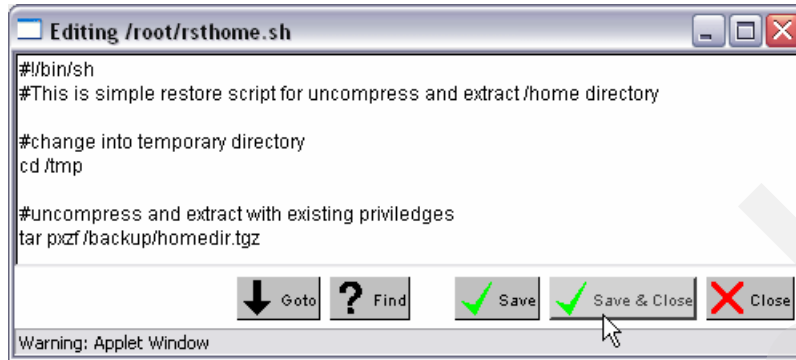


Figure 7-44 Restore script

Restore lost files and delete extracted files from /tmp directory.

### 7.3.5 Linux-oriented restore

To perform a Linux-oriented restore, perform the following steps:

1. Insert tape cartridge with Linux backup which was done on base instruction in 7.3.3, "Linux-oriented backup" on page 278.
2. Vary off tape device in i5/OS partition.
3. Use ssh Putty to issue the following command:  
`cd /tmp && tar pxf /dev/st1`  
This restores /home directory to /tmp/home directory.
4. Eject the tape cartridge and deallocate tape from Linux.  
Vary on tap device on i5/OS.

Archived



## Leveraging virtualization for specific requirements

Virtualization of I/O resources provides greater flexibility in the allocation and management of storage for a Linux instance. This chapter will present information about managing the storage environment including the following:

- ▶ Resizing an existing virtual storage space
- ▶ Allocating new storage to a Linux instance through use of additional storage spaces
- ▶ Using Logical Volume Manager in Linux to build a flexible storage environment

## 8.1 Adding space

Virtualization provides a number of ways of adding space that can be used by a Linux partition and the method depends on a number of factors including availability of the server, as well as the structure of the Linux disk partitions. This section will document two methods, increasing the size of an existing virtual storage space and creating a new storage space for the partition.

## 8.2 Increasing size of storage space

This method for adding additional space to the Linux partition can be used when the Linux file system to be resized is the last file system on the disk (storage space). As an example, imagine that the disk partitions on the storage are configured as shown in Figure 8-1:




Figure 8-1 Linux Disk Structure

**Note:** This structure is the default structure for a Novell / SuSE SLES 9 and SLES10 installation. This structure is not the default structure for RedHat RHEL 4.

Within Linux, the above disk structure is represented by the following devices:

- ▶ /dev/sda1                      PrepBoot
- ▶ /dev/sda2                      swap
- ▶ /dev/sda3                      / (root)

Perform the following steps to resize the storage space and increase the size of the last disk partition (in this case the root disk partition - /dev/sda3):

1. A copy of the existing storage space will be made which requires that the Network Server be varied off. Type the WRKCFGSTS \*NWS command and press **Enter**. The Work with Configuration Status screen (shown in Figure 8-2) will be displayed.

```
Work with Configuration Status G10B8DAE 08/02/06 05:38:06
Position to Starting characters

Type options, press Enter.
 1=Vary on 2=Vary off 5=Work with job 8=Work with description
 9=Display mode status 13=Work with APPN status...

Opt Description Status -----Job-----
 LINUX ACTIVE

Parameters or command
====>
F3=Exit F4=Prompt F12=Cancel F23=More options F24=More keys

Bottom
```

Figure 8-2 Work with Configuration Status

2. Position the cursor in the Opt field of the appropriate Network Server, type 2 (Vary off) and press **Enter**.
3. Once the Network Server is in a Varied Off state, press **F3** to exit.

4. Type the WRKNWSSTG command and press **Enter**. The Work with network Server Storage Spaces screen (shown in Figure 8-3) will be displayed.

```

Work with Network Server Storage Spaces
 System: G10B8DAE
Type options, press Enter.
 1=Create 2=Change 3=Copy 4=Delete 5=Display 6=Print 10=Add link
 11=Remove link

Opt Name Server Seq Link Stg
 Name Server Seq Type Access Path

 LINUX LINUX 1 *DYN *UPDATE

Parameters or command
====>
F3=Exit F4=Prompt F5=Refresh F6=Print list F9=Retrieve
F11=Display disk status F12=Cancel F17=Position to
Bottom

```

Figure 8-3 Work with Network Server Storage Spaces

5. Position the cursor in the Opt field of the Network Server Storage Space to be copied, type 3 (Copy) and press **Enter**.
6. On the Create NWS Storage Space screen (shown in Figure 8-4), enter the following information:
  - **Network server storage space:** Provide a name for the new network server storage space.
  - **Size:** Provide the size of the storage space. This size should reflect the size of the existing storage space plus the additional space to be added.
  - **Text ‘description’:** Provide a description for the new storage space.



```

Create NWS Storage Space (CRTNWSSTG)

Type choices, press Enter.

Network server storage space . . newstorage Name
Size > 10240 *CALC, 1-1024000 megabytes
From storage space > LINUX Name, *NONE
Auxiliary storage pool ID . . . > 1 1-255
ASP device Name
Text 'description' 'Copy of storage space'

Bottom
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

```

Figure 8-4 Create NWS Storage Space

7. When you have completed the fields, press **Enter**. The existing storage space will now be copied into the new storage space.

**Note:** The copy of an existing storage space can be a time consuming task.

8. When the copy of the storage space is completed, the Work with Network Server Storage Spaces screen will be re-displayed.

- Position the cursor in the Opt field of the new Network Server Storage space, type 10 (Link) and press **Enter**. The Add Network Server Storage Link screen () will be displayed as shown in Figure 8-5.

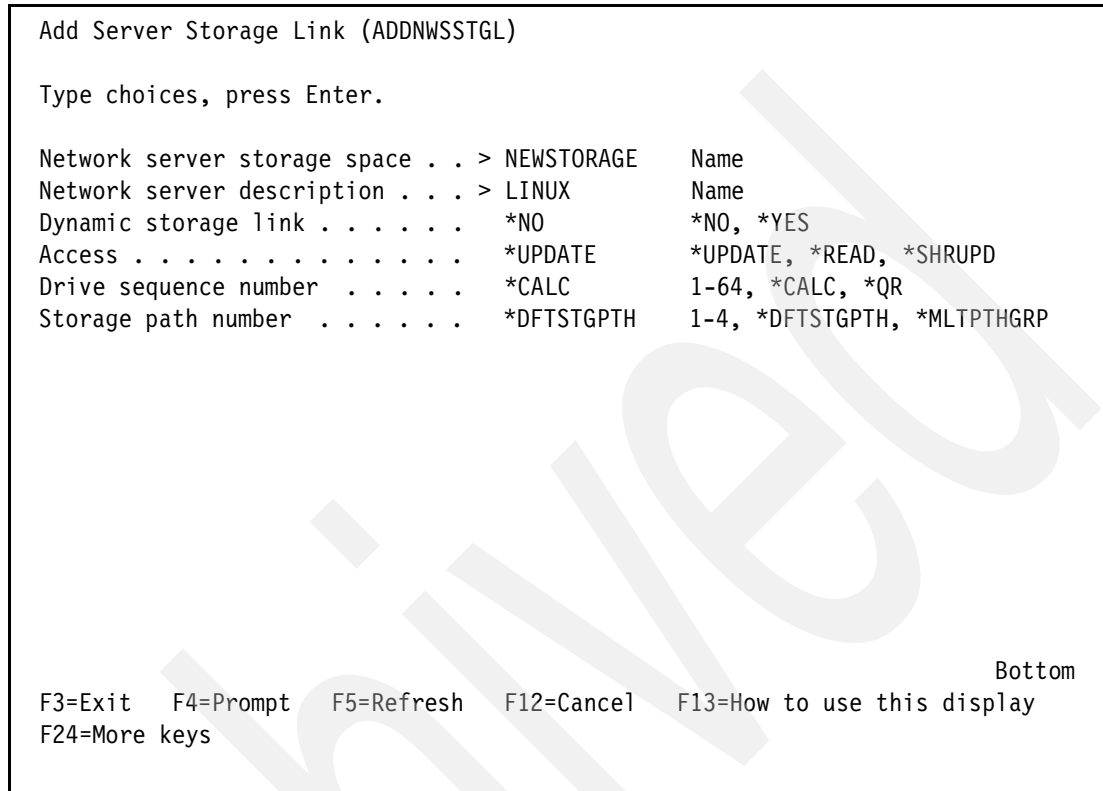


Figure 8-5 Add Network Server Storage Link

- Type the name of the Network Server in the Network server storage space field. Leave the remaining fields at their default values and press **Enter**.
- Press **F3** on the Work with Network Server Storage Space screen.
- Type the WRKCFGSTS \*NWS command and press **Enter**. The Work with Configuration Status window shown in Figure 8-2 will be displayed.
- Position the cursor in the Opt field of the Network Server, type 1 (Vary on) and press **Enter**. At this point the Network Server will be varied on and Linux will be started.
- Once Linux is started, access the console (or terminal session) for Linux and login as the root user.

15. In Linux, enter the `fdisk -l` command to display a list of disk partitions. Output similar to Figure 8-6 will be displayed.

```
Disk /dev/sda: 6440 MB, 6440394240 bytes
199 heads, 62 sectors/track, 1019 cylinders
Units = cylinders of 12338 * 512 = 6317056 bytes

 Device Boot Start End Blocks Id System
/dev/sda1 * 1 1 264+ 41 PPC PReP Boot
/dev/sda2 4 127 755055 82 Linux swap / Solaris
/dev/sda3 127 1020 5510295 83 Linux

Disk /dev/sdb: 10.7 GB, 10742212096 bytes
64 heads, 32 sectors/track, 10244 cylinders
Units = cylinders of 2048 * 512 = 1048576 bytes

 Device Boot Start End Blocks Id System
/dev/sdb1 * 1 1 264+ 41 PPC PReP Boot
/dev/sdb2 24 761 755055 82 Linux swap / Solaris
/dev/sdb3 761 6143 5510295 83 Linux
```

Figure 8-6 Output from the `fdisk -l` command

**Note:** Error messages such as partition does not end on cylinder boundary may be displayed. You can ignore these.

The start and end cylinders of the disk partitions of the two disks may be different. This occurs when the size of the two disks has resulted in a different number of sectors per track (as shown in the above example.) The important thing to realize is that the size of the disk partitions on both disks at this point is exactly the same (as shown in the number of blocks column for the disk partitions of both disks.)

16. The new disk should be the last disk listed (in this case the `/dev/sdb` disk since there are only two disks linked to the network server). Use the `fdisk` command to work with the disk partitioning as shown in Example 8-1.

Example 8-1 The `fdisk` command

```
fdisk /dev/sdb
```

Here `/dev/sdb` is the name of the new (larger) disk device.

17. The first step is to display the current partition table for the disk. At the `fdisk` prompt enter the `p` command and press **Enter**. The partition table will be displayed similar to what is shown in Figure 8-6.

18. Make note of the Start position (cylinder) of the last disk partition (the disk partition that is going to be resized. This value (in this example 761) will be needed in a subsequent step.

19. At this point you are going to delete the last disk partition. Type `d` and press **Enter**.

20. A prompt will be displayed for the Partition number. Type the number of the partition (in our example, the partition number is 3 as indicated by the 3 in the device name) and press **Enter**.

21. Now that the partition has been deleted, it can be re-created. Type `n` and press **Enter**.

22. When prompted for the partition type (as shown in Figure 8-7), type `p` for primary partition and press **Enter**.

```
Command (m for help): n
Command action
 e extended
 p primary partition (1-4)
```

Figure 8-7 `fdisk` - Partition Type prompt

23. When prompted for the partition number, enter the same number as the partition that was deleted.

24. When prompted for the First cylinder, enter the value of the first cylinder that was recorded in a previous step.

25. When prompted for the Last cylinder, the default will be the last cylinder on the disk (which is what we want), therefore just press **Enter**. The prompts and responses should resemble Figure 8-8.

```
Command (m for help): n
Command action
 e extended
 p primary partition (1-4)
p
Partition number (1-4): 3
First cylinder (1-10244, default 1): 761.
Last cylinder or +size or +sizeM or +sizeK (761-10244, default 10244):
Using default value 10244
```

Figure 8-8 `fdisk` - Partition Creation prompts

**Note:** Your values for partition number, first cylinder and last cylinder may be different depending on the disk layout and size of the disk.

26. Once the new partition has been created, type `w` and press **Enter** to cause the disk partition to be written.

27. At this point, the disk partition has been resized. Now the file system on the disk partition should be resized. Use either of the following commands depending on the format of the disk partition:

- ext2 & ext3 file systems: `resize2fs`
- reiserfs file systems: `resize_reiserfs`

As an example, if the new disk partition is the third partition on the second disk and it is formatted as a reiser file system, then the command to resize the disk will be as shown in Example 8-2.

Example 8-2 *Resizing the file system*

---

```
resize_reiserfs /dev/sdb3
```

---

28. Once the resize of the file system is complete, you can shut down the Linux partition and vary off the network server.
29. Back in i5/OS, type the WRKNWSSTG command and press **Enter**. The Work with Network Storage Spaces screen shown in Figure 8-9 will be displayed.

```

Work with Network Server Storage Spaces
System: G10B8DAE
Type options, press Enter.
 1=Create 2=Change 3=Copy 4=Delete 5=Display 6=Print 10=Add link
 11=Remove link

Opt Name Server Seq Link Stg
 Name Server Type Access Path

 LINUX SLES10 1 *DYN *UPDATE
 NEWSTORAGE SLES10 2 *DYN *UPDATE

Parameters or command
====>
F3=Exit F4=Prompt F5=Refresh F6=Print list F9=Retrieve
F11=Display disk status F12=Cancel F17=Position to
Bottom

```

Figure 8-9 Working with network server storage spaces

30. Position the cursor in the opt field of the original storage space (that is, the smaller storage space), type 11 (Remove link) and press **Enter**. The Remove Server Storage Link window shown in Figure 8-10 will be displayed.

```

Remove Server Storage Link (RMVNWSSTGL)
Type choices, press Enter.
Network server storage space . . . > LINUX Name
Network server description . . . > LINUX Name
Renumber link *YES *NO, *YES

Bottom
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

```

Figure 8-10 Remove Server Storage Link

31. Change the value of the 'Renumber link' field to \*YES. This will cause the new storage space to assume the link value of this storage space which will cause it to assume the same Linux device name. Once the field value has been changed press **Enter** to remove the link.
32. Now vary the network server back on to start Linux. Once Linux has been started verify that the resized file system has additional space available.

33. Once you are satisfied that the disk has been successfully resized, the original storage space can be deleted.

## 8.2.1 Optional - Rename Storage Space

The steps in the previous section cause the name of the production storage space to be changed. The following steps will rename the new storage space back to the original name:

**Note:** These steps assume that the original (smaller) storage space has already been deleted.

1. Ensure that the Linux operating system is shutdown and that the Network Server is in a varied off state.
2. Type the WRKNWSTG command and press **Enter**. The Work with Network Server Storage Space screen (shown in Figure 8-11) will be displayed.

```

Work with Network Server Storage Spaces
System: G10B8DAE
Type options, press Enter.
 1=Create 2=Change 3=Copy 4=Delete 5=Display 6=Print 10=Add link
 11=Remove link

Opt Name Server Seq Link Stg
 Type Access Path
NEWSTORAGE LINUX 1 *DYN *UPDATE

Parameters or command
====>
F3=Exit F4=Prompt F5=Refresh F6=Print list F9=Retrieve
F11=Display disk status F12=Cancel F17=Position to
Network server storage space SLES10MSTR deleted.
Bottom

```

Figure 8-11 Work with Network Server Storage Spaces

3. Position the cursor in the Opt field of the Network Server Storage Space to be renamed, type 11 (Remove link) and press **Enter**.
4. When the Remove Server Storage Link screen is displayed, press **Enter**.
5. Press **F3** to exit the Work with Network Server Storage Spaces command.
6. Type the WRKLNK 'QFPNWSSTG' command and press **Enter**.

- On the Work with Object Links screen (shown in Figure 8-12), position the cursor in the opt field of the qfpnwsstg object link, type 5 (display) and press **Enter**.

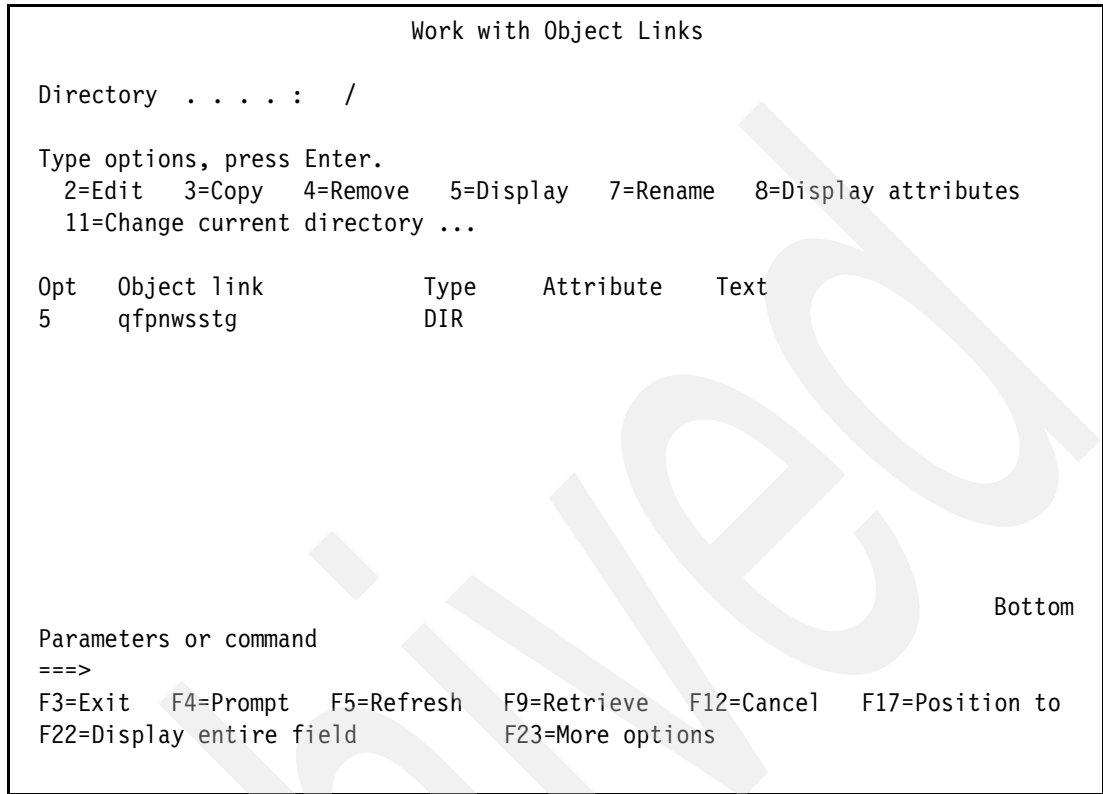


Figure 8-12 Work with Object Links

- When the list of object links is displayed (as shown in Figure 8-13), position the cursor in the Opt field of the storage space to be renamed, type 7 (Rename) and press **Enter**.

```

Work with Object Links

Directory : /qfpnwsstg

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
7 NEWSTORAGE DIR

Parameters or command
====>
F3=Exit F4=Prompt F5=Refresh F9=Retrieve F12=Cancel F17=Position to
F22=Display entire field F23=More options

Bottom

```

Figure 8-13 Work with Object Links - Rename Request



- When the Rename Object screen is displayed (as shown in Figure 8-14), type the new name in the New object field and press **Enter**.

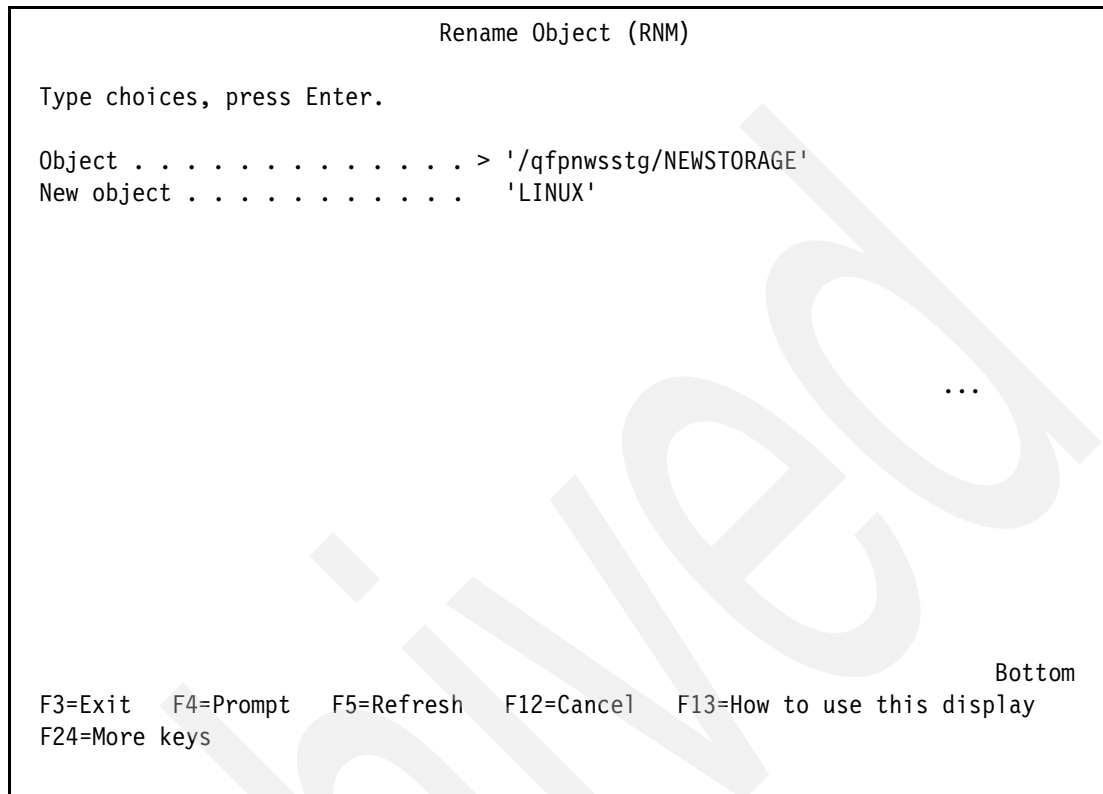


Figure 8-14 Rename Object

- Press **F12** twice.
- Type WRKNWSSTG and press **Enter** to display the list of storage spaces.
- Position the cursor in the 'Opt' field of the renamed storage space. Type 10 (Add link) and press **Enter**.
- When the Add Server Storage Link screen is displayed, enter the name of the Network Server in the Network server description field and press **Enter**.
- At this point, the Network Server can be varied back on.

### 8.3 Adding additional storage spaces

With this method a new Network Storage Space is created in i5/OS and allocated to the Linux Network Server. Linux commands are then used to format and access the disk.

Perform the following steps to add an additional storage space to a Linux partition. This example assumes that there is currently only one (1) disk allocated to the Linux instance:

- In i5/OS issue the following command as shown in Example 8-3.

*Example 8-3 Create Network Server Storage Space syntax*

---

```
CRTNWSSTG NWSSTG(NEWDISK) NWSSIZE(2048) FORMAT(*OPEN) TEXT('Secondary Disk for Linux Partition')
```

---

In Example 8-3:

- The value for NWSSTG is the user-defined name for the Network Server Storage Space
- The value for NWSSIZE is the size (in megabytes) for the new Network Server Storage Space.
- The value for FORMAT is \*OPEN which indicates that a disk partition should not be placed on the disk.
- The value for TEXT is a user-defined description for the Network Server Storage Space.

2. Once the Network Server Storage Space has been created, it needs to be associated with the Network Server as shown in Example 8-4.

*Example 8-4 Add Network Server Storage Link syntax*

---

```
ADDNWSSTGL NWSSTG(NEWDISK) NWSD(LINUX)
```

---

In Example 8-4:

- The value for NWSSTG is the name of the Network Server Storage Space.
- The value for NWSD is the name of the Network Server Description.

Linking a Network Storage Space to a Network Server Description basically places the “disk” on the virtual SCSI bus created by the relationship of the Virtual SCSI Server adapter (in i5/OS) and the virtual SCSI client adapter (in the Linux partition).

Now that the disk has been placed on the SCSI bus, Linux needs to scan the bus to identify the new device.

3. In Linux, issue the following command to scan the bus for the new disk (Example 8-5).

*Example 8-5 Scan scsi bus*

---

```
echo “- -” > /sys/class/scsi_host/host0/scan
```

---

At this point, the SCSI bus has been scanned and a device handle will be created between the Linux operating system and the disk. Now the disk needs to be formatted and made available within the Linux directory tree.

4. The `fdisk` command will be used to place a partition on the disk. In Linux, issue the following command shown in Example 8-6.

*Example 8-6 fdisk*

---

```
fdisk /dev/sdb
```

---

**Note:** This command assumes that the new disk is the second disk linked to the network server. The last letter in the device name (in this case b) indicates the drive sequence number (in this case 2).

The system responds with the following prompt shown in Figure 8-15.

```
Command (m for help):
```

*Figure 8-15 fdisk prompt*

5. To create a new disk partition, type n (for new) and press **Enter**.

The system responds with the following prompt shown in Figure 8-16.

```
Command action
 e extended
 p primary partition (1-4)
```

Figure 8-16 *fdisk - partition type*

6. Type p (for primary) and press **Enter**.

The system responds with the following prompt shown in Figure 8-17.

```
Partition number (1-4):
```

Figure 8-17 *fdisk - partition number*

7. Since this is a new disk, this will be the first partition on the disk, therefore, type 1 and press **Enter**.

The system will respond with the following prompt shown in Figure 8-18.

```
First cylinder (1-1013, default 1):
```

Figure 8-18 *fdisk - starting cylinder*

8. We are going to use the entire disk for a single disk partition, therefore, press **Enter** to accept the default of 1.

The system responds with the following prompt shown in Figure 8-19.

```
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-1013, default 1013):
```

Figure 8-19 *fdisk - last cylinder*

9. Since we are using the entire disk for this partition, press **Enter** to accept the default value (which is the last cylinder on the disk).

The system responds with the following prompt shown in Figure 8-20.

```
Using default value 1013
Command (m for help):
```

Figure 8-20 *Using default value*

10. You can display the resulting partition table by typing `p` (for print) and pressing **Enter**.

The system responds with the following prompt shown in Figure 8-21.

```
Disk /dev/sdb: 2155 MB, 2155023360 bytes
67 heads, 62 sectors/track, 1013 cylinders
Units = cylinders of 4154 * 512 = 2126848 bytes
 Device Boot Start End Blocks Id System
/dev/sdb1 1 1013 2103970 83 Linux
Command (m for help):
```

Figure 8-21 *fdisk - partition table display*

11. At this point the partition table can be written to disk. Type `w` (for write) and press **Enter**.

The system responds with the following prompt shown in Figure 8-22.

```
The partition table has been altered!
Calling ioctl() to re-read partition table.
Syncing disks.
```

Figure 8-22 *fdisk - updating partition table*

Now that the disk partition has been created a file system should be formatted on it.

12. To create an ext3 file system, issue the following command shown in Example 8-7.

*Example 8-7 Create file system - command syntax*

```
mke2fs -j /dev/sdb1
```

The system will respond with the following prompt shown in Figure 8-23.

```
mke2fs 1.38 (30-Jun-2005)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
263296 inodes, 525992 blocks
26299 blocks (5.00%) reserved for the super user
First data block=0
17 block groups
32768 blocks per group, 32768 fragments per group
15488 inodes per group
Superblock backups stored on blocks:
 32768, 98304, 163840, 229376, 294912
Writing inode tables: done
Creating journal (16384 blocks): done
Writing superblocks and filesystem accounting information: done
This filesystem will be automatically checked every 37 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
```

Figure 8-23 *Output from mke2fs command*

**Note:** There are a number of different file systems that are supported in Linux. The following commands are an example of some of the different file systems that can be created:

- ▶ `mke2fs` ext2 file system
- ▶ `mke2fs -j` ext3 file system
- ▶ `mkreiserfs` reiser file system

Once the file system has been created, you should create a mount point.

**Note:** A mount point is a directory within the overall Linux directory tree where the contents of a disk device can be accessed. The directory can be anywhere in the directory tree and it can either be empty or contain other files/directories (although the contents will not be accessible when a disk device is 'mounted' on the directory. By long-standing convention, mount points are usually empty directories under the `/mnt` directory.

13. From Linux, issue the following command to create a mount point called "newdisk" (as shown in Example 8-8).

*Example 8-8 Make directory - command syntax*

---

```
mkdir /mnt/newdisk
```

---

14. Once the mount point has been created the new file system can be mounted (as shown in Example 8-9).

*Example 8-9 Mounting a disk file system*

---

```
mount /dev/sdb1 /mnt/newdisk
```

---

At this point, you can access files from the new disk partition by going through the `/mnt/newdisk` directory.

**Note:** If you want to have the disk automatically mounted every time the Linux partition is started, you should add the following entry to the `/etc/fstab` configuration file:

```
/dev/sdb1 /mnt/newdisk auto defaults 1 1
```

## 8.4 Setting up Logical Volume Manager

Logical Volume Manager (LVM) can be used to combine two or more disks (or disk partitions) into a single Linux file system. Combined with i5/OS Virtual Disk and dynamic disk allocation, LVM provides a unique solution for establishing large and growing file servers, Web servers, and other disk intensive enterprise solutions.

This section gives information about establishing an initial LVM setup by allocating two virtual disks to the Linux partition and using LVM to create a single Linux file system.

## 8.4.1 Creation and allocation of virtual disk

### Creating the virtual disks

The initial LVM environment will be composed of two (2) virtual disks. In i5/OS, use the Create Network Server Storage (CRTNWSSTG) command to create two virtual disks to be used for the initial LVM setup.

1. Create the two initial storage spaces for use with LVM as shown in Example 8-10.

*Example 8-10 Create Network Storage command syntax*

---

```
CRTNWSSTG NWSSTG(DISK2) NWSSIZE(100) FORMAT(*OPEN)
CRTNWSSTG NWSSTG(DISK3) NWSSIZE(100) FORMAT(*OPEN)
```

---

In Example 8-10:

- The value for NWSSTG is the user-defined name for the Network Server Storage space.
- The value for NWSSIZE is the size (in megabytes) for the storage space.
- The value for FORMAT is set to \*OPEN.

### Link the Virtual Disks to the Network Server

Once the virtual disks have been created, they have to be linked to the Network Server Description. The association of virtual disks with the Network Server is accomplished with the Add Network Server Storage Link (ADDNWSSTGL) command with the following steps:

1. Link the storage spaces to the Network Server Description as shown in Example 8-11.

*Example 8-11 Add Network Storage Link command syntax*

---

```
ADDNWSSTGL NWSSTG(DISK2) NWSD(LINUX)
ADDNWSSTGL NWSSTG(DISK3) NWSD(LINUX)
```

---

In Example 8-11:

- The value for NWSSTG is the user-defined name for the Network Server Storage space
- The value for NWSD is the name of the Network Server Description

Once the storage spaces have been linked to the Network Server, the SCSI bus (in Linux) needs to be scanned to make the disks known to the operating system.

2. Issue the following command in Linux to scan the SCSI bus as shown in Example 8-12.

*Example 8-12 SCAN the virtual SCSI bus (Linux)*

---

```
echo "--" > /sys/class/scsi_host/host0/scan
```

---

**Note:** This command assumes that there is only one SCSI client adapter allocated to the Linux partition. If there is more than one SCSI client adapter allocated to the Linux partition, the value host0 will need to change to reflect the correct SCSI client adapter.

At this point the disks are known to the Linux operating system. This can be confirmed through the **fdisk** command:

3. Issue the **fdisk** command in Linux to display the disks/partitions known to the operating system as shown in Example 8-13.

*Example 8-13 fdisk command to list disks/partitions*

---

```
fdisk -l
```

---

You should observe two additional disks listed (they will be /dev/sdc and /dev/sdd they were the second and third disks linked to the Network Server Description.)

**Note:** If the storage spaces were created with a format other than \*OPEN then the resulting disk partition will need to be deleted prior to the next step.

## 8.4.2 Setting up the initial LVM environment

### Initializing the disks

Prior to usage within LVM, a disk unit has to be initialized for LVM usage. You can do this with the **pvcreate** command which initializes the physical disk for LVM in the following steps:

1. Issue the following command to initialize the first LVM physical volume shown in Example 8-14.

*Example 8-14 Initialize the first disk for LVM*

---

```
pvcreate /dev/sdc
```

---

The **pvcreate** command will respond with the message shown in Figure 8-24.

Physical volume “/dev/sdc” successfully created

*Figure 8-24 Output of pvcreate command*

2. Now issue the **pvcreate** command for the next disk as shown in Example 8-15.

*Example 8-15 Initialize the second disk for LVM*

---

```
pvcreate /dev/sdd
```

---

### Creating the volume group

Now that there are two physical volumes ready for LVM usage the volume group itself can be created. Creation of the volume group will result in a device file being created that will then be used to create one or more logical volumes within the LVM. From a logical perspective, you can view the “volume group” as analogous to a physical disk device. Perform the following steps:

1. To create the volume group with two physical volumes, issue the following command as shown in Example 8-16.

*Example 8-16 Create the volume group*

---

```
vgcreate vol1 /dev/sdc /dev/sdd
```

---

In Example 8-16:

- vol1 is the name to assign to the new Volume Group.
- /dev/sdc and /dev/sdd are the two disk devices to initially create the group with.

The command will respond as shown in Figure 8-25.

```
Volume group "vol1" successfully created
```

Figure 8-25 Output from `vgcreate` command

This results in the device file `/dev/vol1` being created. Now that the volume group has been created, one or more logical volumes can be placed into it. Again, you can abstract out that the volume group is a disk and the volume itself is a disk partition.

**Note:** The signature information written to each disk in the LVM set (written by the `pvcreate` command) will reduce the amount of disk space available to the logical volume.

2. You can determine the maximum size available through the use of the `vgdisplay` command as shown in Example 8-17.

Example 8-17 Display volume group information

```
vgdisplay vol1
```

The command will result in outputting a variety of information concerning the volume group. The information that we are particularly interested in is the line shown in Figure 8-26.

```
Free PE / Size 50 / 200.00MB
```

Figure 8-26 Output from `vgdisplay` command

This indicates that there is a maximum of 200MB that can be used for the creation of a logical volume within this volume group. That figure will be used in the next step.

3. The following command will create the logical volume within the logical volume group shown in Example 8-18.

Example 8-18 Create the logical volume

```
lvcreate -L200 -ndata vol1
```

**Note:** The `-L` parameter is the parameter that is used to indicate the size (in megabytes) of the logical volume. The 200 is the maximum size as indicated in the earlier `vgdisplay` command.

The command will respond with the following (Figure 8-27):

```
Logical volume "data" created
```

Figure 8-27 Output from `lvcreate` command

This results in the creation of `/dev/vol1/data` which can be thought of as a “partition” on the `/dev/vol1` “disk”. The size of the volume will be 200MB (there is some overhead in the first physical volume for the LVM metadata). The name of the volume will be “data” and it will reside on the device ‘vol1’ (that is, `/dev/vol1/data`).



## Making the volume available to Linux

### Creating the file system

From a Linux perspective we now have a disk partition (the logical volume) that can be used to host a Linux file system using the following steps:

1. The following command will create an ext3 file system on the volume as shown in Example 8-19.

*Example 8-19 Create file system*

```
mke2fs -j /dev/vol1/data
```

Without any additional options, this command will create an ext3 file system on the volume using defaults for the blocking size and number of inodes to support.

### Accessing the file system

To access the file system, it has to be mounted somewhere within the Linux directory tree by performing the following steps:

1. Create a mount point (which is a directory in the tree) as shown in Example 8-20.

*Example 8-20 Create mount point*

```
mkdir /mnt/data
```

2. Gain access to the “device” (in our case a file system on a logical volume) by mounting the file system on top of the directory as shown in Example 8-21.

*Example 8-21 Mount the new file system*

```
mount /dev/vol1/data /mnt/data
```

3. Validate the size of the file system with the `df` command as shown in Example 8-22.

*Example 8-22 Display the disk utilization*

```
df -h
```

4. Issue the following command shown in Example 8-23 to look at the initial contents of the file-system.

*Example 8-23 List the directory contents*

```
ls -la /mnt/data
```

Observe that three directories are listed as shown in Figure 8-28.

```
.
..
lost+found
```

*Figure 8-28 Output of directory listing*

The `.` directory is a link to the directory itself. The `..` directory is a link to the parent directory (in this case `/mnt`). The `lost+found` directory is a directory created by the `mke2fs` command that is used by the file system integrity checks as a holding area for damaged files and inodes that have somehow become detached from their directory. Typically, the `lost+found` directory will

be empty, however, it is a good idea to check this directory after recovery from a system crash to ensure that no files have been placed into it.

## Increasing the size of LVM

Imagine the size of the file system contained within the logical volume is nearing capacity. With virtual disk, additional disk resources can be provided to the logical volume that can then be used to increase the size of the file system.

### i5/OS steps

#### Creating a new virtual disk

Use the CRTNWSSTG command to create a new virtual disk that will be used by the partition as shown in Example 8-24.

*Example 8-24 Create additional storage space*

---

```
CRTNWSSTG NWSSTG(Team10D) NWSSIZE(100) FORMAT(*OPEN)
```

---

#### Associating the new disk to the Network Server Description

Once the virtual disk has been created, it needs to be linked to the Network Server Description as shown in Example 8-25.

*Example 8-25 Link the storage space to the network server*

---

```
ADDNWSSTGL NWSSTG(Team10D) NWSD(Team10)
```

---

### Preparing the disk for LVM usage

#### Scanning the SCSI bus for the new disk

Issue the following command (in Linux) shown in Example 8-26 to scan the bus for the newly-added disk.

*Example 8-26 scan the SCSI bus*

---

```
echo "--" > /sys/class/scsi_host/host0/echo
```

---

#### Initializing the physical volume

Use the `pvcreate` command to initialize the physical volume for LVM usage as shown in Example 8-27.

*Example 8-27 Initialize the disk for LVM*

---

```
pvcreate /dev/sde
```

---

#### Adding the new physical volume to the logical volume

Now the physical volume needs to be added to the existing logical volume group. This is accomplished with the following command shown in Example 8-28.

*Example 8-28 Extend the size of the Logical Volume Group*

---

```
vgextend vol1 /dev/sde
```

---

The `vgextend` command will respond with the following shown in Figure 8-29.

```
Volume group "vol1" successfully extended
```

Figure 8-29 Output from `vgextend` command

### Increasing the size of the logical volume

Once the physical volume has been added to the logical volume group, its space can be used to increase the size of the logical volume as shown in Example 8-29.

Example 8-29 Extend the size of the Logical Volume

```
lvextend -L+100 /dev/vol1/data
```

The `lvextend` command will respond with the following shown in Figure 8-30.

```
Extending logical volume data to 300.00MB
Logical volume data successfully resized
```

Figure 8-30 Output from `lvextend` command

This command extends the size of the logical volume “data” by increasing its size by an additional 100MB.

### Making the additional space available

#### Unmounting the file system

Prior to increasing the size of the Linux file system to take advantage of the additional space in the logical volume group, the file system should be removed (from access) from the Linux directory tree. This is done with the `umount` command as shown in Example 8-30.

Example 8-30 Unmount the file system

```
umount /mnt/data
```

#### Checking the file system for errors

Prior to increasing the size of a file system, it is a good idea to interrogate the file system for any possible errors and to fix those prior to increasing the size. Use the following command to check the file system as shown in Example 8-31.

Example 8-31 Check file system

```
e2fsck -f /dev/vol1/data
```

#### Increasing the file system size

Now the file system can be extended to take advantage of the additional space available on the hosting logical volume as shown in Example 8-32.

Example 8-32 Resize the file system

```
resize2fs /dev/vol1/data
```

## Remounting the file system

Finally, the Linux file system can be remounted onto the Linux directory tree to make it available to the users again as shown in Example 8-33.

*Example 8-33 Mount the file system*

---

```
mount /dev/vol1/data /mnt/data
```

---

## Verifying the additional space availability

You can validate the availability of the additional space can be validated with the **df** command as shown in Example 8-34.

*Example 8-34 Check disk utilization*

---

```
df -h
```

---

## Virtual CD library

The virtual CD library feature of i5/OS provides support for CD images to reside in the Integrated File System (IFS) and be accessed as a CD through the use of a virtual CD device. While typically used to automate the program temporary fix (PTF) load process in i5/OS, the virtual CD library can also be used to install Linux as well as installation of additional packages after the installation has been completed.

This appendix describes the steps necessary to create a virtual CD environment within i5/OS. Information presented in this section includes:

- ▶ Creation of a virtual CD device
- ▶ Creation of the virtual CD library
- ▶ Loading images into the virtual CD library
- ▶ Working with the virtual CD library

## Virtual CD device creation

The CRTDEVOPT command is used to create the virtual CD device with the following steps:

1. Enter the command and press the **F4** key to display the Create Device Description screen shown in Figure A-1.

```

 Create Device Desc (Optical) (CRTDEVOPT)

Type choices, press Enter.

Device description > OPTVRT01 Name
Resource name > *VRT Name, *NONE, *VRT...
Device type *R SRCNAME *R SRCNAME, 4531, 632A,
632B...
Online at IPL *YES *YES, *NO
Message queue *SYSOPR Name, *SYSOPR
 Library Name, *LIBL, *CURLIB
Text 'description' 'Virtual Optical Device'

 Bottom
F3=Exit F4=Prompt F5=Refresh F10=Additional parameters F12=Cancel
F13=How to use this display F24=More keys
```

Figure A-1 Create Device Description (Optical)

The fields should be completed as follows:

- Device description: This is the name with which the virtual device will be referred. The name is user-generated, however, for virtual devices it is recommended that an indication (such as VRT) be included in the name to indicate that it is a virtual device.
- Resource name: For virtual optical devices, the value should be set to \*VRT.
- Device type: Leave the value set at \*R SRCNAME.
- Online at IPL: Typically, left at the default of \*YES.
- Message queue: Typically, left at the default of \*SYSOPR.
- Text description: A description of the purposes of the device.

2. Once all of the values have been entered, press **Enter** to create the device.

- Now that the device has been created, it should be varied on. Type the VRYCFG command and the Vary Configuration screen appears as shown in Figure A-2.

```

Vary Configuration (VRYCFG)

Type choices, press Enter.

Configuration object > OPTVRT01 Name, generic*, *ANYNW...
 + for more values

Type > *DEV *NWS, *NWI, *LIN, *CTL...
Status > *ON *ON, *OFF, *RESET...
Range *NET *NET, *OBJ
Reset *NO *NO, *YES
Submit multiple jobs *NO *NO, *YES
Job description QBATCH Name
 Library *LIBL Name, *LIBL

Bottom
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

```

Figure A-2 Vary Configuration

Complete the fields of the Vary Configuration command as follows:

- Configuration object: The name of the virtual optical device.
- Type: Set at \*DEV to indicate the vary configuration command will work with a device description
- Status: Set to \*ON to cause the command to vary on the optical device.

- Once the fields have been completed, press **Enter** to cause the device to be varied on.

## Virtual CD library creation

The virtual CD library is the i5/OS component that will hold the CD images. The following steps will create the image catalog:

- The CD images will reside in a directory in the IFS. The directory should exist prior to creation of the virtual CD library. Type the MKDIR command followed by the directory in the IFS. In Example A-1, the CD images are going to be contained in the /home/linuxclg directory.

*Example: A-1 Create directory for library contents*

---

```
mkdir '/home/linuxclg'
```

---

2. Type the CRTIMGCLG command and press the **F4** key to display the Create Image Catalog screen as shown in Figure A-3.

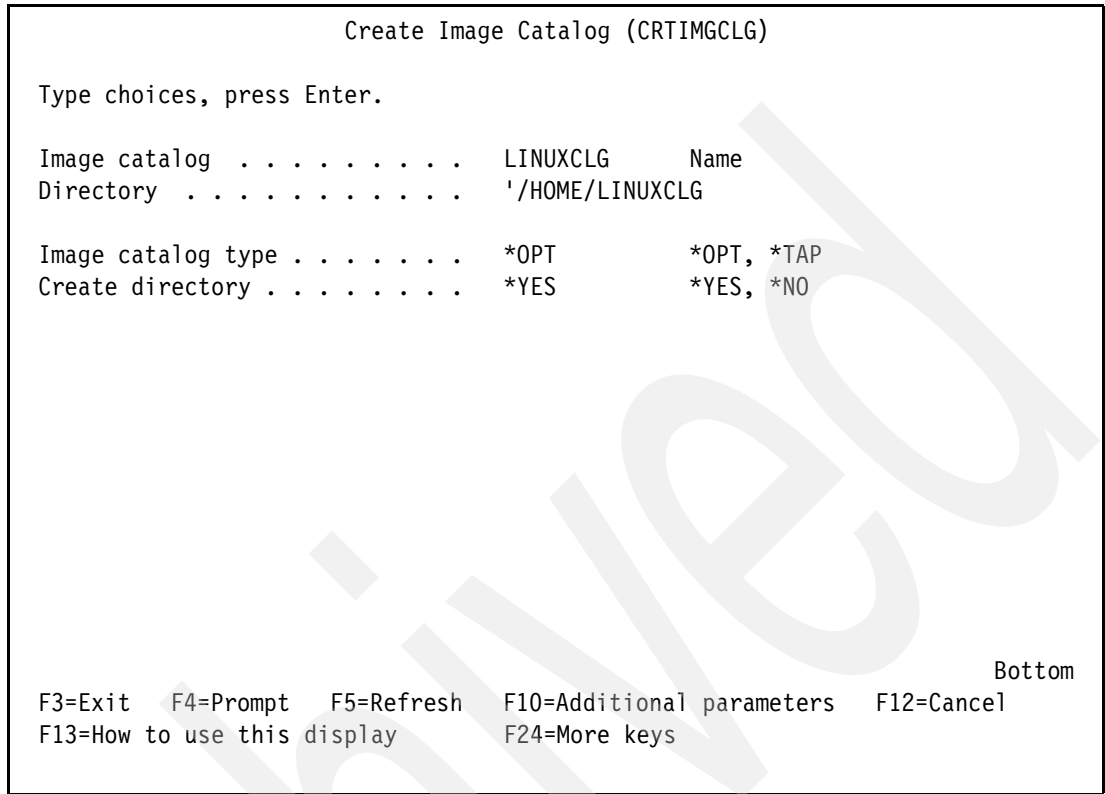


Figure A-3 Create Image Catalog

3. Complete the following fields:
  - **Image catalog:** User-defined name for the image catalog.
  - **Directory:** Location in the IFS where the CD images will be contained.
  - **Image catalog type:** The type of image catalog being created. Possible values:
    - **\*OPT:** An optical library
    - **\*TAP:** A tape library
 For the virtual optical library, the value should be set to \*OPT.
  - **Create directory:** Indicates whether or the directory should be created. If the value is set to \*YES then the directory will be created, if it does not already exist.
4. Once the values have been completed, press **Enter** to create the image catalog.



## Loading images into CD library

The Add Image Catalog Entry command is used to add CD images to the image catalog using the following steps:

1. Type the ADDIMGCLGE command and press **Enter**. The command screen is displayed as shown in Figure A-4.

```

 Add Image Catalog Entry (ADDIMGCLGE)

Type choices, press Enter.

Image catalog LINUXCLG Name
From optical device, or OPT01 Name
From image file

To image file SLES-10-CD1

Image catalog index *AVAIL 1-256, *AVAIL
Replace catalog entry *NO *NO, *YES, *INSERT

 Bottom
F3=Exit F4=Prompt F5=Refresh F10=Additional parameters F12=Cancel
F13=How to use this display F24=More keys
```

Figure A-4 Add Image Catalog Entry

2. Complete the following fields:
  - Image catalog: The name of the image catalog created earlier with the CRTIMGCLG command.
  - From optical device: Indicates the optical device that the CD should be read from.
  - From image file: In addition to creating CD images from a CD, images from an image file (that is, an ISO file) can also be added to an image catalog. This field should be left blank when the From optical device field has been specified.
  - To image file: Specifies the name of the image. This is the name that will be used in the image catalog to work with the CD image. A value of \*GEN will cause the label from the CD to be used.
  - Image catalog index: Typically, left at the default of \*AVAIL, the image catalog index is the index number that will be assigned to the image in the catalog. \*AVAIL will assign the next available index number.
  - Replace catalog entry: Typically, left at \*NO, indicates what action should take place if the specified catalog index already exists.
3. Once the fields have been completed, press **Enter** to cause the image to be added to the catalog.

**Note:** The ADDIMGCLGE entry command reads the CD (or image) and creates an image file. The command may take several minutes to complete based on the size of the CD and other system activity.

## Working with CD library

The ability to be access the contents of one of the images in the CD library requires that the appropriate image be mounted in the image catalog. The following steps will load the image catalog and describe the steps required to mount a volume in the library:

1. To use the images in the image catalog, the catalog must first be loaded on to the virtual optical device. Type the LODIMGCLG command and press **F4**. The Load Image Catalog command entry screen is displayed as shown in Figure A-5.

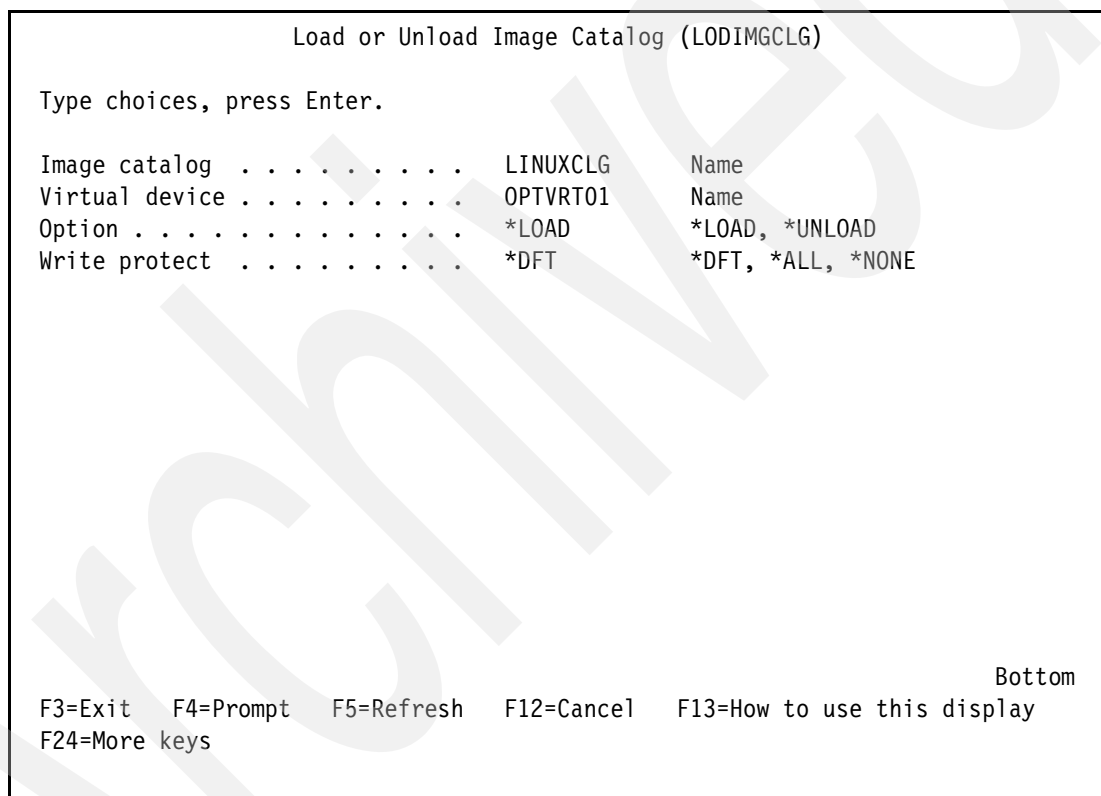


Figure A-5 Load or Unload Image Catalog

2. Once the fields have been completed, press **Enter** to cause the image catalog to be loaded onto the virtual CD device.

- The WRKIMGCLG command is used to work with the entries in the image catalog. Type the WRKIMGCLG command and press **Enter**. The Work with Image Catalogs menu is displayed as shown in Figure A-6.

```

Work with Image Catalogs
System: G10B8DAE
Type options, press Enter.
1=Create 2=Change 4=Delete 8=Load 9=Unload 10=Verify
12=Work with entries

Image
Opt Catalog Status Type Device Device Status
12 LINUXCLG Ready Optical OPTVRT01 Active

F3=Exit F5=Refresh F11=View descriptions F12=Cancel
Bottom

```

Figure A-6 Work with Image Catalogs

- The system can have multiple image catalogs. To work with the images in the catalog, press **Tab** to position the cursor in the Opt field of the appropriate image catalog and then type 12 and press **Enter**.
- The Work with Image Catalog Entries shown in Figure A-7 provides the interface to work with the entries in the image catalog. Access to the contents of a specific entry requires

that it be mounted. Press **Tab** to position the cursor at the Opt field of the desired catalog entry, then type 6 (Mount) and press **Enter**.

```
Work with Image Catalog Entries
 System: G10B8DAE
Catalog . . . : LINUXCLG Status . . . : Ready
Type : Optical Device . . . : OPTVRT01
Directory . . : /home/linuxclg

Type options, press Enter.
 1=Add 2=Change 4=Remove 6=Mount 8=Load 9=Unload
10=Initialize volume 12=Work with volume

Opt Index Status Image File Name
*AVAIL
 1 Loaded SLES-10-CD1
 2 Loaded SLES-10-CD2
 3 Loaded SLES-10-CD3
 4 Loaded SLES-10-CD4

 Bottom
F3=Exit F5=Refresh F6=Load image catalog F7=Verify image catalog
F8=Reorder by index F12=Cancel F24=More keys
Image catalog entry 1 loaded.
```

Figure A-7 Work with Image Catalog Entries

6. Now that the entry is mounted, it is accessible through the normal optical device paths of i5/OS (such as the /qopt path).

**Note:** Since the virtual CD device is seen as an optical device in i5/OS it is made available to Linux partitions that are using the i5/OS partition as a hosting partition for I/O resources.



## Setting up Linux on iSeries models

The following information is presented as reference for creating of the logical partitions (LPARs) for Linux on iSeries systems.

**Note:** For purpose of definition, an iSeries model refers to those models of iSeries that are iSTAR, sSTAR, or POWER4 based that are enabled to run Linux. Examples of these systems include model 800, 810, 825, 870, and 890.

## Freeing up system resources

The first step is to review the current allocation of the system resources and free up resources (memory and processor) from the OS/400 or i5/OS partition using the following steps.

On iSeries models, partition operations are accomplished through System Service Tools.

**Note:** The steps provided here assume that the logical partition being set up for Linux will be using all virtual I/O resources (that is, resources hosted by an i5/OS partition.)

1. From the System Service Tools (SST) menu (Figure B-1) type 5 (Work with system partitions) and press **Enter**.

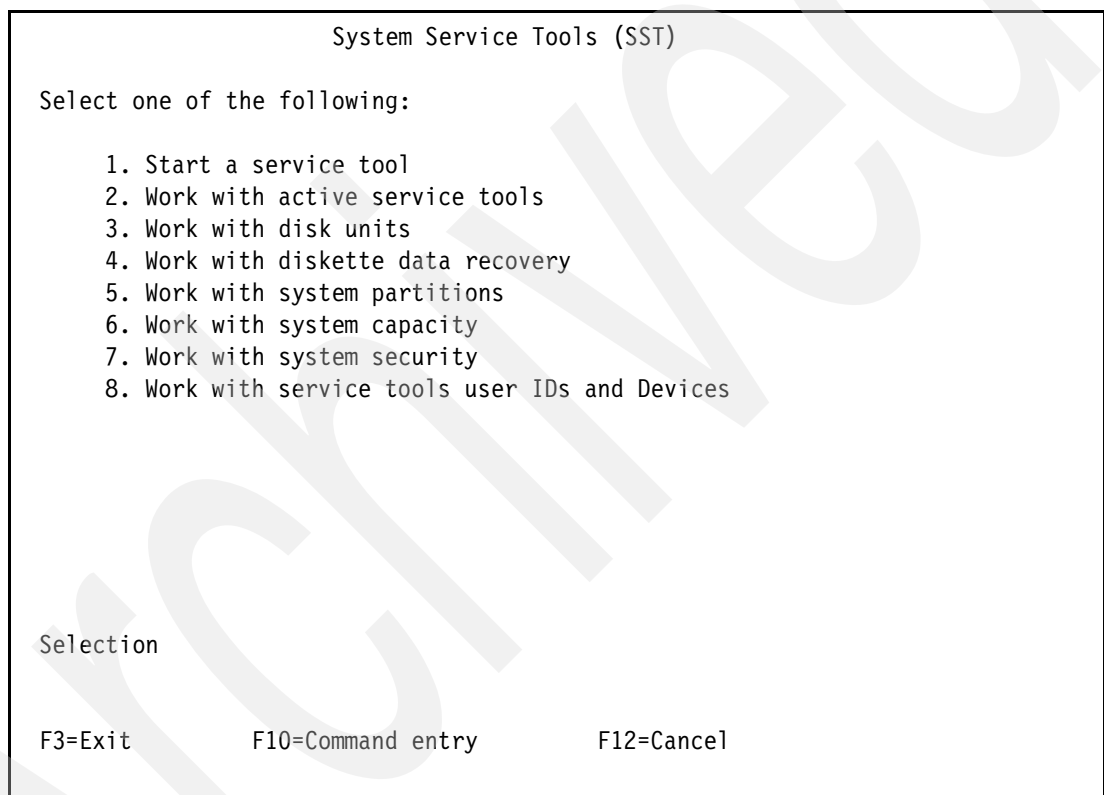


Figure B-1 System Service Tools menu

2. From the Work with System Partitions menu shown in Figure B-2, type 3 (Work with partition configuration) and press **Enter**.

```
Work with System Partitions
System: TECCONF4
Attention: Incorrect use of this utility can cause damage
to data in this system. See service documentation.

Number of partitions : 1
Partition manager release : V5R3M0 L000

Partition identifier : 0
Partition name : PRIMARY *

Select one of the following:

 1. Display partition information
 2. Work with partition status
 3. Work with partition configuration
 4. Recover configuration data
 5. Create a new partition

Selection

F3=Exit F10=IPL system to activate changes F12=Cancel
System IPL may be required to activate changes.
```

Figure B-2 Work with System Partitions

3. In the Work with Partition Configuration screen shown in Figure B-3, position the cursor in the Opt field of the partition that resources are going to be removed from, type 2 (Change partition processing resources) and press **Enter**.

```
Work with Partition Configuration System: TECCONF4

Type option, press Enter.
 1=Change partition name 2=Change partition processing resources
 3=Add I/O resources 4=Remove I/O resources
 5=Change bus ownership type 6=Select load source resource

Option Par ID Name
 0 PRIMARY <

< Indicates partition IPL may be required.
F3=Exit F5=Refresh F9=Work with shared processor pool
F10=Work with Virtual Ethernet F11=Work with partition status
F12=Cancel F23=More options
```

Figure B-3 Work with Partition Configuration



- In the Change Partition Processing Resources screen Figure B-4, press **F9** to display all of the fields.

```

Change Partition Processing Resources
System: TECCONF4
Type changes, press Enter.

Partition identifier and name : 0 PRIMARY
Current / available number of processors . . . : 6 / 0
New number of processors 6
Minimum / maximum number of processors 1 / 6
Use shared processor pool 1 1=Yes, 2=No
 New shared processor pool units 6 . 00
 Minimum / maximum processor pool units 1 . 00 / 6 . 00

Current / available size of main storage (MB) : 7168 / 1024
New size of main storage (MB) 7168
Minimum / maximum size of main storage (MB) . . 1024 / 8192

Current / available interactive feature . . . : 100 / 0 %
New interactive feature 100 %
Minimum / maximum interactive feature 0 / 100 %

F3=Exit F9=Exclude limits F10=Work with shared processor pool
F11=Display partition processing configuration F12=Cancel

```

Figure B-4 Change Partition Processing Resources

- The Change Partition Processing Resources screen shows the current resource allocation for the partition. For creation of the Linux partition, processor resources and memory resources should be made available.

To reduce the processor resources, position the cursor in the New shared processor pool units and type in a new, smaller value. Keep in mind that the value cannot be reduced below the value shown in the Minimum processor pool units field without changing that value as well.

- To reduce the memory resources, position the cursor in the New size of main storage (MB) field and type in a new, smaller value. Keep in mind that the value cannot be reduced below the value shown in the Minimum size of main storage (MB) field without changing that value as well.

Once you have changed the values, press **Enter**. The Confirm Changed Partition screen shown in Figure B-5 appears. Press **Enter** to confirm the change.

```

Confirm Changed Partition
System: TECCONF4
Verify information, press Enter.

Partition identifier and name : 0 PRIMARY
Number of partition processors : 6
Minimum / maximum number of processors : 1 / 6
Use shared processor pool : Yes
 Shared processor pool units : 5.00
 Minimum / maximum processor pool units : 1.00 / 6.00
Size of partition main storage (MB) : 6096
Minimum / maximum size of main storage (MB) : 1024 / 8192
Partition interactive feature : 100 %
Minimum / maximum interactive feature : 0 / 100 %
OptiConnect / Ethernet / HSL OptiConnect : No / Yes / Yes

I/O Resource
Description Type-Model Serial Part
System Bus 1 25CA- 0F-3122008 97P3203
Combined Function IOP * 2844-001 53-3081139 0000097P2694
Communications IOA 287F-001 0F-3122008 97P3203

F10=Display logical address F12=Cancel
More...

```

Figure B-5 Confirm Changed Partition

7. The reduction in processor resources will cause the interactive feature available to the partition to be reduced. The Change Partition Processing Resources screen will be redisplayed with the reduced value for the interactive feature. Press **Enter** to confirm the partition resource change.
8. At this point, the Confirm Changed Partition screen is again displayed. Press **Enter** to confirm.

9. Now that the partition resources have been changed, the Work with Partition Configuration screen shown in Figure B-6 is redisplayed.

```
Work with Partition Configuration System: TECCONF4
Type option, press Enter.
 1=Change partition name 2=Change partition processing resources
 3=Add I/O resources 4=Remove I/O resources
 5=Change bus ownership type 6=Select load source resource

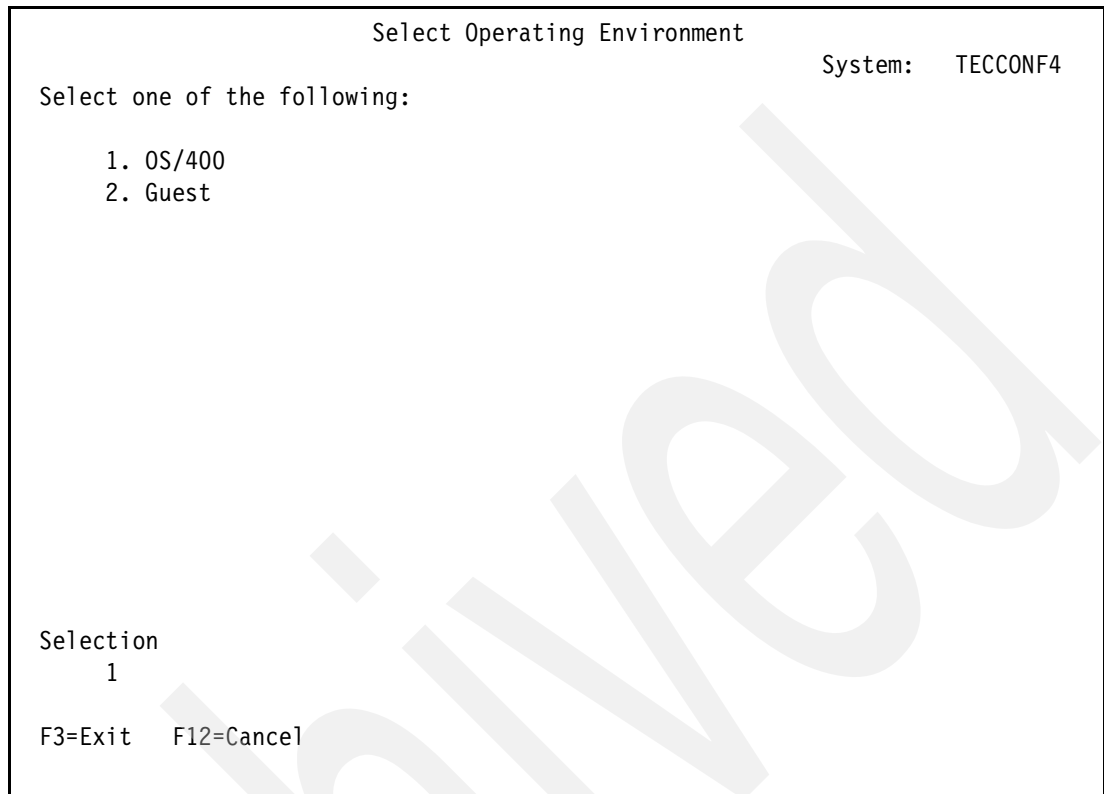
Option Par ID Name
 0 PRIMARY <

< Indicates partition IPL may be required.
F3=Exit F5=Refresh F9=Work with shared processor pool
F10=Work with Virtual Ethernet F11=Work with partition status
F12=Cancel F23=More options
```

Figure B-6 Work with Partition Configuration

10. The < symbol indicates that the resource changes will require an initial program load (IPL) to take affect. However, the logical partition for Linux will be created prior to the IPL. Press **F3**.
11. In the 'Work with System Partitions' screen, type 5 (Create a new partition) and press **Enter**.

12. When the 'Select Operating Environment' menu is displayed as shown in Figure B-7, type 2 (Guest) and press **Enter** to indicate that a guest partition (for Linux) will be created.



*Figure B-7 Select Operating Environment*

13. When the Create New Partition screen is displayed as shown in Figure B-8, press **F9** to show all the fields of the partition configuration.

```

 Create New Partition
 System: TECCONF4
Complete blanks, press Enter.

Partition identifier and name 1
Number of available system processors : 0
Number of partition processors
Minimum / maximum number of processors /
Use shared processor pool 2 1=Yes, 2=No

Size of available system main storage (MB) . . : 2096
Size of partition main storage (MB)
Minimum / maximum size of main storage (MB) . . /

F3=Exit F9=Exclude limits F10=Work with shared processor pool
F11=Display partition processing configuration F12=Cancel

```

Figure B-8 Create New Partition

14. Complete the following page:

- **Partition identifier and name:** Leave the partition identifier as generated by the system. Provide a name for the partition.
- **Number of partition processors:** This is the number of processors you want the workload to be spread across. Typically, for Linux partitions, this value will be set at 1.
- **Minimum / maximum number of processors:** This is the minimum and maximum number of processors that can be assigned to the partition without requiring an IPL of the system.
- **Use shared processor pool:** This indicates if the partition can share processor resources with other partitions. Considering that most Linux partitions use less than a full processor, typically this value is set at 1 (yes).
- **Size of partition main storage (MB):** The amount of memory to be allocated to the Linux partition.
- **Minimum / maximum size of main storage (MB):** Indicates the minimum and maximum amount of memory that can be assigned to the partition without requiring an IPL of the system.

When the fields have been completed, press **Enter**. If the Use shared processor pool value was set to 1 (yes), then two additional fields are displayed to indicate the amount of processor resources to assign to the partition:

- **Shared processor pool units:** Indicates the amount of processor resources to assign to the partition. The value must be equal to or less than the number of partition processors defined for the partition.

- **Minimum / maximum processor pool units:** Indicates the minimum and maximum amount of processor units that can be assigned to the partition without requiring an IPL of the system.

Figure B-9 shows an example of a completed partition configuration. Press **Enter** to create the partition.

```

 Create New Partition
 System: TECCONF4
Complete blanks, press Enter.

Partition identifier and name 1 LINUX

Number of available system processors : 0
Number of partition processors 1
Minimum / maximum number of processors 0 / 1
Use shared processor pool 1 1=Yes, 2=No
 Shared processor pool units 0 . 5
 Minimum / maximum processor pool units 0 . 0 / 1 . 0

Size of available system main storage (MB) . . . : 2096
Size of partition main storage (MB) 1024
Minimum / maximum size of main storage (MB) . . 0 / 2048

F3=Exit F9=Exclude limits F10=Work with shared processor pool
F11=Display partition processing configuration F12=Cancel
Enter value for shared processor pool units.

```

Figure B-9 Create New Partition - Example Values

15. In the Select Communication Options screen shown in Figure B-10, type a 1 (yes) under the virtual ethernet adapter to assign to the partition. Typically, this will be the same virtual

LAN that the i5/OS partition has a virtual LAN connection on. Once you have selected the virtual LAN adapters, press **Enter**.

```
 Select Communication Options
 System: TECCONF4
Partition identifier : 1
Partition name : LINUX

Type changes, press Enter.
1=Yes 2=No

-----Virtual Ethernet Identifiers-----
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
1 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2

F3=Exit F11=Display communication options F12=Cancel
```

Figure B-10 Select Communication Options

16. The Confirm New Partition screen shown in Figure B-11 is displayed providing a summary of the configuration settings for the partition. Press **Enter** to confirm the partition creation.

```

Confirm New Partition
System: TECCONF4

Verify information, press Enter.

Partition identifier and name : 1 LINUX
Number of partition processors : 1
Minimum / maximum number of processors : 0 / 1
Use shared processor pool : Yes
 Shared processor pool units : 0.50
 Minimum / maximum processor pool units : 0.00 / 1.00
Size of partition main storage (MB) : 1024 (21)
Minimum / maximum size of main storage (MB) : 0 / 2048
Virtual Ethernet : Yes

I/O Resource
Description Type-Model Serial Part
 Number Number

F3=Exit F9=Select host partition F10=Display logical address
F11=Add I/O resources F12=Cancel

```

Figure B-11 Confirm New Partition

**Note:** Once the partition has been created, an IPL of the system is required for the partition to be available for use.

The Network Server Description, Network Server Storage Space, and Virtual Networking support on the pre-System i5 models is the same as the System i5™ models. Refer to the information presented earlier in this book about those topics.



## Performance considerations

This appendix presents a number of considerations regarding implementation of Linux-based solutions on the IBM System i platform. The information presented in this appendix is not intended to be the definitive word on performance and will not present specific performance characteristics of a given workload. As with anything involving performance, it is recommended that you plan on prototyping your solutions and be prepared to make changes as you go forward.

The information presented in this appendix can be broken down into three broad categories:

- ▶ **Virtual I/O support:** One of the key strengths to implementing Linux-based solutions on the System i platform is the ability to have the I/O (disk) hosted by an IBM i5/OS partition and, thereby, extend the benefits of single-level store to the Linux operating system. We will discuss considerations regarding performance as well as changes that have occurred in recent versions of i5/OS that you should be aware of.
- ▶ **Virtual network support:** Another advantage of implementing Linux-based solutions on the System i platform is the ability to build virtual Ethernet networks (LANs) inside the managed system (that is, with no physical hardware). We will discuss the importance of the frame sizes of the network adapters as well as when it may be advantageous to implement multiple virtual LANs.
- ▶ **Processor considerations:** The System i platform supports the ability to share processors between multiple logical partitions (LPARs) and to have the firmware (hypervisor) balance workload across the available processors through the uncapped partition support. This section will discuss performance considerations of fractional processor support as well as virtual processors.

## Virtual I/O support

With Virtual I/O support, a Linux partition can be implemented with no physical I/O resources, all of the I/O resources are owned by a host partition (i5/OS) and the Linux partition is provided access to those resources through a virtual I/O driver running in i5/OS. Virtual I/O can, given the appropriate workload, provide significant performance advantages over the same workload implemented on a separate server. It all depends on the workload being implemented as well as proper configuration of both the Linux and host (i5/OS) partitions.

## Multi-path I/O

One of the advantages of a hosted Linux solutions is that multi-path I/O support can be provided to the underlying storage architecture (that is, single-level store) without any driver considerations within Linux.

*Multi-path I/O* provides multiple software threads from the virtual I/O driver in i5/OS to the underlying single-level store architecture. The more threads that are reading or writing data from or to the storage architecture, the faster the I/O throughput is conceptually possible.

At this point, it is important to make a distinction between the System i5 platform and previous POWER iSeries models as well as different versions of i5/OS regarding multi-path I/O (shown in Table C-1) before making any recommendations.

Table C-1 System i models and i5/OS versions for multi-path I/O support

|                         | iSeries models | System i5 |
|-------------------------|----------------|-----------|
| V5R3M0 - Multi-path I/O | Yes            | No        |
| V5R3M5 - Multi-path I/O | Yes            | Partial   |
| V5R4M0 - Multi-path I/O | Yes            | Yes       |

**Note:** In the above table, iSeries models refers to iSTAR, sSTAR, and Power4 systems that are Linux enabled (example, 270, 810, 825, 870, 890.). System i5 refers to Power5 and Power5+ systems (that is, 520, 550, 570, 595).

If your Linux solution is implemented on an iSeries model that is a predecessor to the Power5 and Power5+ based i5 models then there is nothing you need to do to obtain multi-path I/O support. However, if your Linux solution is implemented on the System i5 platform, then you need to be aware of the support for multi-path I/O. If you have a heavy I/O workload (for example, file serving) implemented on a System i5 model and you are currently running at V5R3M0, then you should consider upgrading to V5R4M0 to provide multi-path I/O support. If such an upgrade is not currently possible, then you should consider changing the configuration of the Linux partition to have multiple network Server Descriptions with multiple virtual Small Computer System Interface (SCSI) adapters. This type of configuration will simulate, to a degree, the multi-path I/O provided in V5R4.

**Note:** It is recommended that for systems implementing Linux with heavy I/O requirements, the i5/OS partition that is hosting the I/O for Linux should be upgraded to V5R4M0 to achieve the best possible virtual I/O performance.

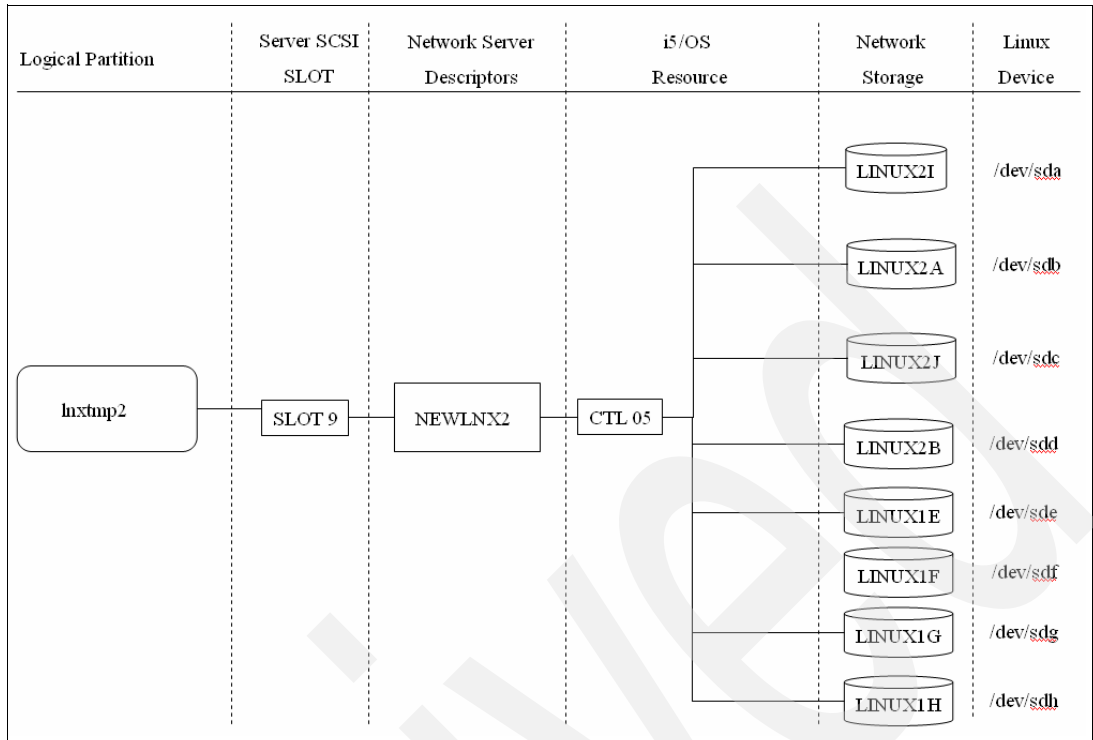


Figure C-1 Virtual I/O example

The configuration illustrated in Figure C-1 provides multiple disk drives (storage spaces) to Linux, however, they are going through a single SCSI adapter being hosted by a single Network Server. On a System i5 model if the I/O for this partition was hosted by a V5R3M0 i5/OS partition, then the I/O to all of the storage spaces would be single threaded. The configuration can be changed to what is illustrated in Figure C-2.

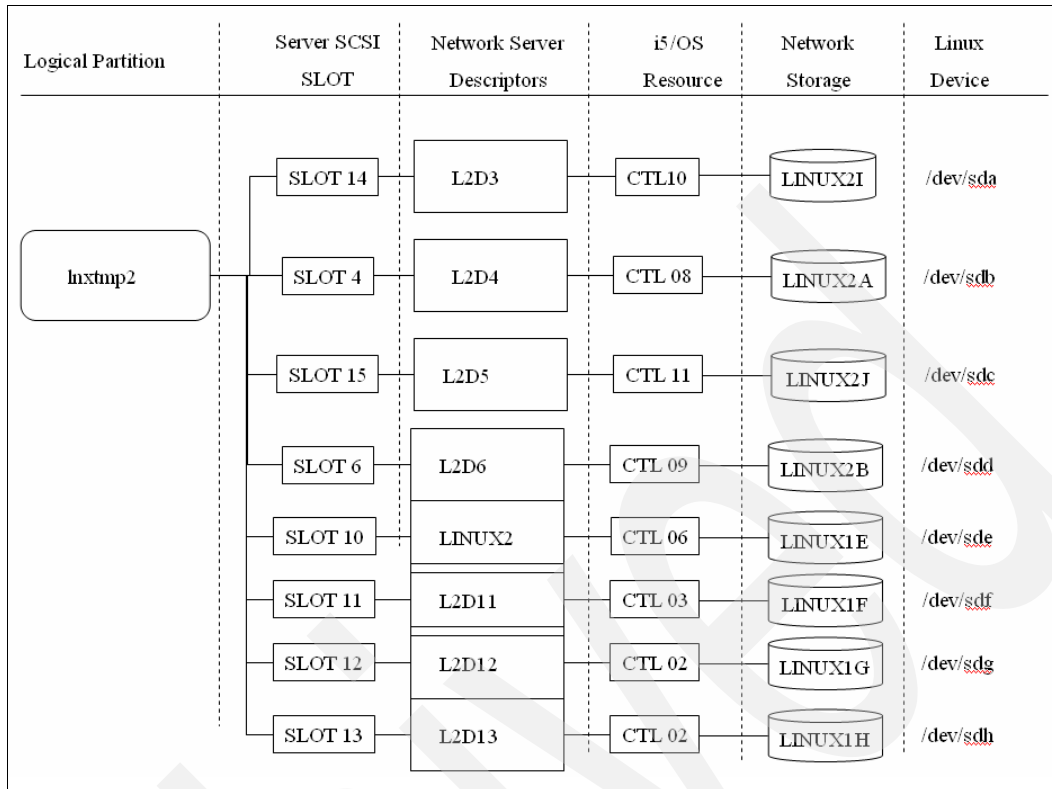


Figure C-2 Virtual I/O example

This configuration provides multiple client/server SCSI pairs with multiple network servers and associated network storage spaces to provide multiple virtual I/O connections.

**Note:** If a configuration is established that uses multiple Network Server Descriptions (NWSDs), you may want to configure all but the first NWSD with the Restrict Device Resources (RSTDDEVRS) value set to \*ALL. This parameter is used to restrict access through the Network Server to I/O devices in i5/OS. If this value is not set to \*ALL then any CD/DVD/TAPE in i5/OS will have multiple device handles created in the Linux partition since they will be presented to the partition on multiple SCSI chains.

Even with a single storage space, V5R3M5 and V5R4M0 hosted I/O provides multi-path I/O.

## Memory pool

Keep in mind that virtual I/O requests from a Linux partition rely on a virtual I/O driver in the i5/OS partition to actually process the request. This is one of the aspects of hosting Linux workloads that can confuse the performance and resource requirements discussions. Heavy I/O workloads can increase the memory requirements on the partition that is hosting the I/O (that is, the i5/OS partition).

What should be understood from a configuration viewpoint is that the virtual I/O driver in i5/OS runs out of the memory pools of the operating system. Reviewing the system status (WRKSYSSTS) of the hosting i5/OS partition can provide insight into the performance

characteristics of the workload. As an example, in Figure C-3, the system status represents I/O bound Linux workload that is being hosted on an i5/OS partition.

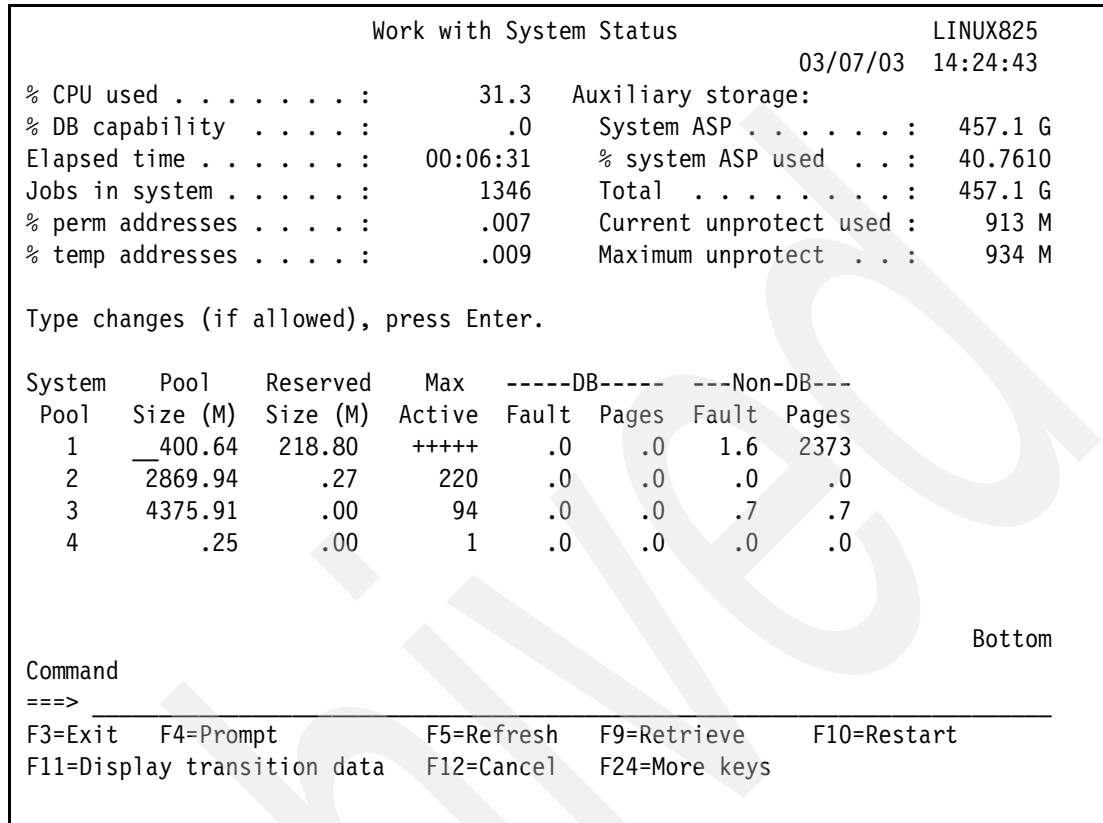


Figure C-3 Work with system status

The information of interest with regards to Virtual I/O is the Non-DB pages. The number of pages shown may be an indication of the number of I/O requests that are waiting to be processed by the virtual I/O driver in i5/OS. Obviously the larger the number the more requests are waiting and the slower the performance of the partitions that have hosted I/O could be. Adding more memory to the memory pool can have an immediate and substantial impact on the processing of the I/O requests.

Take the following factors into consideration:

- ▶ The memory allocated can be changed over time by the system if the QPRFADJ system value is set.
- ▶ The memory pool used by the virtual I/O component was changed in V5R4. Prior to V5R4 the memory pool used was \*MACHINE. In V5R4 the memory pool used is \*BASE.

## Virtual network

*Virtual network* provides the ability to build virtual Ethernet network segments (LANs) inside of the System i platform without requiring any physical hardware. Virtual network is a commonly used feature when establishing Linux partitions and can help provide for robust configurations. There are three performance considerations with regards to virtual network:

- ▶ Frame size
- ▶ Use of multiple network segments to split internal and external traffic
- ▶ Proxy ARP versus Network Address Translation (NAT)

## Frame size

The frame size indicates the size (in bytes) of the network packet that can be supported by the network adapter. Typical frame sizes that will be encountered when implementing Linux on the System i platform include those shown in Table C-2.

Table C-2 Typical frame sizes per network speed and operating system

| Network speed | i5/OS | Linux |
|---------------|-------|-------|
| 100MB         | 1496  | 1500  |
| 1GB           | 8996  | 9000  |

It is important, from a performance perspective, to ensure that the frame size of all of the adapters matches the lowest common denominator. As an example, Figure C-4 shows a configuration with a virtual LAN that has connections made available (or transported) across a real or physical network segments.




Figure C-4 Virtual network example

In this example, the network for the Linux partition (virtual LAN 1) is made available to the external network through a Proxy ARP configuration on the physical interface allocated to the i5/OS partition. Now, assume that the physical interface (CMN43) in i5/OS is configured with a frame size of 1496 and that the virtual LAN interfaces are configured with a frame size of 8996/9000. In this case, every packet intended for the virtual LAN that goes across the physical interface will first have to be fragmented by the i5/OS network interface to allow it to fit in the smaller network packet. Configuring the frame sizes of the virtual LAN interfaces to the smaller frame size (1496/1500) will negate this fragmentation of the network packets and can actually improve performance on heavy network workloads.

## Multiple network segments to separate internal/external traffic

When a workload that has a requirement for substantial network traffic between the Linux and i5/OS partitions (for example, Linux-based Apache Web application accessing DB2/400 data through ODBC) is being implemented, you can increase performance by implementing separate virtual LAN just for that traffic. In the example shown in Figure C-4, if the Linux partition has to access data from the i5/OS partition, it is going to do that across virtual LAN 1 and if the adapters on that virtual LAN have been configured for a slower physical network

segment, then the partition to partition network communications will also be at that slower speed. Refer to the example shown in Figure C-5.




Figure C-5 Virtual network example - Proxy ARP

The second virtual LAN (VLAN 2) can be configured at the faster 1GB frame sizes (8996/9000) and thereby provide faster data transfer between the two partitions.

**Note:** An alternative to the above approach would be the allocation of a second physical network adapter to the Linux partition (rather than using Proxy ARP to forward traffic to the Linux partition.) In this case, a virtual LAN would still be established between the i5/OS and Linux partitions for fast intra-partition communication.

## Proxy ARP versus Network Address Translation

It is possible to make partitions with virtual LAN connections visible to an external network. Two of the more popular approaches are Proxy ARP and Network Address Translation. In both of these methods, an i5/OS partition is allocated connections on both the virtual LAN as well as the external network. With Proxy ARP (refer to the examples shown in Figure C-5), the physical network that the i5/OS partition is on is segmented such that i5/OS essentially becomes the router for a set of addresses in the overall network. The addresses in the smaller segment are then assigned to the partitions on the virtual LAN and the i5/OS network adapter responds to those addresses and broadcasts the resulting traffic on the virtual LAN.

The important thing to note here is that the i5/OS partition does not process the network traffic in any way, it simply re-broadcasts the traffic on the virtual LAN.

With Network Address Translation, a private network segment is established on the virtual LAN and mappings are established that take addresses on the real network and translate them to the private addresses as shown in Figure C-6.

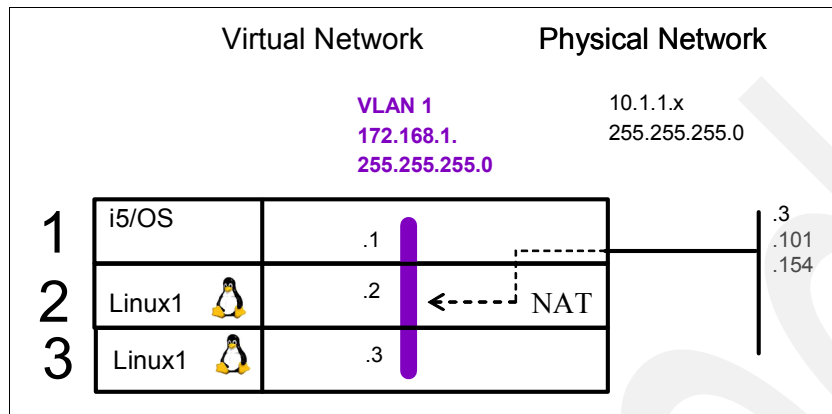


Figure C-6 Virtual network example - Network Address Translation

The addresses on the physical network are allocated as TCP/IP interfaces against the i5/OS adapter. When traffic is seen on one of the addresses, it is routed to the i5/OS partition. At that point, the network stack in i5/OS goes into the packet and changes the IP address in the header to the IP address of the partition on the virtual LAN and then broadcasts the traffic on the virtual LAN. Similarly, as traffic is routed out of the virtual LAN, the network stack in i5/OS will again go into the packet and change the IP address in the packet header to the external address of the partition. For partitions that will see heavy network traffic, the use of NAT can cause performance degradation since the partition relies on i5/OS to translate the address in each packet. The use of Proxy ARP over NAT has been shown to provide better network performance for certain workloads.

## Processor considerations

Most Linux-based workloads implemented on the System i platform leverage support of fractional processor allocation and uncapped partition processing. Fractional processor allocation allows for multiple logical partitions (and their corresponding operating systems) to run on a single POWER processor. Uncapped partition processing allows the hypervisor of the managed system to monitor the performance characteristics of the logical partitions to make additional processor resources available when needed and available.

## Virtual processors

Probably the biggest concern, from a performance perspective, to be aware of regarding processor allocation is the concept of virtual processors. Virtual processors are a representation of a processor thread as presented to the logical partition. It is possible to allocate multiple virtual processors to a logical partition even when less than a full processor has been allocated to that partition. The only rule for the allocation of virtual processors is that each virtual processor has to have at least 1/10th of a physical processor to run against. Therefore, as an example, if 0.40 processor units are allocated to a partition, that processor allocation can be spread across from 1 to 4 virtual processors. The performance consideration here is to understand how that processor allocation is spread across the time slice of the overall system.

The virtual processor setting defines the way that a partition's processor entitlement may be spread concurrently over physical processors. The key here is that the spread is concurrent.



The number of virtual processors is what the operating system thinks it has for physical processors. The Hypervisor dispatches virtual processors onto physical processors.

When considering the effect of the number of virtual processors on performance, it is important to have at least a basic understanding of the dispatch cycle that the Hypervisor employs for allocation of processor units to the logical partitions. The dispatch cycle is 10 milliseconds (ms) and the processor capacity allocated to a partition is taken from that 10 millisecond timing cycle. As an example, if a partition is allocated 0.40 processor units, then it is entitled to 4 milliseconds in the 10 millisecond timing cycle. The number of virtual processors does not change the timing cycle (that is, it is still 10 milliseconds), however, it can have an effect on how quickly the partition can consume its entitled capacity. If the 0.40 processor units is allocated across a single processor, then the time slice allocated to the partition will take 40% of the timing cycle as shown in Figure C-7.




Figure C-7 Processor allocation example

The same partition with the same entitled capacity (0.40 processor units) that has an allocation of 4 virtual processors would have its processor capacity allocated across the 4 processors and that allocation could be spread concurrently as shown in Figure C-8.

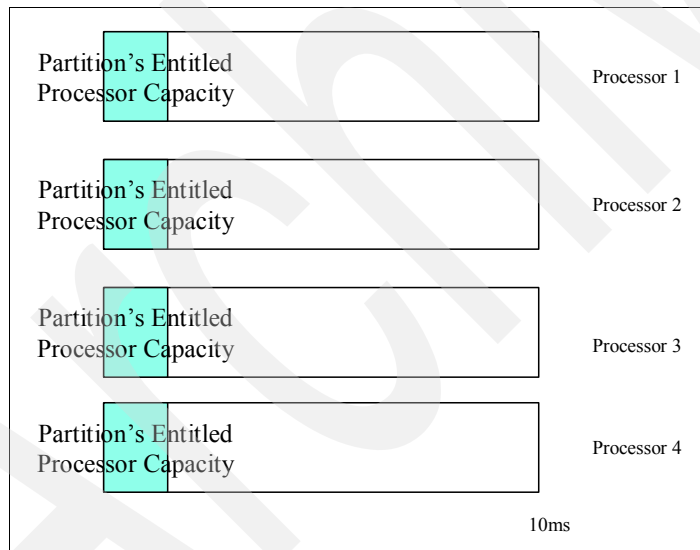


Figure C-8 Processor allocation example

The impact of spreading the processor allocation across multiple virtual processors can be increased gaps in the processing of the partition that have been known to result in noticeable lags in processing (such as terminal access that temporarily halts and then restarts). As a general rule, unless there is a specific workload requirement for a large number of processor threads, limiting the number of virtual processors to the smallest number that can accommodate the processor allocation will minimize any processing lag seen by the partition.

## Uncapped partitions

Note that the number of virtual processors that the processor allocation is spread across will have an impact on the maximum amount of additional processor units that can be made

available to an uncapped partition. With uncapped partitions, if a partition has used its allocation of processing units during the current processor dispatch cycle and additional processor resources are required (and available) the Hypervisor can dispatch additional portions of the cycle up to the number of virtual processors allocated to the partition. Table C-3 provides an example of the maximum amount of processor units that could be made available to an uncapped partition.

Table C-3 Entitled capacity for a number of virtual processors

| Entitled capacity             | 2 Virtual processors | 4 Virtual processors |
|-------------------------------|----------------------|----------------------|
| 1.5 processor units allocated | 2.0                  | 4.0                  |

**Note:** It is important to remember that the uncapped feature of a partition will only affect a partition's processor resources for a given timing cycle (which is 10 milliseconds). At the start of each timing cycle, a partition's processor resources is set to its Entitled Capacity (which is the amount of processor resources allocated to the partition when the partition was started - or changed through Dynamic LPAR).

## Dedicated processors

In addition to shared processors, a partition can also be configured to use dedicated processors. Allocation of a dedicated processor to a partition will provide for the most efficient usage of the processor as the hypervisor will not need to perform task switching for the processor and will be able to ensure memory affinity for the processor. One trade off with use of dedicated processors is that the partition using the dedicated processor must be a capped partition - that is the partition cannot participate in the hypervisor's load balancing of processor resources across a number of partitions. This means that if the operating system is only using a fraction of the dedicated processor allocated to it, the remaining processor allocation is wasted since the hypervisor cannot allocate it during the time slice to other partitions that may need additional processor resources.

**Note:** A partition cannot be configured with both shared and dedicated processors.

## Summary

A number of factors can affect the overall performance of Linux-based workloads hosted on the System i platform. A good understanding of both the characteristics of the workload as well as the configuration aspects will go a long way to obtain the best performance for your workload. Table C-4 provides a high-level summary of performance aspects that you should consider for different types of workloads.

Table C-4 Performance considerations summary

| Workload  | Performance consideration                                                                                                                                                                                                                                                |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Heavy I/O | <ul style="list-style-type: none"> <li>▶ Consider the amount of memory allocated to the memory pool in the hosting i5/OS partition.</li> <li>▶ Consider moving to V5R4M0 or implementing multiple network servers and multiple virtual SCSI adapter pairings.</li> </ul> |

| Workload                              | Performance consideration                                                                                                                                                                                   |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Heavy network                         | <ul style="list-style-type: none"><li>▶ Ensure that the frame size for all adapters is the frame size of the slowest adapter.</li></ul>                                                                     |
| Heavy intra-partition network traffic | <ul style="list-style-type: none"><li>▶ Consider implementing an additional virtual LAN to be used for intra-partition communication. Ensure that the adapters on this LAN are at gigabit speeds.</li></ul> |

Archived

Archived



## Troubleshooting

A number of common issues can occur when first establishing the Linux on IBM System i platform. Most of these problems tend to stem from configuration mis-matches between the definition of the logical partition and the IBM i5/OS components. This chapter presents some of the more common problems, what the typical symptoms of those problems are and how to resolve the issue.

## Virtual I/O - network server configuration issues

The following problems generally occur when there is an error in the configuration of the Network Server in the i5/OS partition that is hosting the I/O resources for the Linux partition.

**Note:** When a vary on a Network Server fails, you can obtain additional detailed information concerning the failure (including error codes) by performing the following steps:

1. Position the cursor on the failure message (you may need to do a DSPMSG QSYSOPR) to display the failure message.
2. Press the **F1** key.


## Attempt to vary on fails - partition not found

### Symptom

An attempt to vary on the Network Server fails with an indication that the partition cannot be found.

### Problem determination

One cause of a vary on failure of a Network Server Description can be the inability to “find” the partition. Review the job lot for more information as shown in Figure D-1.



```
Session A - [24 x 80]
File Edit View Communication Actions Window Help
Additional Message Information
Message ID : CPDB1AD
Date sent : 08/23/04 Time sent : 13:02:58
Message : Network server LINUX1 vary on failed.
Cause : The network server failed to vary on with reason code
00000002. Reason codes and their meanings follow:
--0001: Storage space not found.
--0002: Partition not found.
--0003: Partition not correct type.
--0004: Server already varied on for partition.
--0005: Partition already active.
--0006: A file system error occurred accessing the IPL stream file.
--0007: IPL source not valid.
--0008: IPL of partition failed.
--0009: A host partition is not configured for this partition.
More...
Press Enter to continue.
F1=Help F3=Exit F6=Print F9=Display message details F12=Cancel
F21=Select assistance level
MB a 01/001
```

Figure D-1 Troubleshooting - partition not found

### Corrective action

One cause of this error can be a mismatch between the partition name in the Network Server Description and the partition name in the logical partition definition. These names must match and are case sensitive. As Figure D-2 shows a mismatch of the case sensitivity of the names will cause the partition not found error to occur.

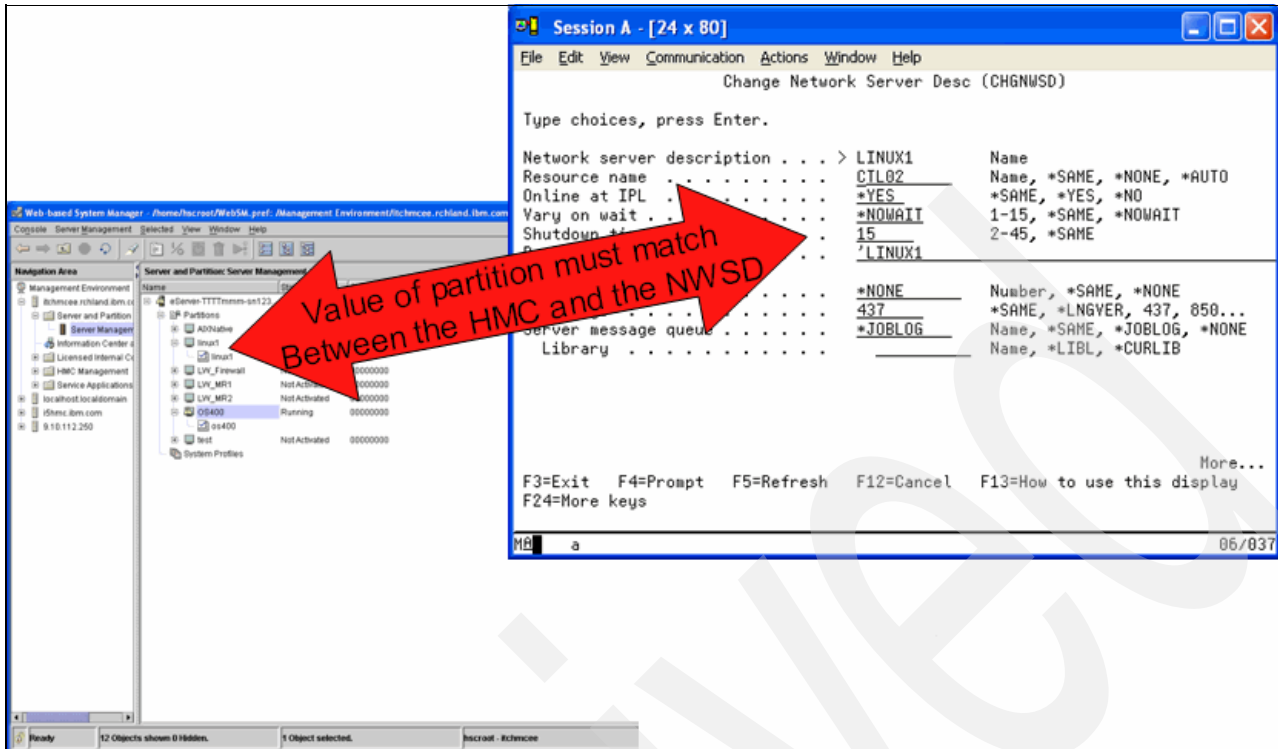


Figure D-2 Troubleshooting - partition not found


The corrective action is to either change the name of the partition in the Network Server Description or change the name of the partition in the partition definition.

**Note:** If you change the name of the partition in the partition definition, then you will need to re-activate the partition from the HMC to register the partition change with Open Firmware.

## Attempt to vary on NWSD fails - SCSI server resource definition problem

### Symptom

An attempt to vary on the Network Server fails and an error code 12 - Invalid resource name is returned as shown in Figure D-3.



```
Session A - [24 x 80]
File Edit View Communication Actions Window Help
Additional Message Information
Message ID : CP0B1AD
Date sent : 08/23/04 Time sent : 10:45:38
--0010: A network server storage space linked to this network server is
damaged.
--0012: Invalid resource name.
--0013: Resource name does not match partition.
--0014: Unable to determine partition for resource name.
--0015: Unknown error occurred. See previous messages in the job log for
details.
Recovery . . . : Correct the error. Try the request again.

Press Enter to continue.

F1=Help F3=Exit F6=Print F9=Display message details F12=Cancel
F21=Select assistance level

01/001
```

Figure D-3 Troubleshooting - invalid resource name

### Problem determination

This problem is typically caused by a value in the resource field of the network server that does not point to the correct virtual Server Small Computer System Interface (SCSI) adapter in the i5/OS partition.

### Corrective action

For most Linux implementations, a single virtual SCSI Server/Client pairing is configured between the Linux partition and the i5/OS partition. For those implementations a setting of \*AUTO can be used in the resource field of the network server.

For configurations with more than one virtual SCSI Server/Client pairing, ensure that the resource identified in the resource field identifies the Virtual SCSI Server adapter that is connected to the virtual SCSI Client adapter in i5/OS that provides the virtual disk with the



operating system installed on it. Remember that you can use the location information of the virtual resource (shown in Figure D-4) to determine the virtual slot number of the adapter.

```
Display Resource Detail
System: G10B8DAE
Resource name : CTL02
Text : Comm Adapter
Type-model : 290B-001
Serial number : 00-00000
Part number :

Location: U9406.520.10B8DAE-V7-C3

Logical address:
SPD bus:
System bus 255
System board 0
System card 0
More...

Press Enter to continue.

F3=Exit F5=Refresh F6=Print F12=Cancel
```


Figure D-4 Display Resource Details

The C# portion of the location field provides the virtual slot number (in this case 3) that resource is associated with.

## Vary on NWS fails - host partition not defined

### Symptom

An attempt to vary on the Network Server fails and an error code 9 - IPL of partition failed (as shown in Figure D-5) occurs.



```
Session A - [24 x 80]
File Edit View Communication Actions Window Help
Additional Message Information
Message ID : CPDB1AD
Date sent : 08/23/04 Time sent : 10:53:20
Message : Network server LINUX1 vary on failed.
Cause : The network server failed to vary on with reason code
00000009. Reason codes and their meanings follow:
--0001: Storage space not found.
--0002: Partition not found.
--0003: Partition not correct type.
--0004: Server already varied on for partition.
--0005: Partition already active.
--0006: A file system error occurred accessing the IPL stream file.
--0007: IPL source not valid.
--0008: IPL of partition failed.
--0009: A host partition is not configured for this partition.
More...
Press Enter to continue.
F1=Help F3=Exit F6=Print F9=Display message details F12=Cancel
F21=Select assistance level
MA a 01/001
```

Figure D-5 Troubleshooting - IPL of partition failed

### Problem determination

This error occurs when the vary on of the network server sends a power control request to the Hypervisor, however, the partition definition has not be defined with the i5/OS partition as a power controlling partition.

## Corrective action

Make sure that the power controlling partition is set in the partition profile to the partition that you are trying to vary on the Network Server from as shown in Figure D-6.

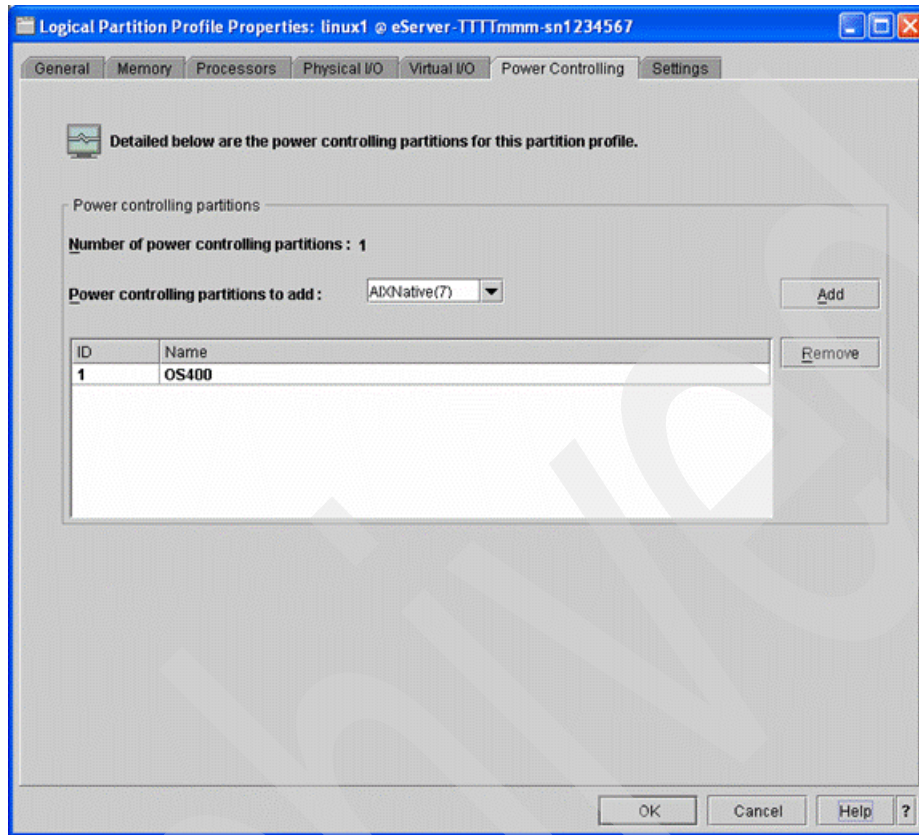


Figure D-6 Troubleshooting - Power control partition


**Note:** This problem will also occur if the Linux partition was not activated (via the Hardware Management Console (HMC)) after it was configured. Remember that an activation of the logical partition from the HMC is required whenever a change to the partition profile is made to register that change with the Open Firmware of the managed system.

If a change is made to the partition profile to correct this problem (that is, the power controlling partition is configured). Remember to activate the partition from the HMC to register the new configuration.

## File system error accessing stream file

### Symptom

An attempt to vary on the Network Server (with IPL Source set to \*STMF) fails with an error code 6 as shown in Figure D-7.



```
Session A - [24 x 80]
File Edit View Communication Actions Window Help
Additional Message Information
Message ID : CP0B1AD
Date sent : 08/23/04 Time sent : 12:47:10
Message : Network server LINUX1 vary on failed.
Cause : The network server failed to vary on with reason code
00000006. Reason codes and their meanings follow:
--0001: Storage space not found.
--0002: Partition not found.
--0003: Partition not correct type.
--0004: Server already varied on for partition.
--0005: Partition already active.
--0006: A file system error occurred accessing the IPL stream file.
--0007: IPL source not valid.
--0008: IPL of partition failed.
--0009: A host partition is not configured for this partition.
Press Enter to continue.
F1=Help F3=Exit F6=Print F9=Display message details
F10=Display messages in job log F12=Cancel F21=Select assistance level
MSE a 01/001
```

Figure D-7 Troubleshooting - error accessing stream file

### Problem determination

Check the value of the IPL stream file (in the Network Server Description) and ensure that it is a valid IFS path or file.

### Corrective action

Often the cause of this error is that the CD is not mounted. Additionally, if you are using a virtual CD catalog, it is possible that the virtual CD device is not varied on, the catalog is not loaded on the device, or the wrong volume is mounted in the library.

## Inability to access the Linux console through i5/OS virtual console support

### Symptom

In addition to virtual console support through the HMC, virtual console support is provided through a TCP/IP enabled virtual console application in i5/OS. One possible problem with this

support is that an attempt to access the virtual console indicates that the Linux partition is not displayed in the list of guest partitions as shown in Figure D-8.

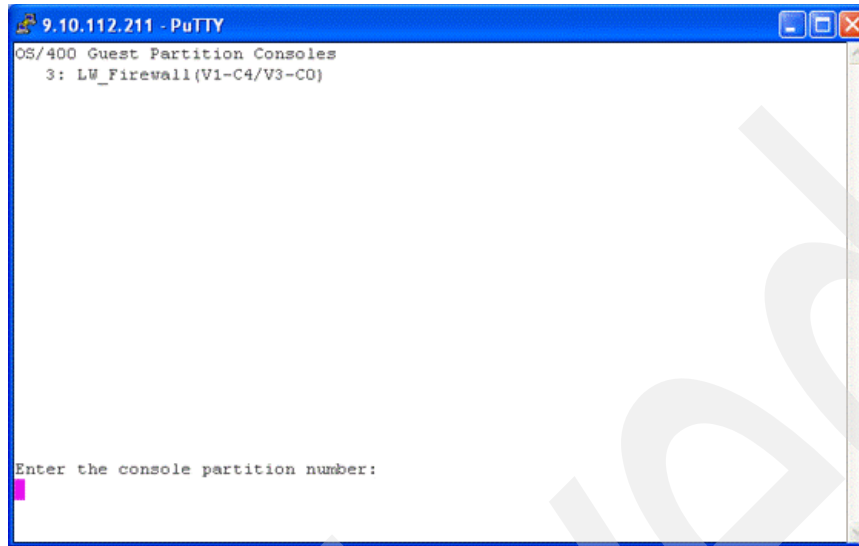


Figure D-8 Troubleshooting - Linux console not listed in virtual console list

### Problem determination

Failure of the Linux partition to be displayed in the list of partitions may be a virtual serial adapter connection has not been established between the Linux and i5/OS partitions.

### Corrective action

Create a “client” serial adapter (as shown in Figure D-9) in the hosting i5/OS partition that connects to the “server” serial adapter in slot 0 of the Linux partition.

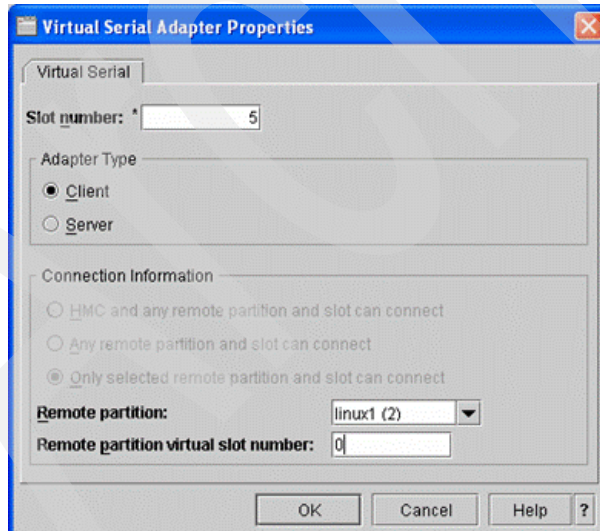


Figure D-9 Troubleshooting - Create Client Serial Adapter

## Virtual console access denied (through i5/OS) - console in use

### Symptom

An attempt to access the Linux console through the virtual console support in i5/OS fails with a message indicating that the console is already in use as shown in Figure D-10.

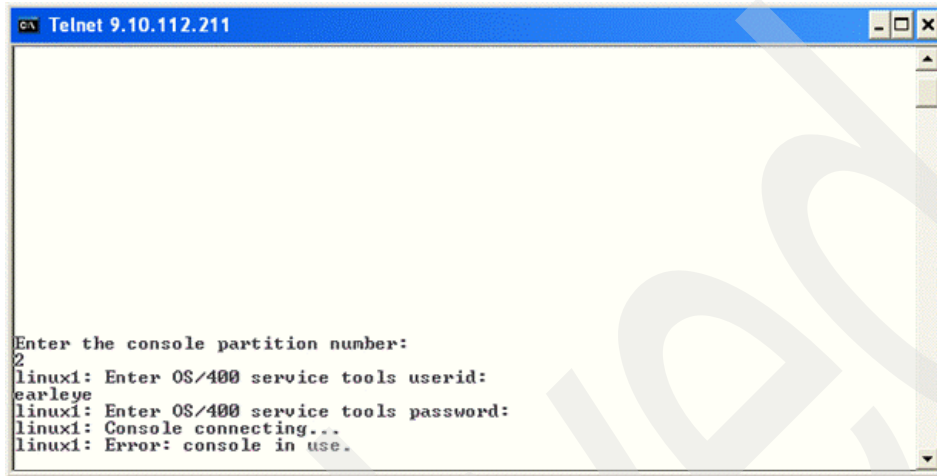


Figure D-10 Troubleshooting - virtual console in use

### Problem determination

This error typically occurs when someone has access to the console for the Linux partition through the HMC.

**Note:** While concurrent access to the console through the HMC and the virtual console support in i5/OS is not permitted, multiple concurrent connections to the console through the virtual console support in i5/OS is allowed. Keep in mind that each connection to the console is actually the same connection and anything that one user types in the console will be displayed in the other console windows.

### Corrective action

Terminate the HMC terminal connection (as shown in Figure D-11) and then re-attempt the virtual console connection.

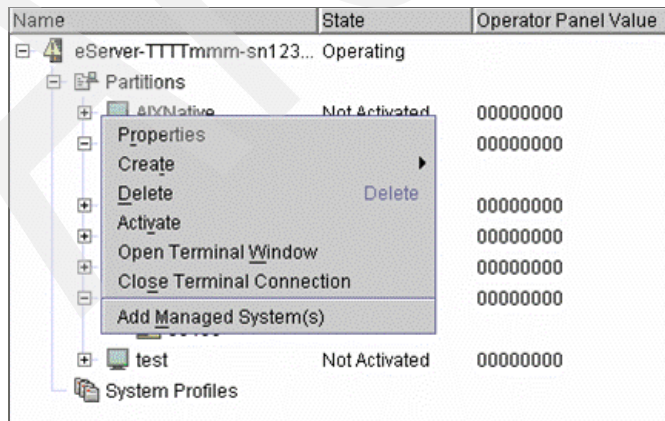


Figure D-11 Troubleshooting - Terminate Console Connection (HMC)

## Inability to access Linux console (via HMC) - console in use

### Symptom

An attempt to access the console through the HMC (or WebSM) results in an error message indicating that the console is in use.

### Problem determination

One cause of this condition could be that the console is already being accessed via a virtual console through the hosting i5/OS partition. Simultaneous access to the Linux console from the HMC and hosting i5/OS partition is not allowed.

### Corrective action

You can use the `netstat *cnn` command within the i5/OS hosting partition to see if there is active virtual console connection (refer to Figure D-12) and terminate the connection.

```
Session A - [24 x 80]
Work with TCP/IP Connection Status
System: S1069BCD
Type options, press Enter.
3=Enable debug 4=End 5=Display details 6=Disable debug
8=Display jobs

 Remote Remote Local
Opt Address Port Port Idle Time State
-- -
* * * as-signon 000:27:00 Listen
* * * as-netd > 000:27:54 Listen
* * * as-tran > 000:27:00 Listen
* * * as-vrtp > 000:26:59 Listen
- 9.10.86.219 3682 as-vcons 000:03:53 Established

Bottom
F3=Exit F5=Refresh F9=Command line F11=Display byte counts F12=Cancel
F15=Subset F22=Display entire field F24=More keys

MB a 09/003
```

Figure D-12 Work with TCP/IP Connection Status

## Partition fails to start - resource conflict

### Symptom

An attempt to activate a partition from the HMC fails and a resource conflict dialog box (shown in Figure D-13) is displayed.

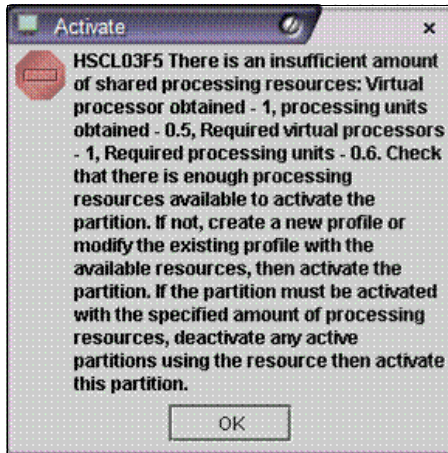


Figure D-13 HMC - Resource Allocation Error

### Problem determination

As the error dialog box suggests, the minimum CPU units (in this example shown in Figure D-13) required by the partition exceeds what is available on the system.

### Recovery

Decrease the minimum CPU requirement in the partition's profile (as shown in Figure D-14).

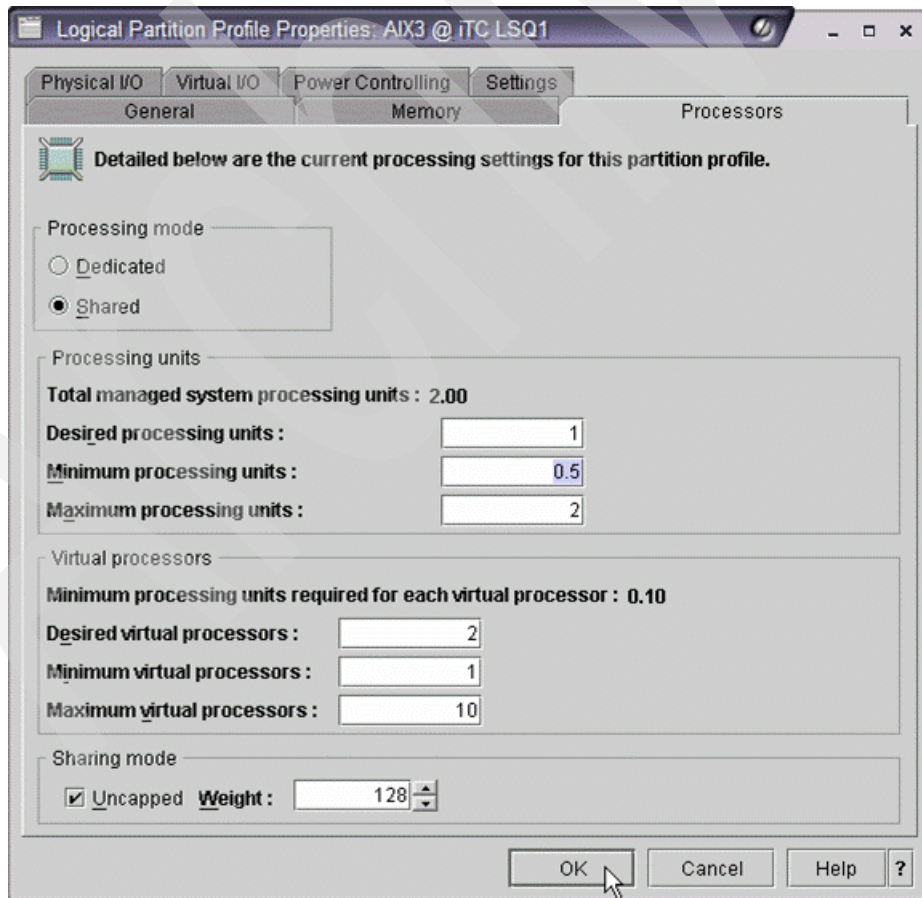


Figure D-14 Partition Profile - Processor Properties



Alternatively, use Dynamic LPAR to free some CPU units from another partition.

## Inability to open Linux console in WebSM

### Symptom

An attempt to access the Linux console through WebSM generates the following error message shown in Figure D-15.




Figure D-15 WebSM Console Access - Error Dialog

### Recovery

Click **HMC Configuration** in the navigator pane, then click **Enable or Disable Virtual Terminal**. Check the Enable check box (as shown in Figure D-16) and click **OK**.

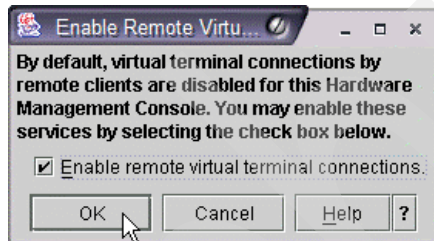


Figure D-16 WebSM - Enable Remote Virtual Console

## Linux does not shutdown when the network server is varied off

### Symptom

When the Network Server is varied off, the Linux operating system does not perform an operating system shutdown.

### Problem determination

The Linux on Power Utilities are not installed.

### Recovery

Download and install the Linux on Power Utilities that supports power off. You require the following packages to support shutdown of Linux:

- ▶ Platform Enablement Library (librtas)
- ▶ Error Log Analysis (diagela)

## An RPA communication error occurs when a DLPAR operation is attempted

### Symptom

When a DLPAR function (such as move of processors) is attempted, the HMC reports that communications with the partition could not be established.

### **Problem determination**

The Linux on Power utilities that supports Dynamic LPAR is not installed.

### **Recovery**

Download and install the Linux on Power Utilities that supports Dynamic LPAR.

You require the following packages for support of Dynamic LPAR:

- ▶ CSM Core (csm-core)
- ▶ RSCT Utilities (rsct.core.utils)
- ▶ SRC (src)
- ▶ RSCT Core (rsct.core)
- ▶ Remote File Distribution (rdist)
- ▶ CSM client (csm-client)
- ▶ Dynamic RM (Dynamic RM)
- ▶ Dynamic Reconfiguration Tools (rpa-dlpar)

# Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

## IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 359. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *IBM AS/400 Printing V*, SG24-2160
- ▶ *Linux on the IBM eServer iSeries Server: An Implementation Guide*, SG24-6232
- ▶ *Logical Partitions on System i5: A Guide to Planning and Configuring LPAR with HMC on System i*, SG24-8000
- ▶ *V5 TCP/IP Applications on the IBM eServer iSeries Server*, SG24-6321
- ▶ *Virtual Partition Manager A Guide to Planning and Implementation*, REDP-4013
- ▶ *Windows-based Single Signon and the EIM Framework on the IBM eServer iSeries Server*, SG24-6975

## Other publications

These publications are also relevant as further information sources:

- ▶ *Backup and Recovery V5R3*, SC41-5304

## Online resources

These Web sites and URLs are also relevant as further information sources:

- ▶ *The Book of Webmin Or: How I Learned to Stop Worrying and Love UNIX*  
<http://www.swelltech.com/support/pdfs/webminguide.pdf>
- ▶ Downloading and Installing Webmin  
<http://www.webmin.com/download.html>
- ▶ General discussion of Webmin is at:  
<http://www.webmin.com>
- ▶ Gentoo Linux home page  
<http://www.gentoo.org>
- ▶ Gentoo Linux Projects: Gentoo Linux PPC64 Development  
<http://ppc64.gentoo.org>
- ▶ Heimdal Kerberos 5  
<http://www.pdc.kth.se/heimdal>

- ▶ IBM @server i5 and iSeries System Handbook for i5/OS V5R3 is available at:  
<http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/ga195486.html?Open>
- ▶ IBM @server iSeries ODBC Driver for Linux: Connection String Keywords and Values  
<http://www.ibm.com/servers/eserver/iseries/linux/odbc/guide/odbcproperties.html>
- ▶ IBM Resource Link  
<http://www.ibm.com/servers/resourceLink>
- ▶ IBM Systems Information Centers  
<http://publib.boulder.ibm.com/iseries/>
- ▶ iSeries Access for Linux  
<http://www.ibm.com/servers/eserver/iseries/access/linux>
- ▶ Kerberos: The Network Authentication Protocol  
<http://web.mit.edu/kerberos/www>
- ▶ The Linux Documentation Project  
<http://www.tldp.org>
- ▶ The Linux Home Page at Linux Online  
<http://www.linux.org>
- ▶ LPAR Overview  
<http://www.iseries.ibm.com/lpar>  
<http://www-1.ibm.com/servers/eserver/iseries/lpar>
- ▶ Novell: SUSE Linux  
<http://www.suse.com>
- ▶ Novell: SUSE Linux Enterprise Server 9  
<http://www.suse.com/sles/documentation/samba>
- ▶ Novell: Why Choose Novell for Linux?  
<http://www.novell.com/linux>
- ▶ OpenOffice.org  
<http://www.openoffice.org>
- ▶ OpenPKG: OpenPGP Key Server  
<http://pgp.openpkg.org>
- ▶ PuTTY: A Free Telnet/SSH Client  
<http://www.chiark.greenend.org.uk/~sgtatham/putty>
- ▶ RealVNC  
<http://www.realvnc.com>
- ▶ Red Hat Linux home page  
<http://www.redhat.com>
- ▶ Service and productivity tools for Linux on POWER  
<https://www14.software.ibm.com/webapp/set2/sas/f/topdiags/home.html>
- ▶ Standalone Diagnostics CD-ROM  
<http://www14.software.ibm.com/webapp/set2/sas/f/diags/home.html>

- ▶ Third-party modules for Webmin  
<http://webadminmodules.sourceforge.net>
- ▶ The unixODBC Project home page  
<http://www.unixodbc.org>
- ▶ Various Linux related information about System i is available at:  
<http://www.ibm.com/servers/eserver/series/linux>
- ▶ WinSCP: Freeware SFTP and SCP client for Windows  
<http://winscp.sourceforge.net/eng>
- ▶ Workload Estimator (WLE) is available at:  
<http://www-912.ibm.com/wle/EstimatorServlet>

## How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

[ibm.com/redbooks](http://ibm.com/redbooks)

## Help from IBM

IBM Support and downloads

[ibm.com/support](http://ibm.com/support)

IBM Global Services

[ibm.com/services](http://ibm.com/services)

Archived

# Index

## Symbols

%CPU 253  
%MEM 253  
\*AUTO 346  
\*BASE 335  
\*CALC 126  
\*DFTSTGPTH 126  
\*GUEST 119  
\*LINUXPPC 119  
\*LNGVER 119  
\*MACHINE 335  
\*PANEL 122  
\*READ 125  
\*SAVF 264  
\*SHRUPD 126  
\*STMF 122  
\*UPDATE 125  
/dev/sda1 288  
/dev/sda2 288  
/dev/sda3 288  
/etc/EXPORTS 269  
/mnt 307  
/proc/cpuinfo 21  
/QFPNWSSTG 264  
/qopt path 318  
/qopt/RED\_HAT/IMAGES/PSERIES 122  
/qopt/SLES100.001/SUSEBOOT/INST64 122  
/qopt/SU90.001/install 122  
/qopt/SU90SP3.P01 122  
/sys/class/scsi\_host/host0/scan 300  
/vdevice/v-scsi 130

## Numerics

268C 21  
802.1Q 23

## A

ACL icon 260  
Add Mount 270  
ADDIMGCLGE 315  
ADDTCPIFC 158, 165  
advanced system management 3  
Apache 124, 336  
ARP 22  
ASCII 278  
ASM 3  
automatic processor movement 4

## B

balance workload 331  
bckhome.sh 273  
boot manager 3

boot option 128  
boot time driver 3  
broadcast 337

## C

Capacity on Demand 21  
capped 18  
case sensitive 344  
case sensitivity 344  
CD images 311  
CD library 316  
CFGOBJ 264  
clone 20  
code page 119  
command line or program name 253  
console 292, 352  
console in use 352  
Contents Area 5  
cost of switching 18  
CPU time 253  
CPU usage 253  
creating a backup script 273  
Cron 275  
CRTDEVOPT 312  
CRTIMGCLG 314  
CRTLIB 264  
CRTLINETH 153  
CRTNWS 118  
CRTNWSSTG  
    i5/OS commands  
        CRTNWSSTG 23  
CRTSAVF 264  
CSM client 356  
CSM Core 356  
csm-client 356  
csm-core 356  
CUoD 5, 20

## D

datagram forwarding 155  
dedicated processor 18, 61  
default gateway 162  
default partition profile 54  
default weight 19  
desired memory 60  
desired processing unit 61  
desired value 54  
device driver 3  
device type 129  
df -h 310  
diagela 355  
direct attached 144  
direct I/O 24  
direct routing 176

- dispatch cycle 339
- Display Disk Array Status screen 138
- DLPAR 2, 21
- DMZ 26
- DNS 165
- DSPSAVF 284
- DST 5, 13
- dynamic logical partitioning 21
- dynamic real-time view 252
- Dynamic Reconfiguration Tools 356
- Dynamic RM 356

## E

- EBCDIC 278
- echo 300
- Enterprise Workload Manager 24
- error code 9 348
- Error Log Analysis 355
- establishing networking 143
- EWLM 24
- EXT icon 260
- ext2 294
- ext3 294

## F

- failure message 344
- fdisk -l 293
- find 279
- firmware 2
- flash memory 2
- flexibility 287
- flexible storage environment 287
- fractional processor allocation 338
- fractions 18
- fragmentation 336
- fragmented 336
- frame size 331, 336
- free 254
- full system partition profile 54

## G

- gateway 144
- Grid Toolbox 24
- group of partitions 58
- group of physical processors 18

## H

- hardware error 55
- hardware management console 4
- hardware requirements 44, 48
- Hardware tab 250
- HMC 4
  - Contents Area 5
  - Navigation Area 5
- HMC fails 353
- HMC terminal connection 352
- hosted 24
- hybrid scenario 49

## I

- I/O bound Linux workload 335
- I/O pool 20, 65
- I/O throughput 332
- I3P 174
- i5/OS commands
  - CRTLIB 264
  - CRTSAVF 264
  - DSPSAVF 284
  - INZTAP 278
  - RST 284
  - RSTCFG 284
  - SAV 265
  - VRYCFG 264
- IBM Director Multiplatform 24
- IBM Grid Toolbox 24
- IEEE 802.11Q 23
- ifconfig 44
- image catalog 315
- Info icon 260
- input/output device loading 254
- installation media 131
- INTNETADR 158
- INZTAP 278
- iostat 254
- IPL stream file 350
- iSeries Navigator 13, 166
- ISO file 315

## K

- kernel version 2.6 12

## L

- layer 2 bridging 23
- LCLIFC 158
- librtas 355
- limited interactive interface 252
- LIND 158
- Linux commands
  - find 279
  - free 254
  - iostat 254
  - top 252
  - vmstat 253
- Linux console options 51
- Linux disk partitions 288
- Linux instance 53
- Linux kernel version 2.6 12
- Linux RAID sets 139
- Linux utilities 255
- list of tasks 252
- LODIMGCLG 316
- Low Level Firmware 3
- LVM 308

## M

- map address translation' 169
- maximum memory 60



- maximum processing unit 61
- maximum value 54
- media type 129
- memory cache 18
- memory pools 334
- memory requirements 334
- memory size 60
- memory usage 253
- micro-partitioning 18, 61
- minimum memory 60
- minimum processing unit 61
- minimum value 54
- mismatch 344
- MKDIR 313
- mke2fs 302
- mount 131, 350
- Mounting the NFS directory to Linux file system 269
- multi-path I/O support 332
- multiple client/server SCSI pairs 334
- multiple logical partitions 18
- multiple network servers 334
- multiple SCSI chains 334
- multiple software threads 332

## N

- NAT 22, 144
- Navigation Area 5
- netstat 353
- Network Address Translation 337
- network address translation 22
- Network Filesystem 270
- network interface 336
- network packet 336
- network server description 23
- network stack 338
- NFS Export 268
- NFS server 267
- NI 253
- NIC 45, 49
- nice value 253
- non-swapped physical memory used 253
- non-hosted 24
- normal mode boot 131
- NWSD 23
- NWSSTG 23, 118
- NWSSTG level backup 264

## O

- ODBC 336
- OF 3
- Open Firmware 3, 56, 345
- open-source applciation 124
- optimal processor assignment 18
- overcommitment 54
- overcommitted resources 55

## P

- Packet Rules 166

- partition ID 54
- partition profile 54
- PCI adapter 3
- performance characteristics 331
- performance degradation 338
- performance perspective 336
- PID 253
- Platform Enablement Library 355
- policy parameter 55
- power controlling partition 349
- POWER Hypervisor 2
- Power RAID Configuration Utility 141
- Power RAID Configuration Utility menu 137
- POWER5 2
- PR 253
- predictable performance 19
- PrepBoot 288
- primary partition 5
- priority 253
- private addresses 338
- Process ID 253
- process manipulation 252
- process status 253
- processing lag 339
- processor thread 338
- Processors and Memory tab 251
- provisioning 24
- Proxy ARP 144, 337
- proxy ARP 22
- PTF load process 311
- PuTTY 257

## Q

- QPRFADJ 335
- QSYSOPR 344

## R

- R/W protected 135
- RAID set 137
- range of addresses 157
- rdist 356
- real network 338
- Red Hat Enterprise Linux AS 3
  - introduction 12
- Redbooks Web site 359
  - Contact us xi
- register 345
- registers 21
- reiserfs 294
- relative weight 19
- Remote File Distribution 356
- remove link 295
- report error 55
- RES 253
- rescue 133
- resize2fs 294
- resizing 287
- resource conflict 353
- resource sharing 20

- restarts 252
- restoring NWSD 284
- Restrict Device Resources 334
- restricted state 24, 42
- root 134
- router 144, 337
- RPA 355
- rpa-dlpar 356
- RSCT Core 356
- RSCT Utilities 356
- rsct.core 356
- rsct.core.utils 356
- RST 284
- RSTCFG 284
- RSTDDEVRSC 334
- RTAS 3
- Rules Editor 166
- Run-Time Abstraction Services 3

## S

- S 253
- SAMBA 124
- SAV 265
- Save file 264
- scheduling backup 274
- SCP 263
- SCSI 125
- Select Disk Units for Disk Array menu 139
- server error 55
- service partition 55
- service processor 2–3
- SFTP 263
- share processors 331
- shared memory size 253
- shared pool 18
- shared processor 18, 61
- shared processor pool 18
- shared processor units 19
- SHR 253
- simultaneous multi-threading 21
- single physical processor 18
- single-level store 331
- SLES 9
  - intro 12
- SLIC 3
- slot 0 351
- smaller network packet 336
- smaller segment 337
- SMS 3
- SMT 2, 21
- SP 2–3
- SRC 356
- src 356
- SSH 257
- SST 13, 51
- static NAT rules 165
- STRTCPIFC 159, 165
- sub processor feature 61
- subnet 157
- subnet mask 144

- SUBNETMASK 158, 166
- SUSE Linux Enterprise Server 9
  - intro 12
- swap 21, 288
- swap memory 254
- swap utilization 253
- sysstat 254
- System Licensed Internal Code 3
- system management services 3
- system policy parameter 55
- system profile 55
- System Services Tool 51
- system summary information 252
- system-config-network-tui 162
- systems provisioning 24

## T

- TCP/IP interface 338
- terminal session 292
- threads 21
- TIME+ 253
- top 252
  - columns header 253
  - message/prompt line 253
  - summary area 253
  - task area 253
- trunk adapter 23

## U

- uncapped 18
- uncapped logical partition 19
- uncapped partition processing 338

## V

- V5R3M0 332
- validation 55
- VE 20
- VIRT 253
- virtual CD 23
- virtual CD device 311
- virtual CD environment 311
- virtual CD library 311
- virtual console 51
- virtual ethernet
  - benefits 22
- virtual I/O 21, 24
- virtual I/O driver 332
- virtual I/O performance 332
- virtual LAN 21
- virtual memory used in the task 253
- Virtual Network Computing 258
- virtual processor 338
- virtual SCSI adapter 23
- virtual serial adapter 23
- virtual tape 23, 278
- virtualization 287
- virtualization engine 20
- vmstat 253

VNC 258  
VNC=1 122  
VPN 175  
VRYCFG 264, 313  
VTerm 51

## **W**

Web-based System Manager 10  
Webmin 258  
WebSM 10, 355  
weight 19  
WinSCP 263  
Work with Disk Arrays menu 138  
Work with Disk Unit Recovery menu 136  
workload 331  
WRKCFGSTS \*CMN 119  
WRKHDWRSC 151  
WRKIMGCLG 317  
WRKSYSSTS 334

## **Y**

YaST 256  
YaST Control Center 159

Archived

Archived



Redbooks

## Implementing POWER Linux on IBM System i Platform

(0.5" spine)  
0.475" x 0.873"  
250 <-> 459 pages







# Implementing POWER Linux on IBM System i Platform



**Planning and configuring Linux servers on IBM System i platform**

**Linux distribution on IBM System i Platform installation guide**

**Tips to run Linux servers on IBM System i platform**

This IBM Redbook is about POWER Linux implementation on the IBM System i platform. With the brand-new technologies of POWER5 and IBM Virtualization Engine, POWER Linux on i5 is even more powerful, scalable, and flexible. We also have new versions of Linux distributions that are available for POWER Linux on IBM System i platform. They are SUSE Linux Enterprise Server (SLES) 10 and Red Hat Enterprise Linux (RHEL) 4.

This book gets you ready to plan, configure, and install Linux on System i platform. It provides information about setting up a network connection between the server i5/OS partition and the client Linux partition over virtual LAN using various methods. It also covers various topics of administration, operations, and tips and techniques.

**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](http://ibm.com/redbooks)

SG24-6388-01

ISBN 0738489689