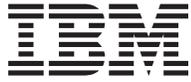




System i
Linux in a Logical Partition

Version 6 Release 1





System i
Linux in a Logical Partition

Version 6 Release 1

Note

Before using this information and the product it supports, be sure to read the information in "Notices," on page 65.

This edition applies to version 6, release 1, modification 0 of i5/OS (product number 5761-SS1) and to all subsequent releases and modifications until otherwise indicated in new editions. This version does not run on all reduced instruction set computer (RISC) models nor does it run on CISC models.

© Copyright International Business Machines Corporation 1998, 2008. All rights reserved.

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

Contents

Linux in a logical partition	1
What's new for V6R1	1
PDF file for Linux in a logical partition	1
Planning to run Linux in a logical partition	2
Linux support on System i products	2
Hosted versus nonhosted partition running Linux	5
Virtual I/O in a logical partition running Linux	5
Directly attached I/O for Linux running in logical partitions on System i i models 8xx and 270	6
System i I/O adapters (IOAs) supported by Linux	6
Directly attached SCSI support for a logical partition	10
Obtaining Linux for System i products	10
Ordering or upgrading a System i product to run Linux in a logical partition	11
Creating a logical partition to run Linux	11
Configuring a logical partition	11
Assigning I/O adapters (IOAs) to a logical partition	12
Creating a network server description and a network server storage space	16
Configuring the LAN console for Linux running a logical partition	17
Connecting the virtual console	17
Installing Linux in a logical partition	18
Installing Linux from CD-ROM images in the integrated file system	18
Starting and stopping the NWSD for a Linux partition	20
Starting device parity protection and creating a disk array	20

Managing Linux in a logical partition	21
Authority for Linux running in logical partitions	21
Displaying and changing hardware configuration	22
Connecting your logical partition to the LAN	23
Configuring a print server	30
Moving shared processing power between logical partitions	31
Dynamically adding disks to a Linux partition	32
Network server description	34
Application support for Linux in a logical partition	38
Backing up and recovering a logical partition	40
Starting a logical partition	56
Troubleshooting Linux running in a logical partition	59
Debugging NWSD error messages for Linux running in a logical partition	59
Debugging a processor multitasking error	60
System reference codes (SRCs) for Linux running in a logical partition	61
Recovering from a virtual Ethernet network failure	62
Related information for Linux in a logical partition	63

Appendix. Notices	65
Programming Interface Information	67
Trademarks	67
Terms and conditions	67
Code license and disclaimer information.	68

Linux in a logical partition

Install, configure, manage, and troubleshoot Linux in logical partitions.

IBM® and a variety of Linux distributors have partnered to integrate the Linux operating system with the reliability of the System i™ product. Linux brings a new generation of Web based applications to the System i platform. IBM has modified the Linux PowerPC® kernel to run on previous System i products and in a logical partition hosted by either the AIX® or i5/OS® operating systems and contributed the kernel back to the Linux® community.

Use this information to manage Linux running in logical partitions on System i models 8xx and 270 and Linux servers that use virtual disks and resources hosted by the i5/OS operating system.

For information about running Linux on model 5xx systems, see Partitioning for Linux with an HMC in the IBM Systems Hardware Information Center.

What's new for V6R1

Read about new or significantly changed information for the Linux in a logical partition topic collection.

Enhanced back up and recovery

The IBM Extended Integrated Server Support for i5/OS licensed program (5761-LSV) provides support for saving and restoring individual files from storage spaces (also known as virtual disks) that are provided by the i5/OS operating system. Support is also added for saving storage spaces for active Linux servers from the i5/OS operating system.

- “Backing up storage spaces for an active Linux server in a logical partition” on page 48
- “Backing up and recovering individual files and directories for Linux servers running in logical partitions” on page 49

How to see what's new or changed

To help you see where technical changes have been made, the information center uses:

- The  image to mark where new or changed information begins.
- The  image to mark where new or changed information ends.

In PDF files, you might see revision bars (|) in the left margin of new and changed information.

To find other information about what's new or changed this release, see the Memo to users.

PDF file for Linux in a logical partition

You can view and print a PDF file of this information.

To view or download the PDF version of this document, select Linux in a logical partition.

Saving PDF files

To save a PDF on your workstation for viewing or printing:

1. Right-click the PDF link in your browser.
2. Click the option that saves the PDF locally.

3. Navigate to the directory in which you want to save the PDF.
4. Click **Save**.

Downloading Adobe Reader

You need Adobe Reader installed on your system to view or print these PDFs. You can download a free copy from the Adobe Web site (www.adobe.com/products/acrobat/readstep.html) .

Planning to run Linux in a logical partition

Before you configure your System i product to run the Linux operating system, you need to carefully plan and evaluate your software and hardware resources.

This section guides you through the planning process before you create a logical partition to run Linux.

For more information, refer to Partitioning for Linux .

Note: It is recommended that you perform a full server back up before changing your server configuration. Refer to Back up your server for instructions on performing a full server back up.

Linux support on System i products

To successfully partition a System i product to run the Linux operating system, the system requires specific hardware and software.

The primary partition must run i5/OS V5R1 or later versions and be updated with the latest PTFs. The latest Linux-related i5/OS PTFs can be found on Linux on the System i platform .

Linux is not supported on a primary partition. Additional information on software function can be found at i5/OS logical partition function by release.

Selected models can run Linux in the shared processor pool configuration. When using shared processors, one processor can be shared among up to 10 i5/OS and Linux partitions. Other models require the use of dedicated processors for a Linux partition. Those same models also require you to disable processor multitasking for the whole server, including the primary partition. See Debug the processor multitasking error for more information.

The Linux operating system supports single processors or multiple processors. This choice is made at the time the logical partition is created. If a Linux kernel built for a single processor is loaded into a partition with multiple processors assigned to it, it will function correctly, but only one processor will be used. If multiple processors are assigned to a partition, Linux built for Symmetric MultiProcessors (SMP) should be used. Any number of available processors may be assigned to a logical partition.

The following tables help you determine if your server supports Linux in a logical partition. For information about Linux support for other models, refer to Partitioning the server.

Table 1. Model 270

Feature code	Number of processors	Supported LPAR	Supports Linux	Linux shared processor	Disable processor multitasking
2248	1	No	No	No	No
2250	1	No	No	No	No
2252	1	No	No	No	No
2253	2	No	No	No	No

Table 1. Model 270 (continued)

Feature code	Number of processors	Supported LPAR	Supports Linux	Linux shared processor	Disable processor multitasking
2422	1	No	No	No	No
2423	1	No	No	No	No
2424	2	No	No	No	No
2431	1	Yes	Yes	Yes	No
2432	1	Yes	Yes	Yes	No
2434	2	Yes	Yes	Yes	No
2452	1	Yes	Yes	Yes	No
2454	2	Yes	Yes	Yes	No

Table 2. Model 800

Feature code	Number of processors	Support LPAR	Support Linux	Linux shared processor	Disable processor multitasking
	1	Yes	Yes	Yes	No

Table 3. Model 810

Feature code	Number of processors	Support LPAR	Support Linux	Linux shared processor	Disable processor multitasking
	1	Yes	Yes	Yes	No
	2	Yes	Yes	Yes	No

Table 4. Model 820

Feature code	Number of processors	Support LPAR	Support Linux	Linux shared processor	Disable processor multitasking
2395	1	Yes	No	No	No
2396	1	Yes	No	No	No
2397	2	Yes	Yes	Yes	Yes
2398	4	Yes	Yes	Yes	Yes
2425	1	Yes	No	No	No
2426	2	Yes	Yes	No	Yes
2427	4	Yes	Yes	No	Yes
0150	1	Yes	Yes	Yes	No
0151	2	Yes	Yes	Yes	No
0152	4	Yes	Yes	Yes	No
2435	1	Yes	Yes	Yes	No
2436	1	Yes	Yes	Yes	No
2437	2	Yes	Yes	Yes	No
2438	4	Yes	Yes	Yes	No
2456	1	Yes	Yes	Yes	No
2457	2	Yes	Yes	Yes	No
2458	4	Yes	Yes	Yes	No

Table 5. Model 825

Feature code	Number of processors	Support LPAR	Support Linux	Linux shared processor	Disable processor multitasking
	3	Yes	Yes	Yes	Yes
	4	Yes	Yes	Yes	Yes
	5	Yes	Yes	Yes	Yes
	6	Yes	Yes	Yes	Yes

Table 6. Model 830

Feature code	Number of processors	Supports LPAR	Supports Linux	Linux shared processor	Disable processor multitasking
2400	4	Yes	Yes	No	Yes
2402	4	Yes	Yes	No	Yes
2403	8	Yes	Yes	No	Yes
2351	4/8	Yes	Yes	No	Yes
0153	4/8	Yes	Yes	No	Yes

Table 7. Model 840

Feature code	Number of processors	Supports LPAR	Supports Linux	Linux shared processor	Disable processor multitasking
2418	12	Yes	Yes	No	Yes
2420	24	Yes	Yes	No	Yes
2416	8/12	Yes	Yes	No	Yes
2417	12/18	Yes	Yes	No	Yes
2419	12/18	Yes	Yes	No	Yes
2461	24	Yes	Yes	Yes	No
2352	8/12	Yes	Yes	Yes	No
2353	12/18	Yes	Yes	Yes	No
2354	18/24	Yes	Yes	Yes	No
0158	12	Yes	Yes	Yes	No
0159	24	Yes	Yes	Yes	No

Table 8. Model 870

Feature code	Number of processors	Support LPAR	Support Linux	Linux shared processor	Disable processor multitasking
	8/16	Yes	Yes	Yes	Yes

Table 9. Model 890

Feature code	Number of processors	Supports LPAR	Supports Linux	Linux shared processor	Disable processor multitasking
0197	24	Yes	Yes	Yes	Yes
0198	32	Yes	Yes	Yes	Yes
2487	16/24	Yes	Yes	Yes	Yes
2488	24/32	Yes	Yes	Yes	Yes

Contact your IBM marketing representative or IBM Business Partner for more information on upgrading existing feature codes on your server.

Refer to Linux on the System i platform  (www.ibm.com/systems/i/os/linux/) for the latest updates to Linux running on a System i product.

Hosted versus nonhosted partition running Linux

The Linux operating system runs in both a hosted and nonhosted partitioned environment on System i products.

A *hosted* partition uses I/O resources that belong to a hosting i5/OS partition. The hosting i5/OS partition can either be a primary or a secondary partition. The I/O resources that a hosted partition can use from a server include disk, CD, and tape devices.

The hosted logical partition must be started from the hosting i5/OS partition by varying on a network server description (NWS D). The NWS D is used to control the logical partition. The logical partition can only be active when the server is both active and out of restricted state. When i5/OS is in restricted state, all NWS Ds are automatically varied off. An i5/OS partition can host multiple logical partitions. You need to perform proper capacity planning to ensure that the server can support logical partitions. Refer to IBM System Planning Tool  for information on the IBM System Planning Tool.

You should not power on a hosted logical partition using the Work with Partitions Status display. If you power on the logical partition running Linux using the Work with Partitions Status display, none of the virtual I/O devices will be available.

A *nonhosted* logical partition is not dependent on a hosting i5/OS partition for any I/O resources. The logical partition has its own disk units or the partition makes use of networking support to do a network start. A nonhosted logical partition can be started even if the primary partition is not fully active. A nonhosted logical partition can be started from the Work with Partitions Status display.

Virtual I/O in a logical partition running Linux

Virtual I/O resources are devices owned by the hosting i5/OS[®] partition that provides I/O function to the logical partition.

The System i Linux kernel and i5/OS support several different kinds of virtual I/O resources. They are virtual console, virtual disk unit, virtual CD, virtual tape, and virtual Ethernet.

Virtual console provides console function for the logical partition through an i5/OS partition. The virtual console can be established to the server or to the primary partition. The use of the virtual console allows the installation program to communicate with the user prior to networking resources being configured. It can also be used for troubleshooting server errors.

Virtual disk unit can only be provided by a server. Virtual disk provides access to NWSSTG virtual disks for Linux. By default, the CRTNWSSTG command creates a disk environment with one disk partition formatted with the FAT16 file system. The Linux installation program will reformat the disk for Linux or you can use Linux commands such as fdisk and mke2fs to format the disk for Linux.

Virtual CD is needed to support the installation of Linux and is only provided by a server. By default, a Linux partition can see all CD drives on the logical partition. You can change an option on the NWS D to restrict Linux from accessing some or all of those drives. See “Creating a network server description and a network server storage space” on page 16 for more information on changing NWS D attributes.

Virtual tape provides access to the i5/OS tape drives from a logical partition. By default, a logical partition can see all tape drives on the hosted partition. You can change an option on the NWSD to restrict Linux from accessing some or all of those drives. See “Creating a network server description and a network server storage space” on page 16 for more information on changing NWSD attributes.

Virtual Ethernet provides the same function as using a 1 Gigabyte (GB) Ethernet adapter. A logical partition can use virtual Ethernet to establish multiple high speed inter-partition connections. i5/OS and Linux partitions can communicate with each other using TCP/IP over the virtual Ethernet communication ports. Refer to “Using virtual Ethernet in a logical partition” on page 23 for more information on virtual Ethernet.

For additional information about how your company might use logical partition with virtual I/O resources, see Logical and guest partition scenarios.

Directly attached I/O for Linux running in logical partitions on System i models 8xx and 270

With directly attached I/O, Linux manages the hardware resources directly and all I/O resources are under the control of the Linux operating system.

A directly attached I/O allows you to allocate disk units, tape devices, optical devices, and LAN adapters to a logical partition running Linux. You must have a network server description (NWSD) object to install Linux in a logical partition. Once Linux is installed, you can configure the partition to start independently. For directly attached hardware, all failure and diagnostic messages will be displayed within the logical partition. For a better understanding of how you might use a logical partition with directly attached I/O resources, read Logical and guest partition scenarios.

Contact your IBM marketing representative or IBM Business Partner for further assistance on using directly attached I/O in a logical partition.

System i I/O adapters (IOAs) supported by Linux

IOAs can be assigned to a logical partition. System i models 8xx and 270 support the following adapters in a logical partition running Linux.

The Linux *acenic*, *olympic*, and *pcnet32* device drivers are available with the Linux operating system. These drivers are compiled into the kernel and are available from your preferred Linux distributor.

The *ibmsis* driver provides an interface for the operating system to communicate with certain tape, optical, and disk adapter controllers. The most current version of this device driver can be found at the

Linux on the System i platform  Web site. All new tape, optical, and disk adapters are controlled using the open-source *ipr* device driver. This device driver is available from your preferred Linux

distributor as a kernel module. Additionally, the *ipr* device driver has a site available at: SourceForge 

The *icom* driver is an open-source device driver available with Linux. This driver can be compiled into the kernel and may be available from your preferred Linux distributor. To order the WAN adapters CCIN 2745 or 2772, it may be necessary to order hardware RPQs 847141, 847142, or 847143 when the order feature codes 0608, 0609, or 0610 are not available. WAN adapter CCIN 2771 is a base adapter for the server and can be used with Linux, but it is not a separately orderable feature.

The *lpfc* driver provides an interface for the operating system to communicate with disk unit sub-servers across fibre channel. The most current version of this device driver can be found at the Linux on the

System i platform  Web site. Supported configurations include direct attach, managed hubs, and

switched fabric topologies. To order the fibre channel adapter CCIN 2766, it may be necessary to order hardware RPQ 847145 when the order feature code 0612 is not available.

The e1000 driver is available from Intel™. This device driver can be compiled into the kernel and may be available from your preferred Linux distributor. The most current version of this device can be found at the <http://support.intel.com/support/network/adapter/1000/index.htm>  Web site.

This list may change between i5/OS release boundaries. Refer to the Linux on the System i platform  Web site.

Adapters (CCIN)	Feature codes	Description	Linux device driver	SubSystem vendor/ SubSystem type	PCI vendor ID/PCI ID
2742	0613	2-line WAN with 2 RVX ports supporting serial communication	icom	1014/021A	1014/0219
2743	0601	1-Gbps Ethernet (fibreoptical)	acenic	1014/0104	12AE/0001
2744	0603	100/16/4-Mbps Token Ring	olympic	1014/016D	1014/003E
2745	0608	2-line WAN with 2 RVX ports supporting multiple asynchronous communication	icom	0000/0000	1014/0031
2748	0605	Ultra-2 SCSI (3 buses; 26-MB write cache; extended adaptive read cache; RAID 5*; supporting internal disk and tape units, CD-ROM, and DVD-RAM)	ibmsis	1014/0099	1014/0096
2757	0618	Ultra-3 SCSI (4 buses; 757-MB write cache; hardware RAID 5**; supporting internal disk and tape units, CD-ROM, DVD-RAM, and DVD-ROM)	ibmsis	1014/0241	1014/0180
2760	0602	1-Gbps Ethernet UTP (unshielded twisted pair)	acenic	1014/00F2	12AE/0001
2763	0604	Ultra-2 SCSI (2 buses; 10-MB write cache; RAID 5*; supporting internal disk and tape units, CD-ROM, and DVD-RAM)	ibmsis	1014/0098	1014/0096

Adapters (CCIN)	Feature codes	Description	Linux device driver	SubSystem vendor/ SubSystem type	PCI vendor ID/PCI ID
2766	0612	Shortwave fibre channel adapter for point-to-point or arbitrated loop topologies; supporting tape and disk attachment in i5/OS(R); supporting disk for Linux	lpfc	10DF/F900	10DF/F900
2771		2-line WAN with 1 RVX port supporting asynchronous communication and 1 RJ11 modem port supporting V.90 56k	icom	0000/0000	1014/0031
2772	0609/0610	2-line WAN with 2 RJ11 modem ports supporting V.90 56k; Available in non-CIM (Complex Impedence Matching) and CIM packages.	icom	0000/0000	1014/0031
2778	0606	Ultra-2 SCSI (3 buses; 78-MB write cache; RAID 5*; extended adaptive read cache; supporting internal disk and tape units, CD-ROM, and DVD-RAM)	ibmsis	1014/0097	1014/0096
2780		Ultra-3 SCSI (4 buses; 757-MB write cache; 1-GB read cache; hardware RAID 0, 5, and 10; supporting internal disk and tape units, CD-ROM, DVD-RAM, and DVD-ROM)	ipr	1014/0264	1014/0180
2782	0619	Ultra-3 SCSI (2 buses; 40-MB write cache; hardware RAID 5**; supporting internal disk and tape units, CD-ROM, DVD-RAM, and DVD-ROM)	ibmsis	1014/0242	1069/B166

Adapters (CCIN)	Feature codes	Description	Linux device driver	SubSystem vendor/ SubSystem type	PCI vendor ID/PCI ID
2793	0614/0615	2-line WAN with 1 RVX port supporting multiple protocols and 1 RJ11 modem port supporting V.92 56k V.92 data modem, V.44 data compression, V.34 FAX modem and FAX functions; Available in non-CIM (Complex Impedence Matching) and CIM packages. Feature code 0165 is a CIM package.	icom	1014/0251	1014/0219
2805	0616/0617	4-line WAN with 4 RJ11 modem ports supporting V.92 56k and V.34 FAX functions; Available in non-CIM (Complex Impedence Matching) and CIM packages.	icom	1014/0252	1014/0219
2838	0607	100/10-Mbps Ethernet	pcnet32	1014/0133	1022/2000
2849	0624	100/10-Mbps Ethernet (unshielded twisted pair)	pcnet32	1014/024C	1022/2000
5700	0620	1000/100/10-Mbps Ethernet (fibreoptical)	e1000	8086/1009	8086/1009
5701	0621	100/100/10-Mbps Ethernet (unshielded twisted pair)	e1000	8086/100F	8086/100F
5702	0624	Ultra-3 SCSI (2 buses; no write cache; no hardware RAID; supporting internal disk and internal or external tape units, CD-ROM, and DVD-RAM)	ipr	1014/0266	1069/B166
5703		Ultra-3 SCSI (2 buses; 40MB write cache; hardware RAID 0, 5, and 10; supporting internal disk and tape units, CD-ROM, DVD-RAM, and DVD-ROM)	ipr	1014/0278	1069/B166

Directly attached SCSI support for a logical partition

Directly attached SCSI drives can be configured in a variety of methods.

Directly attached SCSI disk, tape, and optical devices are controlled by either the `ibmsis` device driver or the `ipr` device driver, depending on the storage adapter type. Refer to the I/O device table to determine which device driver your storage adapter utilizes. Refer to the PCI-X SCSI RAID Controller Reference Guide for Linux (Doc SA23-1327) for full documentation for both the `ibmsis` and `ipr` device drivers. This document contains information on how to configure directly attached disks, configure raid arrays, and perform error recovery and service.

Using the option to Display Disk Hardware Status in `sisconfig`, you see a model type for each directly attached SCSI device. You can determine if Linux has access to the disk by model number. If Linux does not have access to the disk, `sisconfig` can be used to resolve common disk errors. If `sisconfig` does not make the disk accessible to the logical partition, contact your next level of support.

Directly attached SCSI drives in a logical partition displays the following information:

Model	Description	Sisconfig task
020	Disk is unprotected and uncompressed.	Sisconfig is not needed.
050	Disk is unprotected and uncompressed.	Sisconfig is not needed.
060	Disk is compressed.	Disk should be initialized and formatted.
07x	Disk is parity protected.	Sisconfig is not needed.
08x	Disk is compressed and parity protected.	Disk should be initialized and formatted. Device parity protection should also be stopped.

Additional technical information on how to use the `ibmsis` utilities is included in a readme file with each download.

Certain Linux drivers are required before you can use the directly attached SCSI devices. Before you use `ibmsis`, verify you have the following Linux drivers:

Description	Module name
Low-level SCSI device driver	<code>ibmsis</code>
SCSI disk driver in base Linux kernel	<code>sd.o</code>
SCSI tape driver in base Linux kernel	<code>st.o</code>
SCSI CD-ROM driver in base Linux kernel	<code>sr.o</code>
SCSI driver in base Linux kernel	<code>scsi_mod.o</code>

Device parity protection for Linux running in a logical partition

Device parity protection is a hardware availability function that protects data from being lost because of a disk unit failure or because of damage to a disk.

For further information, refer to “Starting device parity protection and creating a disk array” on page 20.

Obtaining Linux for System i products

Linux is an open source operating system. Linux can be obtained in source format and built by individuals or business organizations.

The Linux open source code encourages feedback and further development by programmers. Linux developers are encouraged to design their own specialized distribution of the operating system to meet their specific needs and to freely offer their source code to the Linux community for potential inclusion.

All Linux distributions share a similar Linux kernel and development library. Linux distributors provide custom components that ease the installation and maintenance of Linux servers. Before installing another distributor's version of Linux, verify that the kernel has been compiled for the Power PC® and the System i hardware. Otherwise, your system might be incorrectly configured and will not run Linux in a logical partition.

Many of the distributors will provide a download via the Internet or a CD-ROM version of Linux. Refer to Linux on the System i platform  for links to each distributor.

Ordering or upgrading a System i product to run Linux in a logical partition

The IBM System Planning Tool emulates an LPAR configuration, validates that the planned partitions are valid, and allows you to test the placement of i5/OS and Linux hardware within the system.

Refer to IBM System Planning Tool  for information about the IBM System Planning Tool.

Contact your IBM marketing representative or IBM Business Partner to enter the order. You can enter the order by using the System i configurator. The configurator has been enhanced to support ordering IOAs without IOPs when a Linux partition is being defined.

When placing the order for a server with resources that support Linux in a logical partition, specify feature code 0142.

Creating a logical partition to run Linux

Learn the necessary steps for installing Linux on the System i platform.

Before you start to configure a partition on a System i product, with V5R1 or later i5/OS not managed by an HMC, to run Linux, see “Planning to run Linux in a logical partition” on page 2 for assistance. You should also be familiar with basic logical partitions concepts before you start to create a logical partition.

For information on partitioning for Linux on model 5xx systems, refer to Partitioning for Linux. It is recommended that you perform a full server back up before changing your server configuration. See “Backup options for virtual and directly attached disks” on page 41 for instructions on performing a full server back up.

Configuring a logical partition

A logical partition on a System i product can support the Linux operating system.

Note: To run Linux in a logical partition, you must have i5/OS V5R4 or later without a Hardware Management Console running in the primary partition.

Prerequisites

“Backup options for virtual and directly attached disks” on page 41

Before creating new partitions on the server, refer to instructions on performing a full server back up.

Configuring the service tools server

In order to access logical partition functions, you must first configure the service tools server.

Installing System i Navigator

System i Navigator provides a graphical wizard to help create logical partitions. It comes packaged with System i Access for Windows® and does not require a System i Access license to use. You will find detailed instructions and explanations in System i Navigator help topics.

To create a logical partition, follow these steps:

1. In System i Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the server.
3. Expand **Configuration and Service** and select **Logical Partitions**.
4. Right-click the physical server and select **Configure Partitions**.
5. From the Configure Logical Partition display, right-click the **Physical System** and select **New Logical Partitions** to start the wizard.

For information on software release function, refer to i5/OS logical partition function by release.

For information on HMC-managed servers, refer to Partitioning the server.

Assigning I/O adapters (IOAs) to a logical partition

Learn how to assign your I/O adapters to logical partitions.

Before changing your hardware configuration, use the IBM System Planning Tool. This tool allows you to test the placement of the i5/OS operating system and logical partition hardware within the server to ensure that the placement is valid. The output generated from the IBM System Planning Tool will validate the placement of hardware and ensure that your logical partition is configured properly. Refer to IBM System Planning Tool  for information on the IBM System Planning Tool.

Note: Incorrect placement of adapters for logical partitions could result in unpredictable behavior in the logical partition.

If you want to use real hardware devices like communication lines, disk units, internal tape and CD devices in a logical partition, you must assign the controlling IOAs explicitly to the logical partition. These instructions do not apply to virtual devices like virtual Ethernet, virtual disk, virtual CD, or virtual tape provided by a hosting i5/OS partition.

The logic to identifying IOAs is contained in the IOPs to which they are attached. For logical partitions, IOAs are not connected to IOPs so their identity may not be determined. In fact, card positions used by logical partitions will have one of these descriptions:

Empty position: IOA not yet installed.

Occupied position: IOA installed, but not yet assigned to a logical partition or not supported by the logical partition to which it is assigned.

Communication IOA: resource is a communication adapter and supported by the logical partition to which it is assigned.

Multiple Function IOA: resource is a multiple function adapter and supported by the logical partition to which it is assigned.

Disk Controller: resource is a disk controller and supported by the logical partition to which it is assigned.

Generic adapter: IOA installed and may be supported by the logical partition to which it is assigned.

Storage IOA: resource is a storage adapter and supported by the logical partition to which it is assigned.

Due to this non-specific identification, it is very important to note the logical and physical address of the card positions to be used by a logical partition. A logical partition must have access to the bus on which its IOAs reside. The only restriction is that if the bus is shared with ani5/OS partition, the i5/OS partition must own it, shared (own shared) and the logical partition must use it shared (use shared).

A logical partition must have access to the bus on which its IOAs reside. The only restriction is that if the bus is shared with ani5/OS partition, the i5/OS partition must own it, shared (own shared) and the logical partition must use it shared (use shared).

Changing bus ownership type (for shared bus ownership)

To change the ownership for a bus, follow these steps:

1. In System i Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the server.
3. Expand **Configuration and Service** and select **Logical Partitions**.
4. Right-click the **Logical Partition** and select **Configure Partitions**. You are now working in the Configure Logical Partitions window.
5. Right-click the bus for which you want to change ownership and select **Properties**.
6. Select the **Partitions** page.
7. Select the partition that owns the bus in **Owning logical partition**, and then select the ownership type in **Sharing**. If the ownership type is shared, the partitions that share the bus appear in the list. Click **Help** if you need more information on these options.
8. Click **OK**.

Moving an unassigned IOA to a logical partition

To move an unassigned IOA to a logical partition, follow these steps:

1. In System i Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the server.
3. Expand **Configuration and Service** and select **Logical Partitions**.
4. Right-click the **Logical Partition** and select **Configure Partitions**. You are now working in the Configure Logical Partitions window.
5. Right-click the IOA you want move and select **Move**.
6. Select the logical partition to receive the IOA in **Move to -- Logical partition**.
7. Click **OK**. The IOA you specified is now moved.

Moving an unassigned IOA to an i5/OS partition

This requires the use of the Hardware Service Manager. If you are not familiar with this tool, contact your hardware service representative for assistance. This method should only be used by users who are familiar with the use of concurrent maintenance. Improper use of this method could cause the incorrect configuration of devices.

Note: This procedure should not run if any hardware is marked as failed. It should only be run when all server hardware is completely operational.

To move an unassigned IOA to an i5/OS partition, follow these steps:

1. In System i Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the server.
3. Expand **Configuration and Service** and select **Logical Partitions**.

4. Right-click the **Logical Partition** and select **Configure Partitions**. You are now working in the **Configure Logical Partitions** window.
5. Right-click the IOA you want move and select **Move**.
6. Select the logical partition to receive the IOA in **Move to -- Logical partition**.
7. Click **OK** to move the IOA you specified.
8. Verify the IOP belongs to the desired partition. If the IOP does not belong to the desired partition, move the IOP.
9. At the i5/OS command line type STRSST and press **Enter**.
10. From the Start Service Tools (STRSST) Sign On display, key a Service Tools user ID and password and press **Enter**.
11. From the System Service Tools (SST) display, select option **1** (Start a service tool) and press **Enter**.
12. From the Start a Service Tool display, select option **7** (Hardware service manager) and press **Enter**.
13. From the Packaging Hardware Resources display, select option **9** (Hardware contained within package) next to the System Unit or System Expansion Unit that contains the IOA and press **Enter**.
14. Select option **5** (Display Detail) next to the desired IOA and press **Enter**. Note the Resource name, Frame ID, and Card position.
15. Return to the Packaging Hardware Resource display, select option **3** (Concurrent maintenance) next to the IOA that you want to make available and press **Enter**.
16. From the Hardware Resource Concurrent Maintenance display, select option **9** (Work with controlling resource) next to the IOA selected and press **Enter**.
17. From the Controlling resource display, select option **7** (Assign to) next to the IOA and press **Enter**. Verify type/model, serial number and part number of the IOA.
18. From the Confirm Add I/O resources display, verify the desired IOA has been added by verifying the logical address and press **Enter**.
19. Select **View** to refresh the System i Navigator interface.
20. Right-click **Physical System** and select **Recovery** and **Clear non-reporting partitioning resources**.

Moving an IOA assigned to an IOP to a logical partition

This action should be performed when the IOP can be removed from the logical partition.

To move an IOA assigned to an IOP to a logical partition, follow these steps:

1. Ensure that any devices attached to the I/O processor you want to move are not busy. The devices should be varied off and listed as unavailable hardware.
2. In System i Navigator, expand **My Connections** or your active environment.
3. Select the primary partition of the server.
4. Expand **Configuration and Service** and select **Logical Partitions**.
5. Right-click the Logical Partition and select **Configure Partitions**. You are now working in the **Configure Logical Partitions** window.
6. Select the partition that has the I/O processor you would like to move.
7. Right-click the I/O processor you want move and select **Move**.
8. Select the **Remove hardware resource information from source partition** box.
9. Select **Unassigned Hardware in Move to -- Logical Partition** to receive the IOP.
10. Click **OK** to move the IOP you specified.
11. Verify IOP is identified as Unassigned Hardware.
12. Right-click the IOA you want to move and select **Move**.
13. Select the logical partition to receive the IOA in **Move to -- Logical Partition**.
14. Click **OK** to move the IOA you specified.
15. Repeat steps 12-14 to move multiple IOAs.

16. Move the IOP back to the logical partition that owned it.

Assigning an IOA assigned to an IOP to a logical partition

This action should be performed when the IOP cannot be removed from the logical partition for a Linux server.

This requires the use of the Hardware Service Manager. If you are not familiar with this tool, contact your hardware service representative for assistance. This method should only be used by users who are familiar with the use of concurrent maintenance. Improper use of this method could cause the incorrect configuration of devices.

1. At the i5/OS command line of the partition that owns the IOP to which the desired IOA is assigned, type STRSST and press **Enter**.
2. From the Start Service Tools (STRSST) Sign On display, key a Service Tools user ID and password and press **Enter**.

Note: The Service Tools user must have Administration authority.

3. From the System Service Tools (SST) display, select option **1** (Start a service tool) and press **Enter**.
4. From the Start a Service Tool display, select option **7** (Hardware service manager) and press **Enter**.
5. From the Hardware Service Manager display, select option **1** (Packaging hardware resources) and press **Enter**.
6. From the Packaging Hardware Resources display, select option **9** (Hardware contained within package) next to the System Unit or System Expansion Unit that contains the IOA and press **Enter**.
7. Select option **5** (Display Detail) next to the desired communication port and press **Enter**. Note the Resource name, Frame ID, and Card position.
8. Return to the Packaging Hardware Resource display, select option **3** (Concurrent maintenance) next to the communication port that you want to make available and press **Enter**.
9. Confirm the desired IOP/IOA connection is being relinquished by verifying the resource name and press **Enter**. The IOA is now no longer controlled by the IOP so it is identified as an "Occupied position." Subsequent steps are done using LPAR configuration screens.
10. At the i5/OS command line of the primary partition, type STRSST and press **Enter**.
11. From the Start Service Tools (STRSST) Sign On display, key a Service tools user ID and password and press **Enter**.

Note: LPAR Administration authority is needed to change the LPAR configuration.

12. From the System Service Tools (SST) display, select option **5** (Work with servers) and press **Enter**.
13. From the Work with System Partitions display, select option **1** (Display partition information).
14. From the Display Partition Information, select option **5** (Display server I/O resources).
15. From the Display System I/O Resources display, press **F10** twice to see physical addresses (Frame ID and Card positions).
16. Locate the desired Occupied position using the Frame ID and card position. Press **F10** twice to see ownership information and note it. The bus ownership should be shared. Press **F10** three more times to see the associated logical address and note it.
17. Return to the Work with System Partitions display, select option **3** (Work with partition configuration) and press **Enter**.
18. From the Work with Partition Configuration display, select option **4** (Remove I/O resources) next to the partition that owns the desired Occupied position.

Note: This will be the owner of the bus on which it appears.

19. From the Remove I/O Resources display, select option **2** (Remove and clear hardware resources) next to the desired Occupied position and press **Enter**. The bus ownership should be shared.

20. From the Confirm Remove I/O Resources display, confirm the desired Occupied position has been removed by verifying the logical address and press **Enter**.
21. Return to the Work with Partition Configuration display, select option 3 (Add I/O resources) next to the logical partition to which you wish to add the Occupied position and press **Enter**.
22. From the Add I/O Resources display, select option 1 (Own dedicated) next to the Occupied position to be assigned to the logical partition. If the logical partition does not have access to the bus, select option 3 (Use bus shared) and a 1 next to the **Occupied** position.
23. From the Confirm Add I/O resources display, verify the desired IOA has been added by verifying the logical address and press **Enter**.
24. From the System Service Tools (SST) display, select option 5 (Work with servers) and press **Enter**.
25. From the Work with System Partitions display, select option 4 (Recover configuration data) and press **Enter**.
26. From the Recover Configuration Data display, select option 4 (Clear non-reporting logical partitioning resources) and press **Enter**.

Refer to Linux on the System i platform  (www.ibm.com/systems/i/os/linux/) for the latest updates to Linux running on an System i product.

Creating a network server description and a network server storage space

A network server description (NWS D) is used to give a name to the configuration, provide an interface for starting and stopping a Linux partition, and provide a link between the Linux operating system and its virtual disks.

The default or suggested parameter values are provided within the parentheses. These settings are only relevant to a logical partition. Refer to NWS D parameter and description if you need any information on parameter descriptions.

The following steps will create a virtual disk for a logical partition running Linux:

1. At the i5/OS command line on the host partition, type CRTNWS D and press **F4** for prompts.
2. From the Create Network Server Description display, provide the following information:
 - NWS D (Provide a name for the NWS D)
 - RSRCNAME (*NONE). TYPE (*GUEST)
 - ONLINE (*NO or *YES)
 - PARTITION (Provide name of the Linux partition)
 - CODEPAGE (437)
 - TCPPORTCFG (*NONE)
 - RSTDDEVRS C (for virtual CD and tape devices) (*NONE)
 - SYNCTIME (*TYPE)
 - IPLSRC (*NWSSTG)

Attention: To start an NWS D with a kernel from a stream file, set the IPLSRC parameter to *STMF and set the IPLPATH parameter to point to the kernel. This will only load the kernel. Once the kernel is running, it will need to find a root file system. In an initial installation, the root file system might be a RAM disk that is physically attached to the kernel.

It is possible to store a kernel in a partition of an NWSSTG and start from there. The NWSSTG must have a partition type "PReP Boot" (type0x41) marked as startable.

IPLSTMF (*NONE)

Attention: This is the file to load from if you specify that the kernel start from a stream file. You must have read access to the file and the path leading to the file to use the vary on command.

IPLPARM (*NONE)

3. From the Work with Network Server Storage Spaces (WRKNWSSTG) select option 1 (Create Network Server Storage Space), press **Enter**.
4. Provide the following information:
 - NWSSTG (Name)
 - NWSSIZE (Refer to the installation documentation of your preferred Linux distributor)
 - FROMNWSSTG (*NONE)
 - FORMAT (*OPEN)
5. Press **Enter**.

Attention: A network server storage space can be as large as 64 GB for a single storage space. A maximum of 48 storage spaces can be linked to a server.
6. Select option 10 (Add link) next to the storage space that you just created and press **Enter**, Add Server Storage Link (ADDNWSSTGL), and provide the following information:
 - NWSSTG (Name)
 - NWSD (Name)
 - DYNAMIC (*YES)
 - DRVSEQNBR (*CALC)

Configuring the LAN console for Linux running a logical partition

Operations Console is an installable component of System i Access for Windows. It allows you to use your PC as the server console, perform control panel functions, or both.

With Operations Console, you can administer System i servers that are on your local area network (LAN). You can also remotely administer several System i products at different locations, including other LANs. All traffic between the console and the System i product is encrypted to improve the security of the server.

Once a LAN based operations console has been set up, the same connection can be used to provide the Linux console. This has the advantage that all Linux console traffic will be encrypted between the System i product and the console.

To use the Operations Console LAN connection to provide the Linux console, use the same Telnet commands as described in Connect to the virtual console, however the commands must be run on the PC running the Operations Console, and a TCP/IP address of 127.0.0.1 and port 2301 should be used with the Telnet command. 127.0.0.1 is architected by TCP/IP to be the local server. This will cause the Telnet program to connect to the Operations Console program, which will in turn connect the Telnet client to the Linux console. See "Connecting the virtual console" for more information.

If the Operations Console program is connected to more than one System i product, it is currently not possible to specify which server the Linux console connects to. It is recommended, therefore, that if Linux console support is being used, that the Operations Console client be connected to only one System i product at a time.

Using the Operations Console program to connect to the Linux console is an alternative to using a direct Telnet connection to the System i product.

Connecting the virtual console

The virtual console provides the console function for the Linux server.

The virtual console is used primarily during the initial install of Linux. The virtual console can also be used to view server errors or to restore communication to the LAN. This console connection is used prior to configuring TCP/IP.

Any Telnet client can be used as the Linux console. Multiple Telnet clients can share access to the same virtual console. To connect to a console, you will have to Telnet to port 2301 of the primary or server. TCP/IP must be configured and running on at least one i5/OS partition. The example below uses the IBM Personal Communications client.

Use either of the following methods to do this:

Connecting to the virtual console with IBM Personal Communications

To connect to a virtual console with IBM Personal Communications, follow these steps:

1. At the Start button, select **IBM Personal Communications** and **Start** or **Configure Session**.
2. From the Customize Communication display, select **ASCII** as your type of host and select **Link Parameters**.
3. From the Telnet ASCII display, enter the host name or the IP address of the primary or server, enter port number 2301 of the primary or hosting partition and press **OK**.
4. From the i5/OS Guest Partition Console display, select the partition to which you want to connect as console.
5. Enter the i5/OS service tools ID and password to connect to the Linux partition.

Connecting to the virtual console using the MSDOS command prompt:

To connect to the virtual console using the MSDOS command prompt, follow these steps:

1. From a MS DOS Command Prompt, Telnet to your server and port 2301 (*telnet xxxxxx 2301*).
2. Select the partition to which you want to connect as console.
3. Enter the i5/OS service tools ID and password to connect to the Linux partition.

Remote Panel privilege for the Linux partition is required to use the Linux virtual console. The QSECOFR service tools user ID is prevented from using the Linux console.

For more information on how to configure user profiles, refer to Logical partition authority.

Installing Linux in a logical partition

Configure your System i product before you install the Linux operating system.

See “Configuring a logical partition” on page 11 for more information on configuring Linux in a logical partition.

See “Installing Linux from CD-ROM images in the integrated file system” for information on installing from CD images.

IBM is currently working with IBM Business Partners who support Linux on System i products. Your preferred Linux distributor will provide specific documentation on how to install Linux on the System i products. Refer to the Linux on the System i platform Web site for links to your preferred distributor.

Installing Linux from CD-ROM images in the integrated file system

Most Linux distributors provide ISO CD-ROM images that can be downloaded from their Web site. If your distributor has multiple physical CDs for the install action of Linux, ISO file images are a convenient way to install from a directory on your System i product.

The ISO images used to install Linux will have one of these descriptions:

- *Unloaded*: This option allows you to unload an optical image from the virtual optical device. Only images that have a status of Loaded can be unloaded.
- *Loaded*: This option allows you to load an optical image into the virtual optical device. Only images that have a status of unloaded can be loaded.
- *Mount*: This option allows you to mount an optical image into the virtual optical device and activate it. Only images that have a status of loaded can be mounted.

To install from CD-ROM images stored in the integrated file system, follow these steps:

Attention: The following commands are performed during the initial setup only.

1. At the i5/OS command line, enter the command CRTDEVOPT and press **Enter**.
2. From the Create Device Description (Optical) display, provide the following information and press **Enter**.
 - Device description (Provide a name for the new device description.)
 - Resource name (*VRT)
 - Device type (*RSRCNAME)
3. At the command line, enter the command CRTIMGCLG and press **Enter**.
4. From the Create Image Catalog display, provide the following information and press **Enter**.
 - Image catalog name Directory name
5. At the command line, enter the command WRKIMGCLGE and press **Enter**.
6. From the Work with Image Catalog Entries display select **Add** (option 1) to add an optical image to the image catalog and press **Enter**.

Attention: The following commands need to be performed to use the ISO images.

7. At the command line, enter the command WRKIMGCLGE and press **Enter**.
8. From the Work with Image Catalog Entries display, do the following:
 - a. Select Mount (option 6) on the first CD image.
 - b. Vary on the device and press **Enter**. At the command line, enter the command VRYCFG and press **Enter**.
 - c. From the Vary on the device display, provide the following information:
 - Virtual Optical Device Description
 - Config Type of *DEV
 - Turn the status *ON
9. At the command line, enter the command LODIMGCLG and press **Enter**.
10. From the Load or Unload Image Catalog display, provide the following information and press **Enter**.
 - Image catalog (Specifies the name of the image catalog to be loaded into from the virtual optical device.)
 - Device name (Specifies the name of the virtual optical device that the image catalog is to be loaded into.)
11. At the command line, enter the command WRKOPTVOL and press **Enter**.
12. From the Work with Optical Volumes display, verify the information is correct.

Attention: The following commands will set up the network server description in order to install from the ISO images in IFS.

13. At the command line, enter the command WRKCFGSTS and press **Enter**.
14. From The Work with Configuration Status display, provide the following information: Installation will start. When installation asks for CD #2, then do the following steps.

Attention: The following commands will need to be performed to change the ISO images.

- WRKCFGSTS *NWS
- 8 Network Server Description

- 2 Update
- IPL source = *STMF
- IPL stream file = /qopt/path to start image (this is the same as if you were accessing from CD-ROM)
- IPL parameters = *NONE WRKCFGSTS *NWS
- 1 to vary on the partition

Installation will start. When installation asks for CD #2, then do the following steps.

Attention: The following commands will need to be performed to change the ISO images.

15. At the command line, enter the command `WRKIMGCLGE` and press **Enter**.
16. From the Work with Image Catalog Entries display, provide the following information (You will repeat this step until all CDs have been installed):
 - Unload (option 9) on mounted CD
 - Mount (option 6) the next CD

Starting and stopping the NWSD for a Linux partition

You will have to stop and start (vary off and vary on) the network server description (NWSD) to IPL a logical partition running Linux.

Use either of the following methods to do this:

Stop the NWSD with System i Navigator

To stop the NWSD, do the following:

1. Click **Integrated Server Administration** → **Servers**.
2. Right-click the name of the NWSD that you want to stop.
3. Click **Shut Down**.

Starting the NWSD with System i Navigator

To start the NWSD, do the following:

1. Click **Integrated Server Administration** → **Servers**.
2. Right-click the name of the NWSD that you want to start.
3. Click **Start**.

Stopping the NWSD with CL commands

To stop the NWSD using the Work with configuration status (WRKCFGSTS) CL command, do the following:

1. Type `WRKCFGSTS *NWS` and press **Enter**.
2. Type 2 next to the NWSD that you want to stop and press **Enter**.

Starting the NWSD with CL commands

To start the NWSD using the Work with configuration status (WRKCFGSTS) CL command, do the following:

1. Type `WRKCFGSTS *NWS` and press **Enter**.
2. Type 1 next to the NWSD that you want to start and press **Enter**.

Starting device parity protection and creating a disk array

Start device parity protection/create a disk array when you attach a new disk subserver to your server.

You must start device parity protection/create a disk array before you configure the associated disk units for use under Linux because all data is lost as a result of the action. If you want to install on to a parity

set/disk array, start the logical partition into rescue mode prior to the installation, and follow the steps outlined in the PCI-X SCSI RAID Controller Reference Guide for Linux (doc SA23-1327) to start device parity/create a disk array.

Managing Linux in a logical partition

This topic contains information about managing a partition running Linux.

You will learn how to start a partition running Linux and how to communicate and share information between a logical partition and an i5/OS partition. Additional information on managing a partition can be found at Managing logical partitions. Refer to your Linux distributor for Linux specific task management.

For information on HMC-managed partitions on model 5xx systems, refer to Partitioning the server.

Authority for Linux running in logical partitions

- | Use this information to determine the authorities needed to do tasks for Linux servers running in logical partitions.

Two service tool functional privileges relate to logical partitions. These privileges support basic operations or advanced administration. For more information on obtaining operations or administration authority, refer to Logical partition authority.

The following table describes which authority is required to complete a logical partition task.

Table 10. Logical partition authority. This table describes what authority is required to complete a logical partition task.

Function	Administration authority	Operational authority
Back up NWSD for a partition	X	
Create a logical partition	X	
Change logical partition host information	X	
Delete disk drives for a logical partition	X	
Delete NWSDs for a logical partition	X	
Display operating environment of secondary partitions	X	X
Display reference code history for secondary partitions	X	X
Print server configuration for logical partitions	X	X
Restore NWSDs for logical partitions	X	
Save logical partition configuration data	X	X
Unlink disk drives from a logical partition	X	
Vary off NSWD	X	
Vary on NSWD	X	

Displaying and changing hardware configuration

Use these tasks to display and change your server's hardware configuration.

Displaying the console log for Linux running in a logical partition

Use the display console log for the logical partition to see console information for the logical partition running Linux.

You can perform this procedure from server service tools (SST) or dedicated service tools (DST) on the primary partition. To use the Display Guest Environment Console Log display, follow these steps:

1. On the primary partition, start SST or DST.
2. From SST select option **5** (Work with servers) and press **Enter**.
3. From DST select option **11** (Work with servers) and press **Enter**.
4. Select option **1** (Display partition information) and press **Enter**.
5. Select option **10** (Display logical environment console log) and press **Enter**.

Once you are in Display Guest Environment Console Log, you can see console information for the logical partition.

Displaying the operating environment of secondary partitions

You can use the display partition operating environment display to see which release of the operating system each logical partition is running.

To display the operating system type for each partition, follow these steps:

1. In System i Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the server.
3. Expand **Configuration and Service** and select **Logical Partitions**.
4. Right-click the partition and select **Properties**.
5. Select the **General** page to view the operating system. Click **Help** if you need more information on this field.
6. Click **OK**.

Displaying the reference code history for secondary partitions

A reference code indicates a status or an error condition.

The server records the reference code history (the last 200 reference codes) for secondary partitions.

To display server reference code history for a logical partition, follow these steps:

1. In System i Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the server.
3. Expand **Configuration and Service** and select **Logical Partitions**.
4. Right-click the **Logical Partition** and select **Properties**.
5. Select the **Reference Code** page to view a list of the 200 most recent server reference codes. Click **Help** if you need more information on this field.
6. Click **OK**.

Changing the logical partition host information

A hosted logical partition is dependent on an i5/OS partition for some or all I/O resources. The hosting i5/OS partition can either be a primary or a secondary partition.

To change the logical partition host information, follow these steps:

1. In System i Navigator, expand **My Connections** or your active environment.

2. Select the primary partition of the server.
3. Expand **Configuration and Service** and select **Logical Partitions**.
4. Right-click the logical partition and select **Properties**.
5. Select the **Environment** page to specify the name of an i5/OS partition that will act as the server. Click **Help** if you need more information on this field.
6. Click **OK**.

Connecting your logical partition to the LAN

You can connect a logical partition on your System i product to a LAN using TCP/IP.

You can either use directly attached LAN adapters to attach to the LAN or you can use virtual Ethernet and your server.

For more information related to connecting your logical partition to the LAN, refer to:

Directly attached LAN adapters

A logical partition running Linux can own its own LAN adapters.

Once you allocate the supported LAN adapters to the logical partition, i5/OS does not realize the hardware exist and cannot use the I/O resource. See “Assigning I/O adapters (IOAs) to a logical partition” on page 12 for more information.

You will need to refer to your Linux distributor’s documentation for information on network TCP/IP setup.

Virtual Ethernet

A logical partition can use virtual Ethernet to establish multiple high speed inter-partition connections.

Logical partition software allows you to configure up to 16 different virtual local area networks. Virtual Ethernet provides the same function as using a 1 gigabyte (GB) Ethernet adapter. i5/OS and Linux partitions can communicate with each other using TCP/IP over the virtual Ethernet communication ports.

For more information related to configuring virtual Ethernet, refer to:

Using virtual Ethernet in a logical partition:

A logical partition can use virtual Ethernet to establish multiple high speed inter-partition connections.

Logical partition software allows you to configure up to 16 different virtual local area networks. Virtual Ethernet provides the same function as using a 1 Gigabyte (GB) Ethernet adapter. i5/OS and Linux partitions can communicate with each other using TCP/IP over the virtual Ethernet communication ports.

The enablement and setup of virtual Ethernet requires a restart of the Linux partition, but does not require any special hardware or software. Once a specific virtual Ethernet is enabled for a logical partition, a network device vethXX is created in the logical partition. The user can then set up TCP/IP configuration appropriately to start communicating with another partition. Virtual Ethernet provides the capability to provide multiple communication paths between applications that are run in each of the partitions.

Refer to the Linux on the System i platform  Web site for the latest updates to Linux running on a System i product.

Enabling virtual Ethernet:

Virtual Ethernet provides the capability to provide multiple communication paths between applications that are run in each of the partitions.

The enablement and setup of virtual Ethernet requires a restart of the Linux partition, but does not require any special hardware or software. Once a specific virtual Ethernet is enabled for a logical partition, a network device vethXX is created in the logical partition. The user can then set up TCP/IP configuration appropriately to start communicating with another partition.

To enable virtual Ethernet, follow these steps:

1. In System i Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the server.
3. Expand **Configuration and Service** and select **Logical Partitions**.
4. Right-click **Properties** and select the **Virtual Ethernet** page. View the virtual Ethernet options for the logical partition. Click **Help** if you need more information on this field.
5. Click **OK**.

Create an Ethernet line description for virtual Ethernet:

Creating an Ethernet line description is the first step in configuring i5/OS to use virtual Ethernet.

This configuration allows a logical partition to communicate to an i5/OS partition using virtual Ethernet.

To configure a new Ethernet line description to support virtual Ethernet, complete the following steps:

1. At the i5/OS command line, type WRKHDWRSC *CMN and press **Enter**.
2. From the Work with Communication Resources display, select option 7 (Display Resource Detail) next to the appropriate virtual Ethernet port. The Ethernet port identified as a 268C is the virtual Ethernet resource. There will be one for each virtual Ethernet that is connected to the partition.
3. From the Display Resource Detail display, scroll down to find the port address. The port address corresponds to the virtual Ethernet you selected during the configuration of the partition.
4. From the Work with Communication Resources display, select option 5 (Work with configuration descriptions) next to the appropriate virtual Ethernet port and press **Enter**.
5. From the Work with Configuration Descriptions display, select option 1 (Create) enter the name of the line description and press **Enter**.
6. From the Create Line Description Ethernet (CRTLINETH) display, provide the following information:
 - RSRCTYPE
 - LINESPEED (1G)
 - DUPLEX (*FULL)

Press **Enter**. Make sure the maximum frame size is 8996. By changing the frame size to 8996, the transfer of data across the virtual Ethernet is improved.

7. Press **Enter**.


From the Work with Configuration Description display, you will see a message stating the line description has been created.

Virtual Ethernet connectivity methods:

Virtual Ethernet creates a high speed virtual Ethernet segment that can be used to connect logical partitions in a physical System i product.

This LAN segment is separate from any real LAN that the server might have a connection with. A virtual Ethernet is made up of a virtual line description and i5/OS TCP/IP interface. The Linux TCP/IP interface then has its own IP address but uses the virtual network device for its hardware. If a logical partition is

connected only to a virtual Ethernet segment, and you want your logical partition to communicate with servers on an external LAN, you will need to bridge TCP/IP traffic between the i5/OS external LAN and the virtual i5/OS LAN segment. A logical flow of the IP packets would look something like this (Pre-V5R3 diagram):



IP traffic initiated by the logical partition goes from the Linux network interface to the virtual i5/OS interface. If the virtual interface is associated with the external interface, the IP packets can continue on to the external interface and towards its destination.

There are three methods for bridging the external and virtual Ethernet segments. Each method has differences that make it useful based on your knowledge of TCP/IP and your environment. Choose from one of the following methods:

Proxy ARP:

The Proxy ARP method uses a technique commonly known as transparent subnetting.

While the following information will provide a procedure to create connectivity between your virtual and external networks, you might want to learn more information about transparent subnetting. Here are some useful links:


V4 TCP/IP for AS/400(R): More Cool Things Than Ever 

This pre-V5R3 IBM Redbooks™ publication provides sample scenarios that demonstrate common solutions with example configurations. It also helps you plan, install, tailor, configure, and troubleshoot TCP/IP on your System i product.

TCP/IP routing and workload balancing

This topic provides techniques and instructions for routing and workload balancing.

If you choose to use the Proxy ARP method, you must have a firm understanding of subnetting and TCP/IP. You will need to obtain a contiguous block of IP addresses that are routable by your network. You subnet this block of IP addresses and assign one to the virtual TCP/IP interface and one to the TCP/IP connection in your logical partition as described in the following pre-V5R3 figure.



In this example, a contiguous block of four IP addresses is used, which is 10.1.1.12 - 10.1.1.15. Since it is a block of four IP address, the subnet mask for these addresses is 255.255.255.252. To configure a setup similar to this, use the following instructions:

1. Obtain a contiguous block of IP addresses that are routable by your network. Since there is only one logical partition, you only need four IP addresses. The fourth segment of the first IP address in the block must be divisible by four. The first and last IP addresses of this block are the Subnet and broadcast IP addresses and are unusable. The second and third IP can be used for a virtual TCP/IP interface and a TCP/IP connection in your logical partition. For this example, the IP address block is 10.1.1.12 - 10.1.1.15 with a subnet mask of 255.255.255.252.

You also need a single IP address for your external TCP/IP address. This IP addresses does not have to belong to your block of contiguous addresses, but it must be within the same original subnet of addresses as your block. In the example, the external IP address is 10.1.1.5 with a subnet mask of 255.255.255.0.

2. Create an i5/OS TCP/IP interface for your server. In this example, you would assign the 10.1.1.5 as an IP address and the subnet mask would be 255.255.255.0.
3. "Create an Ethernet line description for virtual Ethernet" on page 24. Note the port of the hardware resource, because you will need it when you set up TCP/IP in your logical partition. In the case of the above example, assume the hardware resource is CMN05 with a port of 0 and although arbitrary, the line description name is VETH0. *CRTLINETH LIND(VETH0) RSRcname(CMN05) LINESPEED(1G) DUPLEX(*FULL)*
4. Add a TCP/IP Interface for your Ethernet line description. Following the example, you could use:
*ADDTCPIFC INTNETADR('10.1.1.13') LIND('VETH0') SUBNETMASK ('255.255.255.252')
LCLIFC('10.1.1.5')*

Attention: If the fourth segment of your virtual TCP/IP interface is greater than the fourth segment of your Proxy ARP IP address, then you need to set AUTOSTART to *NO on this interface.

i5/OS starts its TCP/IP interfaces in numerical order and you must start the proxy ARP interface before you start any virtual TCP/IP interfaces.

5. Turn on IP datagram forwarding. This allows the i5/OS TCP/IP interfaces to pass IP packets between each other. *CHGTCPA IPDTGFWD(*YES)*
6. Start your TCP/IP interfaces. You can use a command like *STRTCPIFC INTNETADR(yourIPaddress)*. For the example, you would issue the commands like:
 - *STRTCPIFC INETADDR('10.1.1.5')* // starts the external proxy ARP interface
 - *STRTCPIFC INETADDR('10.1.1.13')* // starts the virtual Ethernet interface


7. On your logical partition, set up networking using the instructions or tools provided by your Linux distribution. Be sure to use the correct IP address, subnet mask, port, and router IP address.
 During the TCP/IP setup in Linux, you will need to know your net or network device. If you are running a 32-bit kernel version of 2.4.9 or less, the network device is VETH plus the port number from when you configured an Ethernet line description. If you are running a 64-bit kernel or a 32-bit kernel greater than or equal to 2.4.10, the network device is ETH plus the port number. See “Create an Ethernet line description for virtual Ethernet” on page 24 for more information. In the example:
 - the interface IP address is 10.1.1.14
 - the subnet mask is 255.255.255.252
 - the gateway IP address is 10.1.1.13
 - the network device is VETH0 or ETH0 depending on the kernel version.
8. To verify network communications, ping the virtual Ethernet interface and a host on the external network from your logical partition. Then from i5/OS ping the virtual Ethernet interface and the Linux interface.

Network Address Translation (NAT):

NAT can route traffic between your logical partition and the external network using virtual Ethernet.

This particular form of NAT is called static NAT, and it will allow both inbound and outbound IP traffic to and from the logical partition. Other forms of NAT like masquerade NAT would also work if your logical partitions do not need to receive traffic initiated by external clients. Like the TCP/IP Routing and Proxy ARP methods, you can take advantage of your existing i5/OS network connection. Since you will be using IP Packet rules, you must use System i Navigator to create and apply your rules.

The following illustration is an example of using NAT to connect your logical partition to an external network. The 10.1.1.x network represents an external network while the 192.168.1.x network represents the virtual Ethernet LAN.



In this pre-V5R3 example, any existing TCP/IP traffic for the System i hosting partition runs over the 10.1.1.2 interface. Since this is a static map scenario, the inbound traffic gets translated from the 10.1.1.3 interface to 192.168.1.1 interface. The outbound traffic gets translated from the 192.168.1.1 interface to the external 10.1.1.3 interface. The logical partition uses the virtual interface (192.168.1.1) and its own 192.168.1.5 interface to communicate with the server.

To make static NAT work, you need to first set up your i5/OS and Linux TCP/IP communications. Then you will create and apply some IP Packet rules. Use the following procedure:

1. During the configuration of your logical partition, make sure you selected to create virtual Ethernet. This is described in *Configure a logical partition*.
2. “Create an Ethernet line description for virtual Ethernet” on page 24. Note the port of the hardware resource, because you will need it when you set up TCP/IP in your logical partition. In the case of the above example, assume the hardware resource is CMN05 with a port of 0 and although arbitrary, the line description name is VETH0. *CRTLINETH LIND(VETH0) RSRCCNAME(CMN05) LINESPEED(1G) DUPLEX(*FULL)*
3. Add a TCP/IP Interface for your virtual line description. Following the example, you could use: *ADDTCPIFC INTNETADR('192.168.1.1') LIND('VETH0') SUBNETMASK ('255.255.255.0')*
Start your new TCP/IP interface using the command *STRTCPIFC INTNETADR(yourIPaddress)* or *STRTCPIFC INTNETADR('192.168.1.1')*, for the example.
4. On your logical partition, set up networking using the instructions or tools provided by your Linux distribution. Be sure to use the correct IP address, subnet mask, port from step 2, and gateway IP address.
During the TCP/IP setup in Linux, you will need to know your net or network device. If you are running a 32-bit kernel version of 2.4.9 or less, the network device is VETH plus the port number from when you configured an Ethernet line description. If you are running a 64-bit kernel or a 32-bit kernel greater than or equal to 2.4.10, the network device is ETH plus the port number. See “Create an Ethernet line description for virtual Ethernet” on page 24 for more information. In the example:
 - the interface IP address is 192.168.1.5
 - the subnet mask is 255.255.255.0
 - the gateway IP address is 192.168.1.1
 - the network device is VETH0 or ETH0 depending on the kernel version
5. Create another TCP/IP interface that connects to the external network. It should use the same line description as your existing external TCP/IP interface. This interface will eventually perform the address translation for your partition. Check to make sure your new interface communicates with the external LAN properly. In the case of the example, this interface has an IP address of 10.1.1.3 on a line description called ETHLINE.
6. Start the external TCP/IP interface. *STRTCPIFC INTNETADR('10.1.1.3')*
7. You should now verify that your virtual Ethernet connection works. From the logical partition, ping the Linux gateway and from i5/OS, ping the logical partition. If the ping is successful, continue.
8. Turn on IP datagram forwarding. This allows the i5/OS TCP/IP interfaces to pass IP packets between each other. *CHGTCPA IPDTGFWD(*YES)*
9. Connect to the server with System i Navigator. You must not connect to the server with the NAT interface that you just created.
10. Navigate your way to Packet Rules. Use the Packet Rules interface to write at least three rules to enable static NAT. You need to create two New Defined Address rules and a New Mapped Address rule.
 - a. In the Packet Rules window, create a new rules file by selecting New File from the File menu.
 - b. In the **New Rules file** menu, right-click on **Defined Addresses** and select **New Defined Address**.
 - c. Enter an Address name, the IP address of the logical partition, and a type of Trusted. For the example, you would enter:
 - Address Name = LINUXPART
 - Defined Address = IP Address where your IP address equals 192.168.1.5
 - Type = Trusted
 - d. In the **New Rules file** menu, right-click on **Defined Addresses** and select **New Defined Address**.


- e. Enter an Address name, the IP address of the logical partition, and a type of Border. For the example, you would enter:
 - Address Name = SHELL
 - Defined Address = IP Address where your IP address equals 10.1.1.3
 - Type = Border
- f. Expand the **Address Translation** menu item from the New Rules file menu window.
- g. Right-click on **Mapped Addresses** and select **New Mapped Address**.
- h. Enter the Mapped address name, the To address name, and the line name. You can leave the Journaling set to off. For the example, you would enter:
 - Mapped address name = LINUXPART
 - To Address name= SHELL
 - Line = ETHLINE
 - Journaling = OFF
- i. Verify your rules by selecting **Verify** under the File menu.
- j. Save your rules file.
- k. Upon successful verification, select **Activate** from the file menu. Your static NAT rules are now active.

To test outbound communications, ping an external host from your logical partition. Then from that external host, ping your logical partition to test inbound communications.

Configuring TCP/IP Routing:

You can route traffic to your logical partitions through your System i product server with various routing techniques.

This solution is not difficult to configure on your System i product but depending on the topology of your network, it may not be practical to implement. Consider the following pre-V5R3 figure:



The existing TCP/IP interface (10.1.1.2) connects to the LAN. The LAN is connected to remote networks with a router. The Linux TCP/IP interface is addressed as 10.1.10.2 and the virtual Ethernet TCP/IP interface as 10.1.10.1. In i5/OS, if you turn on IP Datagram Forwarding, System i product will route the IP packets to and from the logical partition. When you define your Linux TCP/IP connection, the router address must be 10.1.10.1.

The difficulty of this type of routing is getting the IP packets to the System i product. In this scenario, you could define a route on the router so that it passes packets destined to the 10.1.10.0 network to the 10.1.1.2 interface. That works great for remote network clients. It would also work for the local LAN clients (clients connected to the same LAN as the System i product) if they recognize that same router as their next hop. If they do not, then each client must have a route that directs 10.1.10.0 traffic to the System i 10.1.1.2 interface; therein starts the impracticability of this method. If you have hundreds of LAN clients, then you have to define hundreds of routes.

To configure a setup similar to this, use the following instructions:

1. During the configuration of your logical partition, make sure you selected to create virtual Ethernet. See “Configuring a logical partition” on page 11 for more information.
2. “Create an Ethernet line description for virtual Ethernet” on page 24. Note the port of the hardware resource, because you will need it when you set up TCP/IP in your logical partition. In the case of the above example, assume the hardware resource is CMN05 with a port of 0 and although arbitrary, the line description name is VETH0. `CRTLINETH LIND(VETH0) RSRcname(CMN05) LINESPEED(1G) DUPLEX(*FULL)`
3. Add a TCP/IP Interface for your virtual line description. Following the example, you could use:
`ADDTCPIFC INTNETADR('10.1.10.1') LIND('VETH0') SUBNETMASK ('255.255.255.0')`
Start your new TCP/IP interface using the command `STRTCPIFC INTNETADR (yourIPaddress)` or `STRTCPIFC INTNETADR('10.1.10.1')` for the example.
4. On your logical partition, set up networking using the instructions or tools provided by your Linux distribution. Be sure to use the correct IP address, subnet mask, port, and gateway IP address.
During the TCP/IP setup in Linux, you will need to know your net or network device. If you are running a 32-bit kernel version of 2.4.9 or less, the network device is VETH plus the port number from when you configured an Ethernet line description. If you are running a 64-bit kernel or a 32-bit kernel greater than or equal to 2.4.10, the network device is ETH plus the port number. See “Create an Ethernet line description for virtual Ethernet” on page 24 for more information. In the example:
the interface IP address is 10.1.10.2
the subnet mask is 255.255.255.0
the gateway IP address is 10.1.10.1
the network device is VETH0 or ETH0 depending on the kernel version.
5. Turn on IP datagram forwarding. This allows the i5/OS TCP/IP interfaces to pass IP packets between each other. `CHGTCPA IPDTGFW(*YES)`

Configuring a print server

Do these steps to configure a print server for Linux running in a logical partition.

If you are running i5/OS V5R4 on the primary and all of the other partitions on the server are logical partitions, it is strongly recommended that you print the server configuration for all of your I/O resources.

The primary partition auxiliary storage must be protected using RAID or mirroring to minimize losing the logical partition configuration. Logical partition configuration information is not saved during the save process, and therefore a printout will be required to allocate appropriate resources should you have to recover your server in a disaster recovery scenario.

You will also need to print the server configuration report for all logical partitions should you need to perform an MES (Miscellaneous Equipment Specification) also known as hardware upgrade to your server with logical partitions. This information will assist your IBM Business Partner or IBM Marketing Representative to understand how your server I/O resources are assigned to the logical partitions.

Note: Printing a rack configuration listing through Hardware Service Manager within SST will only provides you with a configuration listing of the resources that are allocated to that specific partition. This report will not provide you with details for the entire server. For this reason, you should use the steps outlined below using the primary partition.

Complete the follow steps to print the server configuration.

1. From the primary partition start SST or DST. **NOTE:** These steps only are valid when the system is IPLed to a running i5/OS and just IPLed to DST. i5/OS Service Functions manual SY44-55902-02 has information on attaching SCS printer for DST printouts.
2. From SST select option 5 (Work with servers); from DST select option 11 (Work with servers) and press **Enter**.
3. Select option 1 (Display partition information).
4. Select option 5 (Display server I/O resources).
5. At the Level of detail to display field, type *ALL to set the level of detail to ALL.
6. Press **F6** to print the server I/O configuration.
7. Select option 1 and press **Enter** to print to a spooled file.
8. Press **F12** to return to the Display Partition Information display.
9. Select option 2 (Display partition processing configuration).
10. Press **F6** to print the processing configuration.
11. Press **F12** to return to Display Partition Information display.
12. Select option 7 (Display communications options).
13. Press **F6** to print communication configuration.
14. Select option 1 and press **Enter** to print to a spooled file.
15. Return to an i5/OS command line and print these three spooled files.

Moving shared processing power between logical partitions

You can dynamically move shared processor units between Linux partitions.

The ability to dynamically move processor power becomes important when you need to adjust to changing workloads. Processors have minimum and maximum values associated with them. These values enable you to establish a range within which you can dynamically move the resource without varying off the logical partition. For shared processors, you can specify a minimum value equal to the minimum amount of processing power needed to support the logical partition. The maximum value must be less than the amount of processing power available on the server. If you change either the minimum or maximum value, it will require you to vary off the partition.

To move shared processing power from one logical partition to another, follow these steps:

1. In System i Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the server.
3. Expand **Configuration and Service** and select **Logical Partitions**.
4. Right-click **Logical Partition** and select **Configure Partitions**. You are now working in the Configure Logical Partitions window.
5. From the **Configure Logical Partition** window, select the partition from which you would like to move shared processing power.
6. Right-click **Shared pool processors** and select **Move**.

7. Select the units with which you want to work in **Display units in -- Processor pool**. If you need more information on this field, click **Help**.
8. Specify the amount of shared processing power to move in the **Amount to move** field. This value must be less than the current amount of shared processing power available to the logical partition. The Amount after move values for the source must fall within the minimum and maximum ranges specified for shared processing power for both partitions.
9. Select the partition to receive the dedicated processors in **Move to -- Logical partition**.
10. Click **OK** to move the shared processing power you specified.

Dynamically adding disks to a Linux partition

Do these steps to add a virtual disk to a Linux server that is running in an i5/OS logical partition.

Virtual devices simplify hardware configuration on the System i product because they do not require you to add additional physical devices to the server in order to run Linux.

You can allocate up to 64 virtual disks to a Linux partition, with a maximum of twenty disks visible to Linux, depending on which distribution you use. Each disk supports up to 1000 GB of storage. Each virtual disk appears to Linux as a single physical unit. However, the associated space in the i5/OS integrated file system is spread across all of the disk arms available on the System i product. This provides the benefits of RAID without the overhead of RAID support in the Linux kernel.

i5/OS provides the ability to dynamically add virtual disks to a Linux partition. You can allocate disk space in the integrated files system and make it available to Linux without restarting i5/OS or Linux. The Linux server administrator can also configure the newly allocated disk space and make it available, again without interrupting the server.

To dynamically add virtual disks to a Linux partition, do the following:

1. Create a disk drive using System i Navigator.
 - a. Click **Integrated Server Administration**.
 - b. Right-click the **All Virtual Disks** folder and click **New Disk**.
 - c. In the Disk drive name field, specify a disk drive name.
 - d. In the Description field, specify a description of this disk.
 - e. If you want to copy data from another disk, select **Initialize disk with data from another disk**. Then specify the source disk to copy data from.
 - f. In the Capacity field, specify the disk capacity.
 - g. In the Planned file system field, select **Open source**.
 - h. Click **OK**.
2. Link a disk drive using System i Navigator.
 - a. Click **Integrated Server Administration** → **All Virtual Disks**.
 - b. Right-click an available disk drive and select **Add Link**.
 - c. Select the server that you want to link the disk to.
 - d. Select one of the available link types and the link sequence position.
 - e. Select one of the available data access types.
 - f. Click **OK**.
3. Determine the device name of the new disk. As far as i5/OS is concerned, the space is available to the Linux partition. Now you need to run through a couple of steps within Linux to partition, format, and access the disk. The name depends on two factors:
 - The distribution of Linux
 - The number of disks currently allocated

For example, with the SUSE distribution with a single disk allocated, the first disk device is known as `/dev/hda`. The second disk (the newly allocated virtual disk) is known as `/dev/hdb`.

4. Create partitions on the new disk with the `fdisk` command in Linux. The `fdisk` command is a standard Linux command and is the same for all distributions. You must have superuser (also known as `root`) privileges to execute this command.
 - a. From the command line, type `fdisk /dev/hdb` and press **Enter**.

The command responds with the following prompt: Command (m for help):

5. Type **p** (for print) at the prompt to display the current partition table for the disk. By default, a new virtual disk shows a single partition formatted as FAT16. For example,

```
Disk /dev/hdb: 64 heads, 32 sectors, 200 cylinders
Units = cylinders of 2048 * 512 bytes
```

```
Device Boot      Start         End      Blocks   Id  System
/dev/hdb1        1           199     203760    6   FAT16
```

6. Delete the partition. The FAT16 partition is not required. You need to delete the current partition and then create a new one.
 - a. To delete the partition, Type **d** at the command prompt.

The `fdisk` command responds with the following prompt: Partition number (1-4):

7. Enter the partition number (in this case, 1), and press **Enter**. The `fdisk` prompt returns a message indicating that the deletion was successful.
8. Create a new partition.

- a. Type the command **n** for a new partition. The `fdisk` prompt returns the following:

```
Command action
E   extended
P   primary partition (1-4)
```

- b. Type the command **p** and press **Enter**. The `fdisk` prompt returns the following:

```
Partition number (1-4):
```

- c. Since this is the first partition on the disk, enter **1** and press **Enter**. The `fdisk` prompt returns the following:

```
First cylinder (1-200, default 1):
```

- d. Type **1** and press **Enter**. The `fdisk` prompt returns the following:

```
Last cylinder or +size or +sizeM or +sizeK (1-200, default 200):
```

- e. Type **200** and press **Enter**. The `fdisk` prompt appears again, indicating that the partition creation was successful.

Note: The type of the partition defaults to Linux. If you need a different disk type (like LVM, or Linux Extended), use the **t** (for type) command to change the type of the partition.

9. Type **w** (for write) to commit the changes and press **Enter**. Up until this point, you have made no changes to the disk structure. After you commit the changes, `fdisk` responds with the following diagnostic messages:

```
The partition table has been altered!
```

```
Calling ioctl() to re-read partition table.
Syncing disks.
```

10. Type `mkfs` to format the new partition. The `mkfs` command is a standard Linux command that should be available on each Linux distribution. There are a number of optional parameters for `mkfs`, but typically the defaults satisfy most disk uses. To format the partition created in the previous steps, enter the following command (again, as with the `fdisk` command, you need to be logged in as `root` to execute the `mkfs` command):

```
mkfs /dev/hdb1
```

Remember that a single partition exists on the second disk. Hence the name `/dev/hdb1` (the `hdb` indicates that it is the second disk, and the `1` indicates that it is partition 1). The following diagnostic messages appear:

```
mke2fs 1.28 (31-Aug-2002)
Filesystem label=
OS type: Linux
Block size=1024 (log=0)
Fragment size=1024 (log=0)
51200 inodes, 204784 blocks
10239 blocks (5.00%) reserved for the super user
First data block=1
25 block groups
8192 blocks per group, 8192 fragments per group
2048 inodes per group
Superblock backups stored on blocks:
 8193, 24577, 40961, 57345, 73729
```

```
Writing inode tables: done
Writing superblocks and filesystem accounting information: done
```

```
This filesystem will be automatically checked every 29 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
```

11. Create a new directory. Finally, within the Linux file tree, you need to create a directory that you can use to access the new file system. You also need to allocate the new filesystem to that point within the directory tree. To do this, we create a directory and mount the physical resource as follows:
 - a. Type `mkdir /mnt/data` and press **Enter**.
 - b. Type `mount /dev/hdb1 /mnt/data` and press **Enter**.
12. Add an entry. You now have new virtual disk space available to the running Linux partition. To facilitate automatic mounting of the resource across Linux starts, add an entry to the `/etc/fstab` file (this is the file that stores mount definitions). For this example, the entry in `/etc/fstab` would appear as follows:

```
/dev/hdb1 /mnt/data ext2 defaults 1 1
```

Network server description

Use the following topics to learn how to link and delete network server descriptions.

A *network server description* (NWSGD) is used to give a name to the configuration, provide an interface for starting and stopping a Linux partition, and provide a link between Linux and its virtual disks.

Sharing NWSSTG between Linux partitions on System i

Do these steps to link a network server storage space (NWSSTG) to the network server description (NWSGD) for a Linux server that is running in a logical partition.

Guest partitions or network server descriptions can share data among other network server descriptions. By sharing the data among network server descriptions, you are allowing multiple network server descriptions the ability to have read-only access to data. The option to share data among multiple network server descriptions is beneficial when you have applications that multiple partitions need to access. There is no limit on the number of network server descriptions that can read the shared data.

You can also have shared-update network server storage spaces. Multiple network server descriptions would have access to the data however; only one network server description can update the data at a time. You will not be able to update the data if other network server descriptions are reading data from the drive. This option is beneficial when shared-data changes affect applications shared by multiple partitions.

To link NWSSTG to NWSGDs, follow these steps:

1. At the i5/OS command line, enter the command `ADDNWSSTGL` and press **Enter**.

2. From the Add Server Storage Link display, provide the following information:
 - NWSSTG (Name)
 - NWSD (Name)
 - DYNAMIC (*YES)
 - DRVSEQNBR (*CALC)
3. Press **F10** (Additional Parameters).
4. Enter the type of access the storage space will have and press **Enter**.

Deleting NWSDs for a logical partition

Do these steps to unlink and delete network server descriptions (NWSD) from Linux servers running in System i logical partitions.

Before you delete a NWSD, you are required to unlink storage spaces that are associated with that NWSD. Then you can delete the NWSD. See “Unlinking disk drives from a logical partition” on page 54 for more information on how to unlink NWSDs.

To unlink the storage space for the server drive for NWSDs complete the following steps:

1. On the i5/OS command line, type `RMVNWSTGL NWSSTG(nwsdname1) NWSD(nwsdname)`.
2. Press **Enter**.

Unlinking storage spaces:

To unlink the storage space for the install source drive complete the following steps:

1. On the i5/OS command line, type `RMVNWSTGL NWSSTG(nwsdname2) NWSD(nwsdname)` and press **Enter**.
2. Any user defined storage spaces that have been linked to the NWSD can also be removed at this time using the following command as often as needed
 - a. On the i5/OS command line, type `RMVNWSTGL NWSSTG(nwsstgname) NWSD(nwsdname)`.
 - b. Press **Enter**.

Deleting the NWSD:

To delete the network server description (NWSD) for a logical partition complete the following steps:

1. On the i5/OS command line, type `WRKNWSD` and press **Enter**.
2. Type **8** in the Opt field to the left of the Network Server and press **Enter**. The Work with Configuration Status display appears.
3. If the status of the NWSD is not varied off, type **2** in the Opt field to the left of the Network Server and press **Enter**. Otherwise, continue to the next step.
4. Press **F3** to return to the previous screen.
5. Type **4** in the Opt field to the left of the Network Server and press **Enter**.
6. Press **Enter** in the Confirm Delete of Network Server Descriptions display.

Deleting virtual disks for Linux running in logical partitions

Do these steps to delete virtual disks with System i Navigator.

Before you can delete a disk drive, you must unlink it from the NWSD. Once you have unlinked it, you can delete it. See “Unlinking disk drives from a logical partition” on page 54 for more information.

Deleting disk drives for a logical partition with System i Navigator:

To delete the disk drive, do the following:

1. Click **Integrated Server Administration** → **All Virtual Disks**.
2. Right-click the disk drive that you want to delete.
3. Click **Delete**.
4. Click **Delete** on the confirmation panel.

Deleting disk drives for a logical partition with CL commands:

Follow these steps to delete the disk drive using the Delete Network Server Storage Space CL command, DLTNWSSTG:

1. Type DLTNWSSTG and press **F4**.
2. Type the name of the disk drive in the Network server storage space field. Press **Enter**.

Network server description (NWSD) parameters and descriptions

A *network server description* (NWSD) is used to give a name to the configuration, provide an interface for starting and stopping a Linux partition, and provide a link between Linux and its virtual disks.

Before you create an NWSD for a logical partition, you should understand each parameter description. See “Creating a network server description and a network server storage space” on page 16 for more information.

The following table will describe each parameter value and how each parameter is used when running a logical partition.

Table 11. NWSD parameters and descriptions

Prompt	Parameter	Parameter description
Network server description	NWSD	Specifies the name that you give the network server description. It is recommended to use the same name as your partition. You can create multiple NWSDs that point to the same partition, however only one NWSD can be active for a partition at a time.
Resource name	RESCUE	Specifies the resource name that identifies the hardware that the description represents. This option does not apply to a *GUEST partition.
Network server type	TYPE	Specifies which type of network description to create.
Online at IPL	ONLINE	Specifies whether this object is automatically varied on at initial program load (IPL).
Vary on wait	VRYWAIT	Specifies whether the network server description is varied on asynchronously or synchronously. For a synchronous vary on, specifies how long the server waits for the vary on to complete.
Partition	PARTITION	Specifies the name of the partition given during the configuration of a logical partition. If you provide the incorrect name of a partition here, it will not be detected until you vary on the NWSD.
Code page	CODEPAGE	Specifies the ASCII code page represents the character set Linux uses and that i5/OS assumes the Linux console is using.

Table 11. NWSD parameters and descriptions (continued)

Prompt	Parameter	Parameter description
TCP/IP port configuration	TCPPORTCFG	Specifies the setting the TCP/IP information does not cause any configuration to occur on i5/OS. All information entered here is made available to Linux in the /proc/iSeries/config file. This information can be used by Linux to automatically configure TCP/IP information. There is a feature that prevents the TCP/IP route configuration from showing up on the CRTNWSD screen. Select option F9 to make the configuration options appear.
TCP/IP route	TCPRTE	Allows the user to identify routes to remote destination servers or networks to TCP/IP configuration for the network server. A route specification has three elements, the route destination, the subnet mask, and the next hop internet address. Up to 24 route specifications can be entered.
TCP/IP host name	TCPHOSTNAM	Specifies the short form of the host name to be associated with the network server. The host name can be a text string having 2 through 63 characters. The following characters are allowed in host names: <ul style="list-style-type: none"> • Alphabetical characters A through Z • Digits 0 through 9 • Minus sign (-)
TCP/IP domain server	TCPDMNNAME	Specifies the local domain name associated with the network server. A domain name can be a text string having 2 to 255 characters. Domain names consist of one or more labels separated by periods. Each label can contain up to 63 characters. The following characters are allowed in domain names: <ul style="list-style-type: none"> • Alphabetical characters A through Z • Digits 0 through 9 • Minus sign (-) • Period (.). Periods are only allowed when they separate labels of domain style name (refer to RFC 1034).

Table 11. NWSD parameters and descriptions (continued)

Prompt	Parameter	Parameter description
TCP/IP name server	TCPNAMESRV	Specifies the internet address of the name server that is used by the network server. Typically, this is the same value as it is for the System i product.
Restricted device list	RSTDDEVRSC	Specifies a logical partition running Linux has access to all optical and tape devices in the server. To restrict Linux from using any of these devices, use this feature.
Synchronize date and time	SYNCTIME	Specifies whether the System i product should synchronize the network server date and time with the System i product date and time. As with the TCP/IP configuration, this parameter is just reflected in the /proc/iSeries/config file.
IPL source	IPLSRC	Specifies where the Linux kernel will be loaded from.
IPL stream file	IPLSTMF	Specifies the path to the Linux kernel. The user running the vary on command must have read access to the file and the path leading to the file.
IPL parameters	IPLPARM	Specifies a string of characters that will be passed to the Linux kernel at start time. It consists of commands or configuration information for the kernel.
Authority	AUT	The server determines authority for the object using the value specified for the create authority prompt for the library in which the object is created.
Text	TEXT	The text that briefly describes the network server. The text description must be no more than 50 characters and enclosed in apostrophes.

Application support for Linux in a logical partition

- | Use the following topics to learn how to effectively use System i Access for Linux, Samba, and IBM i5/OS Extended Integrated Server Support in a logical partition.

System i Access for Linux support for an ODBC driver and 5250 session

System i Access for Linux includes an ODBC driver to access System i database data from Linux applications written to the ODBC API, and a 5250 emulator to access the System i product. The ODBC driver is based on the System i Access for Windows ODBC driver.

For further information, see System i Access for Linux.

i5/OS NetServer support for Samba

i5/OS NetServer provides support for Linux running in a System i logical partition. This support allows a Linux client running Samba to connect to i5/OS NetServer.

Samba is an Open Source client and file system compatible with Microsoft® Networking that comes with many current distributions of Linux. For more information on i5/OS NetServer support for Samba, refer to Linux on the System i platform  .

How to use SAMBA and NFS to access files in the Integrated File System

Users and/or applications of a Linux partition can use a variety of tools to access files stored in the integrated file system. Two of these tools are Network File System (NFS) and SAMBA.

For more information on i5/OS NetServer support for Samba, refer to Linux on the System i Platform  .

Installing Extended Integrated Server Support

Do these steps to install the IBM i5/OS Extended Integrated Server Support licensed product.

Extended Integrated Server Support is required for “Configuring file-level backup for Linux in a logical partition” on page 50 and “Backing up storage spaces for an active Linux server in a logical partition” on page 48.

This function is supported for the following distributions of Linux running in a logical partition:

- Red Hat Enterprise Linux 5 (RHEL 5)

- SuSE Linux Enterprise Server 10 (SLES 10)

1. Vary off the network server description (NWSD) for the Linux server.

2. Install these i5/OS options and licensed products:

- 5761-SS1 option 12 Host Servers

- 5761-SS1 option 29 Integrated Server Support

- 5761-SS1 option 34 Digital Certificate Manager

- 5761-SS1 option 35 CCA Cryptographic Service Provider

- 5761-LSV IBM Extended Integrated Server Support for i5/OS

3. Configure the i5/OS user QFPAD.

- a. Enable the user account.

- b. Create a password.

4. Vary on the NWSD. Make sure that you select the generate path certificates (GENPTHCERT) option.

For example, enter `VRYCFG CFGOBJ(NWSDNAME) CFGTYPE(*NWS) STATUS(*ON) GENPTHCERT(*YES)` where *NWSDNAME* is the name of the network server description (NWSD) object for the Linux server. This option will enable secure sockets layer (SSL) communication between the i5/OS operating system and the Linux server.

5. Sign on to the Linux server as the root user.

6. Install the following software on the Linux server.

- samba

- openssl

- libacl

- libattr

See your Linux documentation or contact your distributor for more information.

7. Create the directory `/etc/ibmlsv` for the credentials file. This file contains information that allows the Linux server to sign on to i5/OS as user QFPAD to access software updates.

- | Run the following command to create the directory:
- ```
| mkdir /etc/ibmlsv
```
- | 8. Edit the `/etc/ibmlsv/credentials` file.
    - | a. Add `username=qfpad`.
    - | b. Add `password=password`

| where *password* is the password for the i5/OS user QFPAD.

| You can also create this file by entering the following command at the Linux command prompt:

```
| echo -e "username=qfpad\npassword=password" >> /etc/ibmlsv/credentials
```
  - | 9. Configure the Linux server to access updates to the Extended Integrated Server Support software.
    - | a. Create the directory `/mnt/ibmlsv` to use as mount point for the Extended Integrated Server Support licensed program. For example, enter run the following command:
 

```
| mkdir /mnt/ibmlsv
```
    - | b. Mount the file share for software included in IBM Extended Integrated Server Support for i5/OS. Use type `smbfs` if your Linux operating system will support it. Otherwise use type `cifs`. The following examples show how to mount the directory from the Linux command line where *systemihostname* is the host name (or IP address) of the System i product and *nwsdname* is the name of the network server description (NWS) for the Linux server.
      - | • On SLES 10:
 

```
| mount -t smbfs -o credentials=/etc/ibmlsv/credentials,rw //systemihost/nwsdname /mnt/ibmlsv
```
      - | • On RHEL5:
 

```
| mount -t cifs -o credentials=/etc/ibmlsv/credentials,file_mode=0755,dir_mode=0755 //systemihost/nwsdname /mnt/ibmlsv
```
  - | 10. If there is a firewall on the Linux system, verify that it will permit Samba traffic.
  - | 11. Run the `ibmsetup` utility at the Linux command prompt. Enter `/mnt/ibmlsv/Install/ibmsetup.sh address` where *address* is the host name or IP address of the System i product.

## | **Updating Extended Integrated Server Support**

| Do these steps to update the Extended Integrated Server Support licensed product.

| The Extended Integrated Server Support licensed product includes components that run on both i5/OS and Linux operating systems. Use PTFs to update the i5/OS components and the `ibmlsvupdt` utility to update the Linux software.

- | 1. Apply the IBM Extended Integrated Server Support PTFs to i5/OS. The latest PTFs are listed on the System i integration with BladeCenter and System x  ([www.ibm.com/systems/i/bladecenter/ptfs.html](http://www.ibm.com/systems/i/bladecenter/ptfs.html)) Web site.
- | 2. Sign on to the Linux server as the root user.
- | 3. Run the `ibmlsvupdt` command at the Linux prompt. The syntax is
 

```
| ibmlsvupdt <userid> [<address>]
```

| where *userid* is an i5/OS user profile and *address* is the IP address or host name of the hosting i5/OS partition.

| The Linux operating system prompts for the i5/OS user password.

- | 4. Type the i5/OS user password and press Enter.

## **Backing up and recovering a logical partition**

Use these tasks to back up and recover a Linux server that is running in a System i logical partition.

System i integration of a logical partition combines two operating systems (i5/OS and the logical operating system). You may be able to use i5/OS or Linux utilities or a combination of both to manage back ups.

When you are planning your back up strategy, refer to Backup and Recovery, as well as Linux documentation. For information on backing up and recovering Linux partitions on model 5xx systems, see Partitioning for Linux with an HMC in the IBM Systems Hardware Information Center.

## Backup options for virtual and directly attached disks

Read to learn about backup options for virtual and directly attached disks for Linux and i5/OS utilities.

There are two types of options for backing up and recovering data involving Linux partitions on the System i platform:

- Utilities that run within Linux
- Commands that run within i5/OS

### Utilities that run within Linux

Utilities that run within Linux benefit those users who have Linux installed to a native disk and those users who cannot vary off their Linux partition while they are backing up data. The most common data backup utility in Linux is the tar (a contraction of tape archive) utility. The tar utility archives files and directories into one file that you can save in either of two ways:

- You can write the file directly to a virtual or directly attached tape device. For information about this alternative, see “Saving and restoring Linux data with System i tape devices” on page 45.
- You can write the file to a file in the partition’s file system. Following is an example of this alternative.

**Example:** `tar -cvf /tmp/etc.tar /etc`

Where, in the above example:

| Variable            | Description                                                   |
|---------------------|---------------------------------------------------------------|
| <i>c</i>            | create a tar file                                             |
| <i>v</i>            | verbose (show the files that are being added to the tar file) |
| <i>f</i>            | name of the tar file follows                                  |
| <i>/tmp/etc.tar</i> | name of the tar file                                          |
| <i>/etc</i>         | object to be added to the tar file                            |

**Note:** Since */etc* is a directory, the utility adds all the contents of the directory and its subdirectories to the tar file.

After you create the tar, you can save it to an offline medium in several ways, including writing it to a virtual or directly attached tape device or copying it to the IFS for inclusion in a subsequent save/restore operation.

You can tar the data on a Linux partition during normal server usage, and you can easily automate and start the tar utility by using the partition’s cron daemon (a scheduling mechanism for Linux; cron is an abbreviation of chronology). You can also use the at utility to schedule a single backup request. For example, if you want to use the tar utility to back up the */etc* directory at 10 p.m. on the 19th of September, you could enter the following Linux command:

```
at 10pm Sep 19 -f tar.command
```

**Note:** For additional information about the tar, at, and cron utilities, use the Linux man (an abbreviation for manual) command. For example, man tar

## **i5/OS commands**

Users of virtual disks also have powerful tools to use in i5/OS for backup and recovery. They can use the Save Object (SAV) and Restore Object (RST) commands to save and restore entire virtual disks in their current state.

The SAV command saves the directory that has the same name as the virtual disk under the QFPNWSSTG directory in the integrated file system. This method of backup and recovery is most effective if the Linux kernel is saved in a PrEP start partition on the virtual disk. On most Linux distributions, this normally occurs as part of a default installation.

If you save the Linux kernel in a PrEP partition, you can restore and start the Linux partition after a total i5/OS reinstallation. You can also transport and restore saved virtual disks to other System i servers via FTP and tape.

## **Using the shared, read-only storage space**

Do these steps to read data from a shared storage space for a Linux server in a System i logical partition.

Version 5 Release 2 added the ability for multiple Linux partitions to share virtual disks. Multiple partitions can now read data from a shared storage space concurrently. The ability to share virtual disks is useful whenever two or more Linux partitions need to use the same file. For example:

- Two or more Linux partitions used as Web servers could point to the same Web page.
- Two or more Linux partitions could want to read documentation from the same file.
- Two or more Linux partitions could want to run the same copy of an Apache Web server program.

To use the storage space, you must link to the virtual disk and provide users access. Use either of the following methods to use shared, read-only storage:

### **Using System i Navigator:**


To link a disk drive using System i Navigator, follow these steps:

1. Click **Integrated Server Administration** → **All Virtual Disks**.
2. Right-click an available disk drive and select **Add Link**.
3. Select the server that you want to link the disk to.
4. Select one of the available link types and the link sequence position.
5. Select the **Shared - Read** access type.
6. Click **OK**.

### **Using the Add Network Server Storage Link command:**

To link a disk drive using the Add Network Server Storage Link command, follow these instructions:

1. Type the Add Network Server Storage Link (ADDNWSSTGL) command. The storage link for virtual disks that Linux partitions are going to share concurrently must have an access type of read (\*READ).



On the Add Network Server Storage Link (ADDNWSSTGL) display, the value you enter in the Dynamic storage link field can be either \*YES or \*NO. To display the Access option, press the F9 key. If multiple Linux partitions share disks, the partitions must access (that is, mount) the disks from Linux with read-only access.

2. You can specify that Linux has read-only access to the disk two ways:
  - By using an option on the mount command
  - By specifying read-only as a parameter in the /etc/fstab file

The following is an example of the using the mount command to mount a file system as read-only:

```
mount -o ro /dev/hdb1 /mnt/data2
```

An entry in the /etc/fstab file to mount a disk resource as read-only would resemble the following:

```
/dev/hdb1 /mnt/data ext2 ro 1 1
```

#### Updating and providing users access:

To update a shared storage space and provide users access, follow these steps:

1. Give all users \*SHRUPD access to the storage space.
2. Have all users open the storage space concurrently for read-only access.
3. Have all users stop accessing the data on that storage space and close (unmount) that device from within Linux.
4. Have one user open the device for read-write access, update the data, and close the device.
5. Have all users reopen the device for read-only concurrent access.

Other restrictions:

- Version 5 Release 2 or later of i5/OS is required, but it does not have to be in the primary partition.
- The maximum size supported for a storage space is 1000 GB.

- The maximum number of storage spaces supported is 64. However, while 64 storage spaces may be linked to the server, only the first 20 storage spaces are visible to Linux, depending on the Linux distribution.

## Sharing access to disks in read-write mode

Learn how to share access to disks in read-write mode for Linux servers in a System i logical partition.

Although simultaneous access to a virtual disk by multiple logical partitions is limited to read-only access, V5R2 enhancements make it possible for multiple logical partitions to have the same virtual disk linked to their network storage descriptor (NWSD). However, only one partition can access the disk in read-write mode at any given time. See “Using the shared, read-only storage space” on page 42 for more information.

Use either of the following methods to access disks in read-write mode:

### Using System i Navigator:

To access a disk drive using System i Navigator, follow these steps:

1. Click **Integrated Server Administration** → **All Virtual Disks**.
2. Right-click an available disk drive and select **Add Link**.
3. Select the server that you want to link the disk to.
4. Select one of the available link types and the link sequence position.
5. Select the **Shared - Update** access type.
6. Click **OK**.

*Using the Add Network Server Storage Link command:*

To link a disk drive using the Add Network Server Storage Link command, follow these instructions:

1. Type the Add Network Server Storage Link (ADDNWSSTGL) command. This command links the network storage space to the network server descriptor. You must set the access type to shared update (\*SHRUPD).

```

Session A - [24 x 80]
Add Server Storage Link (ADDNWSSTGL)

Type choices, press Enter.

Network server storage space . . . > DATA2 Name
Network server description . . . > LSAM1 Name
Dynamic storage link *YES *NO, *YES
Network server type *NWSD Character value
Drive sequence number *CALC 1-64, *CALC, *QR

Additional Parameters

Access *SHRUPD *UPDATE, *READ, *SHRUPD

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

MIC a 13/043

```

2. Specify either \*YES or \*NO in the Dynamic storage link field.

3. Press the **F9** key to display the Access field.

After you link the network storage to the partition, Linux can access the storage by using the mount command. For example, to mount a virtual disk in read-write mode, you could use the following mount command:

```
mount /dev/hdb1 /mnt/data2
```

If no other partition has the disk mounted, the mount is successful. Otherwise, Linux generates the following error:

```
mount: /dev/hdb5: can't read superblock
```

In order to grant access to the partition, the first partition (that is, the partition that is currently accessing the resource) must first relinquish access. To accomplish this result, use the Linux `umount` command as follows:

```
umount /dev/hdb1
```

After the first Linux partition relinquishes the resource, the second Linux partition can mount and access it.

## Cloning Linux servers in System i partitions

The use of virtual devices (disk and network) reduces the number of hardware devices that you need to attach directly to a Linux partition in order to clone a Linux partition on one System i product and restore it to the same or another System i product without reconfiguring the hardware in Linux.

To copy a virtual disk, follow these steps:

1. Sign on to **Dedicated Service Tools**.
2. Run the Create a Network Server Description (CRTNWS) command.
3. Run the Work with Network Server Storage Space (WRKNWSSTG) command.
4. Enter 3=Copy in the Option field next to the network server storage space that you want to copy. The Create NWS Storage Space display appears. The input fields for the Size parameter and the From storage space parameter already contain values.
5. Specify a name and description for the new virtual disk and press the **Enter** key. At this point, the server creates the new virtual disk as a copy of the existing virtual disk.
6. Repeat steps 1-5 for each additional virtual disk you want to clone.
7. Link each network server storage space to the appropriate network storage description by using the Add Network Server Storage Link (ADDNWSSTGL) command.
8. Vary on the new logical partition by using the VRYCFG command.

To save additional disk storage, split a virtual disk into two disks: a larger one for read-only access and a smaller one for read/write access. You can now copy the read/write disk and share the read-only disk.

## Saving and restoring Linux data with System i tape devices

To save hosted Linux data to a shared tape drive and restore the data from the tape drive, you can use either the Linux `tar` command or the i5/OS Save Object (SAV) and Restore Object (RST) commands.

If you use `tar` to save data, the only way you can restore that data is by using `tar` again. Similarly, if you use SAV to save data, the only way you can restore that data is by using RST. The two methods of backing up and restoring data are not compatible.

The following restrictions apply:

- The names of the tape drives in RedHat and Suse are stored in `/dev/iseres/vt0`, `vt1`, etc. TurboLinux has its own naming convention.

- To use the tape device from Linux, you must vary the tape off under i5/OS. To use the i5/OS SAV or RST command to save or restore the NWSD, Linux must be inactive (that is, the NWSD must be varied off).
- Saving the storage space is typically much faster than saving using tar, but it does not provide file-level backup and recovery.
- Linux does not support switching tapes in a library device. You can only use the tape that is currently in the device.
- You cannot save i5/OS data and tar data on the same tape volume.

## Saving and restoring hosted Linux files by using Linux tar

Backups using Linux tar are file-level backups. They save only the files and directories that the tar command specifies. Therefore, you cannot use tar to save Linux data that is not in the fileserver. For example, you cannot save a kernel residing in the PReP start partition by using the tar command.

### Saving files by using Linux tar

The Linux command to save hosted Linux files to the host's shared tape drive is the following:

```
tar -b 40 -c -f /dev/iseries/vt0 files
```

where

|                     |                                                |
|---------------------|------------------------------------------------|
| tar                 | = command name (contraction of "tape archive") |
| -b 40               | = block size in sectors                        |
| -c                  | = command action (create)                      |
| -f /dev/iseries/vt0 | = virtual tape device and number               |
| files               | = names of files to be saved                   |

Linux normally treats tape as a "character device" that it can quickly read from or write to in long streams of data, but cannot quickly access to find specific data. By contrast, Linux treats a disk or CD as a "block device" that it can read from or write to quickly at any point on the device, making it suitable for the mount command. The -b 40 argument specifies that Linux should write the archive stream in blocks of 40 sectors (20 kilobytes). If you do not specify a value for this argument, the default value is 20 sectors (10 kilobytes), which does not perform as well over virtual tape as a value of 40 does.

The -c argument specifies that the tar command creates a new archive or overwrites an old one (as opposed to restoring files from an archive or adding individual files to an existing archive).

The -f /dev/iseries/vt0 argument specifies that the command uses virtual tape 0 on the iSeries server. After the tar command runs, Linux closes the tape device and rewinds the tape. If you want to save more than one archive on the tape, you must keep the tape from rewinding after each use, and you must position the tape to the next file marker. To do this, specify the nvt0 (nonrewinding virtual tape) device instead of vt0.

The files argument specifies the names of the files and directories that you want to save.

### Restoring files by using Linux tar

The Linux command to restore hosted Linux files from the host's shared tape drive is the following:

```
tar -b 40 -x -f /dev/iseries/vt0 files
```

where the -x (extract) argument replaces the -c (create) argument in the tar command used to save files to tape.

## Saving and restoring hosted Linux data by using i5/OS SAV and RST

Backups of hosted Linux data from i5/OS are drive-level. This means that Linux backs up the entire contents of a virtual disk or network storage space, rather than individual Linux files. Thus, the correct SAV command backs up any information on the drive, including a kernel in the PReP start partition.

### Saving Linux hosted data by using i5/OS SAV

On i5/OS your data is in a network server storage space. The i5/OS commands to save hosted Linux files to the host's shared tape drive include the following:

1. Save Object (SAV), where you save the network server storage space to the save file. On the Save Object display, enter the following parameter values:
  - In the Device field, enter the associated i5/OS device description. For example, if your tape device is named TAP01, enter `/qsys.lib/tap01.devd`. To save to a save file in a library like QGPL, enter `/qsys.lib/qgpl.lib/myfile.file`.
  - In the Objects: Name field, enter the IFS location of the network server storage space. For example, if your network server storage space is named TEST1, enter `/qfpnwsstg/test1`.
2. Display Saved Objects - Save File (DSPSAVF) to verify that the changed save file exists. In the Option field to left of the new save file name, enter `5=Display objects in subdirectory`. You see a list of the stream files in the save file.

### Restoring Linux hosted data by using i5/OS RST

The i5/OS command to restore hosted Linux files from the host's shared tape drive is Restore Object (RST). On the Restore Object display, enter the following parameter values:

- To restore from a real tape device, enter the associated i5/OS device description in the Device field. For example, if your tape device is named TAP01, enter `/qsys.lib/tap01.devd`. To restore from a save file in a library like QGPL, enter `/qsys.lib/qgpl.lib/myfile.file`.
- In the Objects: Name field, enter the IFS location of the network server storage space, for example, `/qfpnwsstg/test1`.

### Troubleshooting Linux virtual tape

If errors occur while you access Linux virtual tape, examine the file `/proc/iSeries/viotape`. It describes the mapping between i5/OS device names and Linux device names and records the last error for each tape device. Common errors and recovery scenarios include the following:

#### Device unavailable

Make sure the device is varied off under i5/OS.

#### Not ready

Retry the operation. If the operation still fails with the same description in `/proc/iSeries/viotape`, verify that the correct medium is in the tape drive.

#### Load failure or Cleaning cartridge found

Verify that the correct medium is in the tape drive.

#### Data check or Equipment check

Verify that you are using a supported block size to read or write the tape. All known IBM-supported tape devices can use a block size of 20 kilobytes (supplied by the `-b 40` argument to `tar`).

#### Internal error

Contact your service representative.

### Saving logical partition configuration data

Logical partition configuration data is automatically maintained for the life of the physical server.

Each i5/OS partition load source contains a copy of the configuration data. A logical partition start devices do not contain partition configuration data. Systems comprised solely of a single primary i5/OS partition and one or more logical partitions should have mirrored i5/OS load source devices in the primary partition. For more information on how to save logical partition data for one or more logical partitions, refer to Save all configuration data.

Only disaster recovery to a different physical server would require that you rebuild the configuration from the beginning. You should print your server configuration when you make changes to your logical partition configuration. This printout will help you as you rebuild the configuration. See “Configuring a print server” on page 30 for more information.

During a save operation, the configuration data for the logical partition is not saved to the media volume. This allows data to be restored to a server whether or not it has logical partitions. However, you can work with the configuration data for logical partitions as needed for recovery purposes.

**Note:** Logical partitions that you keep turned off for extended periods should be restarted at least once after any change to the logical partition configuration. This allows the server to update the changes on that logical partition’s load source.

### **Backing up the NWSD and disk drives associated with a logical partition**

Do these steps to back up the network server description (NWSD) and disks associated with a Linux server in a System i logical partition.

When you install the logical partitions with virtual disk on System i, the hosting i5/OS partition creates a network server description and disk drives for your logical operating system that you need to back up. Some of the disk drives are server-related (the installation and server drives); others are user-related. Because your logical operating system may consider them a unified server, you need to save all the disk drives and the network server description to restore properly.

The implementation of a logical partition for System i allows you to save and restore virtual disk as i5/OS network server storage space objects. These objects are saved as part of the i5/OS server when you perform a full i5/OS server backup. You can also specifically save the network server description and storage spaces that are associated with a logical partition on a System i product. Daily backup of the server drive is a good idea.

To back up the network server description and the disk drives that are associated with a logical partition on System i, see “Backing up NWSDs for a logical partition” on page 54.

### **Backing up storage spaces for an active Linux server in a logical partition**

- | Use the Save (SAV) command to back up storage spaces for an active Linux server in a logical partition.
- | Use of this feature requires the IBM Extended Integrated Server Support for i5/OS licensed program (5761-LSV).
- | The disks that you create for your Linux servers are in the integrated file system. To save these storage spaces from i5/OS, you use the Save (SAV) command.
- | The i5/OS operating system saves the changes that are made to the storage space during a save operation. This information is stored in a temporary file that can be up to 25% of the total size of the storage space. This default setting should work for most configurations. For information about customizing the backup process, see “Customizing storage space backup for an active Linux server” on page 49.
- | To save disks from i5/OS, do these steps.
  - | 1. Install the IBM Extended Integrated Server Support licensed program. See “Installing Extended Integrated Server Support” on page 39.

2. If you are saving to tape, mount a tape that is formatted for i5/OS.
3. For network server storage spaces created in an independent disk pool, verify that the auxiliary storage pool (ASP) device is varied on before saving the storage space.
4. On the i5/OS command line, type `SAV` and press F4.
5. Specify a value for the **Device** field.
  - If you are saving the storage space to tape, specify the name of the tape device. For example, type `/QSYS.LIB/TAP01.DEVD` in the *Device* field.
  - If you are saving the storage space to a save file instead of to tape, specify the path to the save file as the device or specify the name of your device. For example, to use a save file named `MYSAVF` in library `WINBACKUP`, you would specify: `'/QSYS.LIB/WINBACKUP.LIB/MYSAVF.FILE'` for the device.
  - Otherwise, use the name of your device (for example, `/QSYS.LIB/TAP01.DEVD`).
6. In the **Name** field under **Objects**, specify `'/QFPNWSSTG/stgspc'` where *stgspc* is the name of the network server storage space.
7. Specify `*YES` for the **Save active** parameter. This option allows the storage space to be saved while it is still being used the system.
8. Specify `*NWSSTG` for the **Save active option** parameter. This option allows network server storage spaces in directory `'/QFPNWSSTG'` to be saved when they are active.
9. Specify values for any other parameters that you want and press Enter to save the storage space.

#### Customizing storage space backup for an active Linux server:

Use the freeze and thaw scripts to configure storage space backup for an active Linux server in a logical partition.

The default settings should work for most environments. Use the freeze and thaw scripts if you receive a message that too much space is being used by the process that tracks changes. You can also use the scripts if you know that applications on your Linux server will make frequent read and write requests to the storage space during the backup.

- The `ibmlsvfreeze.sh` script runs when i5/OS starts to back up a storage space. Use this script to stop applications that might fill the temporary storage space.
- The `ibmlsvthaw.sh` script runs when i5/OS finishes backing up a storage space. Use this script to start any applications that you stopped with the `ibmlsvfreeze.sh` script.

Do the following steps to customize storage space backup for a Linux server.

1. Copy the freeze and thaw scripts to the `/etc/ibmlsv` directory and rename them. You can use the following commands at the Linux console.
  - a. `cp /mnt/ibmlsv/service/ibmlsvfr.sh /etc/ibmlsv/ibmlsvfreeze.sh`
  - b. `cp /mnt/ibmlsv/service/ibmlsvth.sh /etc/ibmlsv/ibmlsvthaw.sh`
2. Edit the scripts. See your Linux documentation for more information about editing shell scripts.
3. Use the save (`SAV`) and restore (`RST`) commands to save the storage space.

#### Backing up and recovering individual files and directories for Linux servers running in logical partitions

Use these tasks to back up individual Linux server files and directories for servers that use virtual disks that are hosted by the i5/OS operating system.

The IBM Extended Integrated Server Support licensed product provides support for file-level backup on Linux servers. You can use the i5/OS save (`SAV`) and restore (`RST`) commands to save files to System i tape, disk, or optical devices.

| File-level backup for Linux has the following restrictions:

- | • The maximum size of file that can be saved or restored is 4GB.
- | • Hard-linked files will be restored as separate copies, not as linked files.
- | • The files saved must reside in the ext2, ext3, ext4, JFS/JFS2, ReiserFS, or XFS file system.
- | • Files in the /dev, /sys, /proc and /swap file systems cannot be backed up or restored.
- | • Files must be saved from and restored to the same operating system. For example, a file saved from a Linux server cannot be restored to a Windows server.

### | **Configuring file-level backup for Linux in a logical partition:**

| Do these steps to configure file-level backup for Linux in a logical partition.

| Use of this feature requires the IBM Extended Integrated Server Support for i5/OS licensed program (5761-LSV).

- | 1. Install the IBM Extended Integrated Server Support for i5/OS licensed program. See “Installing Extended Integrated Server Support” on page 39.
- | 2. Use the ping utility to verify that i5/OS can contact the Linux server host name or IP address. The host name of the Linux server is typically the same as the NWSD name. See Ping for information about using the ping utility.
  - | a. If the Linux host name is the same as the NWSD name and the Linux host name can be resolved using DNS, then take no action.
  - | b. If there is no DNS entry for the NWSD name, then add the NWSD name to DNS or use the Add TCP/IP Host Table Entry (ADDTCPHTE) command to add the IP address of the Linux server to the System i host table. See Add TCP/IP Host Table Entry (ADDTCPHTE).

- | 3. Create an account on the Linux server for file-level backup.

| The i5/OS operating system will use this user to sign into the server. This user must have access to the files that you want to save. The user ID and password must match the i5/OS user and password that will be used for backing up files. If the user ID and password do not match, the system will not find the files on the Linux server.

- | 4. Create a Samba password for the Linux user. See the Samba documentation for more information about how to create a Samba password.

| For example, enter `smbpasswd -a userid` where *userid* is the name of the Linux user.

- | 5. Create Samba shares for the data that you want to save. See the Samba documentation for more information about creating shares.
- | 6. Ensure the Samba services can be accessed through your server’s firewall. See your Linux distribution documentation for information about how to configure your firewall to allow Samba traffic.
- | 7. Start Samba. For example, enter `smbd -D` at the Linux command line.
- | 8. Use the Work with Object Links (WRKLNK) CL command to verify that the i5/OS operating system can access the Samba shares on the Linux server through the QNTC file system.

### | **Adding members to the QAZLCSAVL file:**

| Do these steps to add members to the QAZLCSAVL file to back up files for Linux running in logical partitions

| Create a member for each share that you want to back up. *nwsdname* is the name of the network server description (NWSD) for the server.

- | 1. On the i5/OS command line, use the Add Physical File Member (ADDPFM) command to add a file member. Type `ADDPFM FILE(QUSRSYS/QAZLCSAVL) MBR(nwsdname) TEXT('description') EXPDATE(*NONE) SHARE(*NO) SRCTYPE(*NONE)`.

2. In the file member that you just created, list all the shares that you want to be able to save. List each share name that you defined for the server on a separate line. Share names can have embedded blanks. For example, if you defined cshare, dshare, eshare, fshare, gshare, and my share as shares on LINSVR1, your member name LINSVR1 would look like this:

```

QUSRSYS/QAZLCSAVL LINSVR1
0001.00 cshare
0002.00 dshare
0003.00 eshare
0004.00 fshare
0005.00 gshare
0006.00 my share

```

**Note:** If you specify multiple share names that point to the same directory, the i5/OS operating system saves the data multiple times for a "save all" request. To avoid duplicating data when you save it, do not include multiple shares that include the same directory or data.

### Saving and restoring files for Linux in a logical partition:

Use the Save (SAV) CL command to save your files.

To be able to restore a directory or file by share name, you must specify that file or share name on the SAV command.

**Note:** To avoid duplicating data, specify each share only once. If you specify multiple share names that point to the same directory on the Linux server, i5/OS saves the data multiple times.

Do these steps to save your files.

1. Ensure that the Linux server is active. See in "Starting and stopping the NWSD for a Linux partition" on page 20).
2. Also ensure that the QSYSWRK subsystem, QSERVER, and TCP/IP are active. You can use the Work with Active Jobs (WRKACTJOB) command).
3. On the i5/OS command line, type SAV and press F4.
4. In the Device field, specify the device on which you want i5/OS to save the data. For example, 'QSYS.LIB/TAP01.DEVD' saves the data to tape.
5. In the Object field, specify what you want i5/OS to save in the form '/QNTC/*servername*/*sharename*'. You can use wildcard characters. See "Examples: Saving files for Linux in a logical partition" for how to specify particular parts of the Linux server.
6. Use the Directory subtree field to specify whether you want to save subtrees under a directory. The default is to save all directories.
7. To specify that you want to save changes since the last save, specify \*LASTSAVE in the Change period field. You can also specify a specific range of dates and times.
8. Press Enter to save the shares that you specified.

*Examples: Saving files for Linux in a logical partition:*

These examples show how to use the save (SAV) or restore (RST) commands for specific parts of a Linux server in a logical partition.

Here are examples for server *server1*, where *server1* is the name of the Linux server.

| To save or restore this:         | Specify this:                                 |
|----------------------------------|-----------------------------------------------|
| All server objects.              | OBJ('/QNTC/*') SUBTREE(*ALL)                  |
| All objects for <i>server1</i> . | OBJ('/QNTC/ <i>server1</i> /*') SUBTREE(*ALL) |

| To save or restore this:                                                                                                                                                    | Specify this:                                                                               |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| All objects for <i>server1</i> that changed since you last saved the files.                                                                                                 | OBJ('/QNTC/server1/*') SUBTREE(*ALL) CHGPERIOD(*LASTSAVE)                                   |
| All objects for <i>server1</i> that changed during a certain period (in this case between 10/19/99 and 10/25/99).                                                           | OBJ('/QNTC/server1/*') SUBTREE(*ALL) CHGPERIOD('10/19/99' '00:00:00' '10/25/99' '23:59:59') |
| All directories, files, and shares to which a particular share (for example, 'fshare') refers. i5/OS does not save and restore the directory over which the share is built. | OBJ('/QNTC/server1/fshare/*') SUBTREE(*ALL)                                                 |
| Only files to which the specified share (for example, 'fshare') refers that match the specified pattern (pay*). i5/OS does not save directories nor shares.                 | OBJ('/QNTC/server1/fshare/pay*')                                                            |
| Only directories and shares (no objects) for 'fshare' and its immediate children.                                                                                           | OBJ('/QNTC/server1/fshare') SUBTREE(*DIR)                                                   |
| Directories, shares, and files for 'terry' and its subtrees (not directory 'terry').                                                                                        | OBJ('/QNTC/server1/fdrive/terry/*') SUBTREE(*ALL)                                           |
| Only the specific file 'myfile.exe'.                                                                                                                                        | OBJ('/QNTC/server1/gdrive/myfile.exe')                                                      |

## Building a rescue image on an NWSSTG

One rescue solution for a logical partition is to create a small network storage space (NWSSTG) that can be left on the IFS solely for the purpose of rescuing logical partitions.

First, a quick explanation of what a Linux Rescue server normally entails. Many distributors include a rescue image on their installation disks that will start into a minimal version of their distribution that contains all the basic diagnostic tools, drivers and other utilities that would be useful for checking a previously existing logical partition. You can simplify this process by creating a network storage space with a rescue server in it at the time you install your logical partition.

Before creating a rescue image on network storage it is important to document the configuration information for each of your logical partitions. You will want to document the drive configuration which can be obtained from the `/etc/fstab` file. You will also want to capture the networking information that is reported when you run the command `"ifconfig"`. It would also be good to create a list of the modules that are needed by each partition. You can see which modules are in use by using the `"lsmod"` command from within Linux. It is recommended that you take the information obtained from the commands and files listed above and put them into files that can be stored on your rescue network storage space.

To create the rescue storage space, first consult your Linux documentation to see what amount of space is required for a minimum installation. Create a network storage space (CRTNWSSTG) that is slightly larger than the value listed in the distributor's documentation. For example if the documentation says a minimum server installation is 291 megabytes then you would want to create a storage space of 425 megabytes. Making the storage space larger allows space to create a swap partition, a PrEP start partition and allows room for you to install any extra software that you would like to have available in your rescue image. You may want to make a note in the storage space's description field telling what distribution was used to make the rescue image and warn that it should be saved. Next, link this storage space to a Network Storage Descriptor (NWSD). You don't need to create a new NWSD for this step. You could unlink an existing storage space and temporarily link your rescue storage space to any of your existing NWSD's.

Next, start your distributor's installation server as described in the documentation and follow the prompts. Make sure that you create a PrEP start partition if you choose to manually partition your install. When you get to the distributors package group selection display you are going to want to select the minimum number of packages supported. The name for the package group varies by distribution. Finally, let the installer complete its package installation and configuration.

After installation has finished, the installer will start up the server for you. At this time you should take some time to verify that the rescue server has all the utilities that you are going to need. For a logical partition, you are going to want to do an "rpm -qa | grep ibmsis" to make sure that the utilities to work with native disk are available. You should also make sure that the device drivers (i.e., pcnet32 for Ethernet and Olympic for token ring) that your partitions may require are installed. The kernel modules that have been compiled can be found in the /lib/modules/kernel version/kernel/drivers directory structure. You should also install any other special drivers or software packages that your logical partitions may require and ftp the files with the configuration information for your other logical partitions onto the rescue server network storage space at this time.

On some Linux Distributions, you will have to manually install the kernel after installation. Consult the appropriate installation documentation for details regarding installing the kernel.

Since you will be starting the rescue network storage space from network storage you need to make note of the path to the root partition on the rescue storage space. You can determine the root partition by doing the command 'cat /etc/fstab'. The partition that has "/" in the second column is your root partition. The distributor's documentation should also provide assistance in determining the root partition.

You are now done creating your rescue image. Shutdown your logical partition using "shutdown -h now" and vary the partition off after the shutdown has completed. After the partition has varied off unlink the rescue storage space and relink the NWSSTG's normal storage space.

## Using a rescue image from an NWSSTG

Do these steps to use a rescue image for a network server storage space (NWSSTG) that is associated with a Linux server in a System i logical partition.

A rescue image contains the minimal version of all the basic diagnostic tools, drivers and other utilities that would be useful for checking a previously existing logical partition.

To use the rescue image that you created on NWSSTG, use the following steps:

1. Disconnect the failed partition's virtual storage space (if applicable) through WKRNNWSSTG.
2. Connect your rescue storage space as the first drive to the NWSSTG and reconnect the original storage space (where applicable) as the second drive.
3. Edit the NWSSTG for the failed partition to start from IPL Source \*NWSSTG. Also edit the IPL Parameters Field to reflect the root partition on the rescue storage space. For most distributions this is a parameter like 'root=/dev/hda3' or 'root=/dev/vda1' Refer to your Linux documentation for assistance or look at other running partitions using the command 'cat /proc/iSeries/mf/B/cmdline'.
4. Boot the partition.
5. If the existing root partition is on native disk, you may need to insert the ibmsis driver using the command "insmod ibmsis".
6. Create a mount point to which you will mount the root partition of the network storage space you are trying to rescue. You can use a command like "mkdir /mnt/rescue".
7. Now mount the root partition of the network storage space you are trying to rescue. You mount a drive using the command "mount -t your partition's type partition's location mount point" where the partition's type is the format of the partition like ext2 or reiserfs, the partition's location is something like /dev/hdb3 (for non-devfs), /dev/hd/disc1/part3 (for devfs server) or /dev/sda2 in the case of a partition on native disk. It is important to note that the drive you are trying to rescue,

when using virtual disk, will be the second drive rather than the first drive. (i.e. If the drive was /dev/hda3 when the partition was running normally, it will be /dev/hdb3 in the rescue server.) Here again, your Linux documentation or the configuration files you created when you created the rescue NWSSTG will help you determine the device for the root of the partition you are trying to rescue. Finally, your mount point will be something like /mnt/rescue if you use the example from above.

After following the steps listed above, you can either use the rescue tools provided in your rescue storage space against the mount point you created or you can change root to that partition using "chroot mount point" to work on the partition from within its own storage space. See "Building a rescue image on an NWSSTG" on page 52 for more information.

## Backing up NWSDs for a logical partition

When you save the storage space objects that are associated with a logical partition virtual disks on a System i product, you also need to save the network server description (NWSd). Otherwise, a logical partition may not be able to reestablish items such as the partitions file system permissions.

To save the NWSd, you use the Save Configuration (SAVCFG) command:

1. On the i5/OS command line, type SAVCFG.
2. Press **Enter** to have i5/OS save the NWSd configuration.

**Attention:** The save configuration command (SAVCFG) saves the objects associated with an NWSd and the current static network server storage spaces. It does not save the links associated with the dynamically added storage spaces. These need to be added manually once the configuration and the dynamically linked storage spaces have been restored.

## Restoring NWSDs for a logical partition

Do these steps to restore a network server description (NWSd) for a Linux server in a System i logical partition.

In a disaster recovery situation, you would restore all the configuration objects, which include the network server description (NWSd) for your logical partition.

In some situations, for example when you migrate to new hardware, you need to specifically restore the NWSd. To have i5/OS automatically relink disk drives within the integrated file system to the restored NWSd, restore those disk drives first.

To restore the NWSd, you use the Restore Configuration (RSTCFG) command:

1. On the i5/OS command line, type RSTCFG and press **F4** (Prompt).
2. In the Objects field, specify the name of the NWSd.
3. In the Device field, specify the device name if you are restoring from media. If you are restoring from a save file, specify \*SAVF and identify the name and library for the save file in the appropriate fields.
4. Press **Enter** to have i5/OS restore the NWSd.

When you are done restoring the NWSd and all its associated storage spaces, start (vary on) the logical partition. See "Starting and stopping the NWSd for a Linux partition" on page 20 for more information.

## Unlinking disk drives from a logical partition

Use these tasks to unlink Linux virtual disk drives (network server storage spaces) from a logical partition.

Unlinking disks disconnects them from the logical partition, making them inaccessible to users. You also need to unlink disk drives if you are uninstalling a logical partition.

### Unlinking disk drive with System i Navigator:

To unlink disk drives from a logical partition, do the following:

1. Vary off the NWSD for your logical partition. See “Starting and stopping the NWSD for a Linux partition” on page 20 for more information.
2. Click **Network** → **Integrated Server Administration** → **All Virtual Disks**.
3. Right-click the name of the disk drive that you want to remove.
4. Click **Remove Link**.
5. Select a server from the list of linked servers.
  - a. If you are unlinking a disk drive that you plan to relink, uncheck **Compress link sequence**. You need to relink the disk drive as the same link sequence number before you vary on the server. By preventing compression of the link sequence values, you avoid having to unlink and relink all the disk drives to get them in the correct sequence.
6. Click **Remove**.

**Unlinking disk drive with CL commands:**

To unlink disk drives from a logical partition using the Remove Server Storage Link (RMVNWSSTGL) CL command, do the following:

1. Vary off the NWSD for your logical partition. See “Starting and stopping the NWSD for a Linux partition” on page 20 for more information.
2. Type RMVNWSSTGL and press **F4**.
3. Type the name of the storage space that you want to unlink in the Network server storage space field. Press **Enter**
4. Type the name of the server from which you want to unlink the storage space in the Network server description field. Press **Enter**
  - a. If you are unlinking a linked disk drive that you plan to relink, specify \*N0 in the Renumber field. You need to relink the disk drive as the same sequence number before you vary on the server. By preventing automatic renumbering, you avoid having to unlink and relink all the disk drives to get them in the proper sequence.
5. Press **Enter**

If you are uninstalling a logical partition, your next step is to delete the disk drive. See “Deleting virtual disks for Linux running in logical partitions” on page 35 for more information. Otherwise, vary on the NWSD for your logical partition. See “Starting and stopping the NWSD for a Linux partition” on page 20 for more information.

**Objects to save and their location on i5/OS**

Use this list to determine which i5/OS objects to save for a Linux server in a System i logical partition

Many objects are created as a result of installing a logical operating system in a partition and using virtual storage. Some of these objects are server-related, others user-related. You need to save them all if you want to restore properly. You can save these objects by using options of the i5/OS GO SAVE command in the server. Option 21 saves the entire server. Option 22 saves server data (which includes objects in QUSRSYS). Option 23 saves all user data (which includes objects in QFPNWSSTG).

If you want to save a particular object, use this table to see the location of that object on i5/OS and the command to use. Manually saving parts of your server has more information about using the save commands.

*Table 12. Objects to save*

| Object content                                  | Object name | Object location | Object type | Save command |
|-------------------------------------------------|-------------|-----------------|-------------|--------------|
| <b>For logical partitions with virtual disk</b> |             |                 |             |              |

Table 12. Objects to save (continued)

| Object content                              | Object name | Object location | Object type                                            | Save command                                                                                  |
|---------------------------------------------|-------------|-----------------|--------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| Guest partition and virtual disk drive      | stgspc      | /QFP NWSSTG     | User-defined network server storage spaces in server   | GO SAVE, option 21 or 23<br>SAV OBJ('/QFPNWSSTG /stgspc')<br>DEV('/QSYS.LIB/ TAP01.DEVD')     |
|                                             |             |                 | User-defined network server storage spaces in user ASP | SAV OBJ((' /QFPNWSSTG /stgspc') ('/dev/QASPnn /stgspc.UDFS'))<br>DEV('/QSYS.LIB /TAP01.DEVD') |
|                                             |             |                 | <b>For all logical partitions with a server</b>        |                                                                                               |
| Messages from the logical partition         | Various     | Various         | Server message queue                                   | GO SAVE, option 21 or 23<br>SAVOBJ OBJ(msg)<br>LIB(qlibrary)<br>DEV(TAP01)<br>OBJTYPE(*MSGQ)  |
| i5/OS config objects for logical partitions | Various     | QSYS            | Device config objects                                  | GO SAVE, option 21, 22, or 23<br>SAVOBJ DEV (TAP01)                                           |
| Various                                     | Various     | QUSRSYS         | Various                                                | GO SAVE, option 21 or 23<br>SAVLIB LIB(*NONSYS) or LIB(*ALLUSR)                               |

## Starting a logical partition

Use the following topics to learn how to start a logical partition.

### Deciding what IPL type to use when running Linux in a logical partition

The IPL type parameter on the network server description (NWS D) determines the source from which a logical operating system is loaded.

The first thing to understand is that on a System i product the initial program load (IPL) comes from one of four locations managed by the LPAR configuration. These four locations are called A, B, C, and D. The initial program (which in the case of Linux is the Linux kernel) can be loaded into these locations two ways:

- From Linux itself
- The C location can be loaded by i5/OS using a network server description (NWS D).

The NWS D itself has a parameter called IPLSRC that specifies from where the logical partition is to be loaded. This parameter can have the following values:

| IPLSRC value | Description                                                                                                                                                                            |
|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| A            | Load the logical partition from location A. Location A must have been previously loaded with a kernel from Linux. IBM recommends using the A location to store a stable, known kernel. |
| B            | Load the logical partition from location B. Location B must have been previously loaded with a kernel from Linux. IBM recommends using B for testing new kernels                       |
| D            | IBM support reserves this type of IPL for future use.                                                                                                                                  |
| *Panel       | The partition is started from the source indicated on the operator's panel.                                                                                                            |

| IPLSRC value                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| *NWSSTG (network server storage space) | This IPL type is used to start a partition from a virtual disk. i5/OS will find the kernel in the virtual disk and load it in location C. Then the partition will be set to IPL from C. i5/OS searches the first virtual disk connected to the NWSD for a partition marked startable, and of type 0x41 (PReP start.) The contents of this location are loaded in C. If a partition of this type does not exist, the partition will fail. |
| *STMF (stream file)                    | This IPL type is used to start a partition from a kernel loaded in the i5/OS Integrated File System. Note that the Integrated File System includes files on the Optical (CD) drive on i5/OS. i5/OS will load the specified file into the C location and the partition will be set to IPL from C. This is typically how initial Linux installations are performed.                                                                        |

## Booting from the LPAR configuration display

The LPAR configuration display can be used to set a logical partition to IPL from four locations. These locations are A, B, C, and D. Currently D is reserved for future use. The IPL type determines which copy of programs your server uses during the initial program load (IPL). These parameter can have the following values:

| IPLSRC values | Description                                                                                                                                                                                      |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| A             | Load the logical partition from location A. Location A must have been previously loaded with a kernel from Linux. IBM recommends using the A location to store a stable, known kernel.           |
| B             | Load the logical partition from location B. Location B must have been previously loaded with a kernel from Linux. IBM recommends using B for testing new kernels.                                |
| C             | This IPL type is typical used when a partition is loading from a network server description (NWSD) or a stream file (STMF). Refer to *NWSSTG for more information on how the C IPL type is used. |
| D             | IBM support reserves this type of IPL for future use.                                                                                                                                            |

## Booting from the A and B IPL source for Linux running in a logical partition

A logical partition running Linux on System i models 8xx and 270 can start from either the A or B IPL source.

Linux must be installed on the server and previously loaded from a different source before this option can be used. Administration authority must be obtained before performing this task. For more information on how to configure user profiles, refer to Logical partition authority for more information. To copy a kernel into the A or B IPL source, the /proc file system is used. The command used to copy the kernel into the A IPL source is `dd if=/path/to/vmlinux of=/proc/iSeries/mf/A/vmlinux bs=4096`. The command used to copy the kernel into the B IPL source is `dd if=/path/to/vmlinux of=/proc/iSeries/mf/B/vmlinux bs=4096`. The advantage of using this start source is that Linux starts faster. A disadvantage is that the start source cannot be saved or restored. It is also difficult to tell which IPL source has stored the kernel. If you need to pass arguments to the kernel during start, you can modify the kernel command line parameters either through the Work with partitions service tool on your i5/OS primary partition or through the /proc file system under Linux.

Follow these steps to set the kernel command line:

1. At the i5/OS command line, type STRSST and press **Enter**.
2. From the Start Service Tools (STRSST) Sign On display, provide the Service Tools user ID and password and press **Enter**. For more information on how to configure user profiles, refer to Logical partition authority for more information.
3. From the System Service Tools (SST) display, select option 5 (Work with servers) and press **Enter**.

4. From the Work with System Partitions display, select option 3 (Work with partition configuration) and press **Enter**.
5. From the Work with Partition configuration display, select option 14 (Change logical commands) next to the logical partition you want to modify and press **Enter**.
6. From the Change Guest Environment Command Line Parameters display, enter the new kernel command and press **Enter**.

To set the kernel command line through the /proc file system for the A IPL source, use a variation of the following command: `echo -n "root=/dev/iSeries/vda1" >> /proc/iSeries/mf/A/cmdline`

To set the kernel command line through the /proc file system for the B IPL source, use a variation of the following command: `echo -n "root=/dev/iSeries/vda1" >> /proc/iSeries/mf/B/cmdline`

## Recovering directly attached disks for servers running Linux in a logical partition

Do these steps to recover disks that are directly attached to a server running Linux in a logical partition on System i models 8xx or 270.

When you install Linux in a logical partition with directly attached disk resources, the kernel is loaded to the A and B slots, not to \*NWSSTG. When you recover a Linux partition in this configuration (for example, when you recover primary partition configuration data after you scratch the primary partition), you must change the network server description of the Linux partition in order to start from a kernel in the integrated file system (\*STMF). After you start the partition, you can use Linux commands to rewrite the Linux kernel and enable the partition to be startable.

To start a Linux partition on a directly attached disk and then write the kernel so that you can start the system on subsequent re-starts, follow these steps:

1. Load a kernel in your integrated file system.
2. Run the Work with Configuration Status (WRKCFGSTS) command, and specify network server (\*NWS) as the configuration type. A list of network server descriptions appears.
3. Type 8 to display the description in the Option column next to the Linux partition you want to start. A description of that network server appears.
4. Type 2 to edit in the Option column next to Start source, Start stream file, and Start parameters.
5. Edit those fields so that they contain the following values:
  - IPL source = \*STMF
  - IPL stream file = *<IFS path to your kernel>*, (for example, /home/kerne1s/vmlinux64)
  - IPL parameters = `root=/dev/sdaX`, where *X* is equal to the disk partition that your root file system is on. This is often sda1 or sda2.
6. Save these changes.
7. Connect to the virtual console. Enter `telnet <your machine> 2301`.
8. Vary on the partition.
9. Log in to the partition.
10. If your server does not have a copy of the kernel that you want to install, FTP it on to your server. You could even NFS share it off of the i5/OS.
11. Write the kernel to your B slot by using the following command:
 

```
dd if <path to your kernel name> = of=/proc/iSeries/mf/B/vmlinux bs=4096
```

This command may take a while to complete.

12. Set the command line to the same value you used for the Start parameters above. First, you may want to echo a large amount of space into the cmdline file to make sure that any previous parameters are cleared. To do so, use the following command:
 

```
echo root=/dev/sdax > /proc/iSeries/mf/B/cmdline
```

13. Echo your command line:  

```
echo root=/dev/sdaX > /proc/iSeries/mf/B/cmdline
```
14. Verify that the command line was successfully entered by using the following command:  

```
cat /proc/iSeries/mf/B/cmdline
```

You should see the same thing you echoed in above.
15. Shut down the partition by using the following command:  

```
shutdown -h now
```
16. Edit your network server description by using the Work with Configuration Status (WRKCFGSTS) command as described in steps 2 and 3 above.
17. Set the Start source = B and the Start stream file = \*NONE.
18. Now you can vary on your partition and use the kernel installed in the System i server slot B.

---

## Troubleshooting Linux running in a logical partition

Use this topic to understand how to analyze and resolve errors specific to Linux running in a logical partition.

If your error is specific to logical partitions, refer to Troubleshooting logical partitions for additional assistance. Linux specific troubleshooting problems require the assistance of your Linux distributor.

## Debugging NWSD error messages for Linux running in a logical partition

Understand and debug your NWSD error messages for Linux running in logical partitions.

You could encounter error messages when you try to vary on a Linux partition. These error messages will appear if you provide information when you create your NWSD that does not apply to a logical partition running on the server. All error messages related to the NWSD should appear in QSYSOPR indicating a description of the problem and a resolution to the issue.

*Table 13. NWSD error messages*

| Reason codes | Code explanation                                                                                                                                                                                                                                                                                                                                |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 00000001     | *NWSSTG was specified as the IPL source, but no storage space was found.                                                                                                                                                                                                                                                                        |
| 00000002     | The partition specified in the PARTITION parameter was not found.                                                                                                                                                                                                                                                                               |
| 00000003     | The partition specified in the PARTITION parameter is not a GUEST partition.                                                                                                                                                                                                                                                                    |
| 00000004     | There is already an NWSD in the i5/OS partition that is active and using the partition specified in the PARTITION parameter of the NWSD.                                                                                                                                                                                                        |
| 00000005     | The partition specified in the PARTITION parameter of the NWSD is turned on (perhaps through the LPAR configuration interface or from another i5/OS partition.)                                                                                                                                                                                 |
| 00000006     | The partition is set to start from a stream file (stmf) and for some reason that didn't work. You should note that the user performing the vary on need read access to the IPL STMF parameter.                                                                                                                                                  |
| 00000007     | The NWSD is set to start from a storage space (NWSSTG), but for some reason the kernel could not found. Some common reasons are no type 0x41 partition or not marked startable.                                                                                                                                                                 |
| 00000008     | The partition would not start. There are a variety of reasons why the partition will not start. You could have a corrupt kernel or the processor feature code does not support the shared processor pool. If the kernel and processor are not the problem, you need to look at the information for this partition and start reviewing the SRCs. |

Table 13. NWSD error messages (continued)

| Reason codes | Code explanation                                                                                                                                                                                                                                                                                                              |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 00000009     | The partition identified as the logical partition is not configured. To select a host partition, use the Work with Partition Status display and press <b>F11</b> (Work with partition configuration). From the Partition Configuration display, select option <b>13</b> next to the partition to assign the partition a host. |
| 00000010     | Contact your next level of support to find a proper solution to the problem.                                                                                                                                                                                                                                                  |
| 00000011     |                                                                                                                                                                                                                                                                                                                               |
| 00000080     |                                                                                                                                                                                                                                                                                                                               |
| 00001088     | The kernel does not appear to be valid. This error is frequently caused if you do not FTP the kernel in binary mode.                                                                                                                                                                                                          |
| 00001089     |                                                                                                                                                                                                                                                                                                                               |
| 0000108A     |                                                                                                                                                                                                                                                                                                                               |
| 0000108B     | The kernel does not appear to be compatible with the version of i5/OS in the primary partition.                                                                                                                                                                                                                               |
| 0000108C     |                                                                                                                                                                                                                                                                                                                               |
| 000010A3     | There is an insufficient amount of processors assigned to the partition or there are not enough shared processors available.                                                                                                                                                                                                  |
| 000010A9     |                                                                                                                                                                                                                                                                                                                               |
| 000010AA     |                                                                                                                                                                                                                                                                                                                               |
| 000010A4     | There is an insufficient amount of memory available for the partition.                                                                                                                                                                                                                                                        |
| 000010A5     |                                                                                                                                                                                                                                                                                                                               |
| 000010AE     | This error will occur on servers that only support dedicated processors when either you have specified a shared processor for a Linux partition or you have the QPRCMLTTSK server value set to 1.                                                                                                                             |

## Debugging a processor multitasking error

In order to run Linux on certain System i products, processor multitasking must be disabled in the i5/OS operating system.

Processor multitasking causes the System i processor to cache information when switching between tasks. This function is not supported by Linux in certain System i products. An IPL of the server is required to activate the change. To disable processor multitasking from an i5/OS command line, change the QPRCMLTTSK server value to 0 (CHGSYSVAL QPRCMLTTSK '0') in the primary partition. If a logical partition is started without disabling the multitasking function, the IPL of the logical partition will fail with server reference code (SRC) B2pp 8105 000010AE.

Before creating a logical partition, disable processor multitasking on servers with the following feature codes:

| Servers | Feature codes requiring QPRCMLTTSK (0) |
|---------|----------------------------------------|
| 820     | 2397<br>2398<br>2426<br>2427           |
| 830     | 2400<br>2402<br>2403<br>2351           |

| Servers | Feature codes requiring QPRCMLTTSK (0) |
|---------|----------------------------------------|
| 840     | 2418<br>2420<br>2416<br>2417<br>2419   |

## System reference codes (SRCs) for Linux running in a logical partition

The following list contains SRCs and suggested corrective actions for Linux running in logical partitions and suggested corrective actions.

To display SRCs, see “Displaying the reference code history for secondary partitions” on page 22 for instructions on how to list the last 200 reference codes for a partition. Additional SRCs for logical partitions can be found at Resolve server reference codes for logical partitions. If an SRC is not listed, it may not be related to logical partitions. You should consult the Troubleshooting documentation or your next level of service.

System reference codes (SRCs) for logical partitions explains what SRCs are and how to identify them.

### Reference codes

#### B2pp 1270 (pp equals the partition ID)

**Cause:** A problem occurred during the IPL of a secondary partition running Linux. The partition cannot IPL because the primary partition must be in a full paging environment.

**Recovery:** IPL the primary partition past the Storage Management full paging IPL step.

**Problem analysis procedure:** Check the IPL mode of the primary partition and the failing secondary partition. If the primary partition is in C mode the logical partition will not IPL.

#### B2pp 6900 (pp equals the partition ID)

**Cause:** The size fo the secondary partition’s kernel exceeds the size allocated for the load area by the secondary partition.

**Recovery:** Verify the size of the secondary partition’s memory allocation is large enough to load the kernel. Ensure the kernel is correct.

**Problem analysis procedure:** Identify the values for words 3 and 4 to determine the cause of the error. The values for the words are as follows:

- Word 3: Allocated size of the secondary partition
- Word 4: Required size of the kernel

#### B2pp 6905 (pp equals the partition ID)

**Cause:** A problem occurred during the IPL of a secondary partition running Linux. The kernel is not valid for the specified IPL mode for the logical partition. The kernel is not available for the IPL mode.

**Recovery:** Verify that the kernel specified to load is valid and the IPL mode specified is where the kernel is located.

**Problem analysis procedure:** Check processor and memory allocations to the secondary partition. You should make sure there are enough functioning processors and memory resources in the server for the partition.

#### B2pp 6910 (pp equals the partition ID)

**Cause:** A problem occurred during the IPL of a secondary partition running Linux. A storage management problem occurred while loading the logical partition.

**Recovery:** Collect the detailed hexadecimal data from the server and contact your next level of support.

**Problem analysis procedure:** Check processor and memory allocations to the secondary partition. You should make sure there are enough functioning processors and memory resources in the server for the partition.

**B2pp 6920 (pp equals the partition ID)**

**Cause:** A problem occurred during the IPL of a secondary partition running Linux. A problem occurred while loading the logical partition.

**Recovery:** Collect the detailed hexadecimal data from the server and contact your next level of support.

**Problem analysis procedure:** Review the SRC history of the server.

**B2pp 6930 (pp equals the partition ID)**

**Cause:** A problem occurred during the IPL of a secondary partition running Linux. A problem occurred while loading the logical partition.

**Recovery:** Collect the detailed hexadecimal data from the server and contact your next level of support.

**Problem analysis procedure:** Check processor and memory allocations to the secondary partition. You should make sure there are enough functioning processors and memory resources in the server for the partition.

**B2pp 8105 (pp equals the partition ID)**

**Cause:** Initialization of secondary partition main storage data structures failed. The IPL is ended.

**Recovery:** This is likely a problem with the load source media being corrupt or invalid. Re-installation of the secondary partition is probably required to recover. If it continues to fail, contact your service provider.

**Problem analysis procedure:** Identify the reason code from word 13 of the SRC. The reason code value for word 13.

- 000000AE: Processor multitasking must be disabled in the i5/OS operating system.

## Recovering from a virtual Ethernet network failure

Learn how to recover from a virtual Ethernet network failure after a kernel upgrade for Linux running in logical partitions.

If you have upgraded to a 64-bit kernel (from a 32-bit kernel less than 2.4.10) or to a 32-bit kernel level greater than 2.4.10, your network device information may have changed. The notation of the virtual network devices in Linux has changed from vethXY to ethXY.

To understand what network devices relates to a respective virtual Ethernet line description, you can view the display message log using this command:

```
dmesg | fgrep veth | less
```

The output of this command should produce messages similar to the following:

```
veth.c: Found an Ethernet device eth0 (veth=0) (addr=c000000000ff2800)
```

In this case, the message is telling you that the i5/OS virtual line description veth0 now relates to a Linux net device of eth0, and veth=0 correlates to VLAN0 in i5/OS. If for some reason your display message log is overrun, you can also analyze the network devices in the proc filesystem with the following command:

```
cat /proc/iSeries/veth/[netdevice]
```

Running this command with a valid netdevice will produce output similar to this:

```
Net device: c00000000ff2800
Net device name: eth0
Address: 0201FF00FF01
Promiscuous: 0
All multicast: 0
Number multicast: 0
```

This file is telling you that the Linux net device eth0 maps to the first virtual LAN on the i5/OS which is VLAN0.

---

## Related information for Linux in a logical partition

Product manuals, IBM Redbooks publications, Web sites, and other information center topic collections contain information that relates to the Linux in a logical partition topic collection. You can view or print any of the PDF files.

### IBM Redbooks publications

Linux on the IBM iSeries™ Server: An Implementation Guide 

This IBM Redbooks publication begins with an overview of Linux, defines what open source means, and explains why using Linux on the System i platform is beneficial. It discusses the basic server administration tasks and Linux application development to help you manage your server and develop Linux applications on the System i platform. It also introduces a wide range of services, such as Firewall, Apache, Samba, and e-mail, and explains the capabilities of each.

### Web sites

- Linux on the System i platform  ([www.ibm.com/systems/i/os/linux/](http://www.ibm.com/systems/i/os/linux/))  
This page contains information and links to install Linux on your System i product.
- System i Access for Linux.  ([www.ibm.com/systems/i/software/access/linux/](http://www.ibm.com/systems/i/software/access/linux/))  
The i5/OS ODBC Driver for Linux allows you to access the i5/OS database from a Linux client. This drive is included in System i Access for Linux

### Other information

- Logical partitions



---

## Appendix. 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.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation  
Licensing  
2-31 Roppongi 3-chome, Minato-ku  
Tokyo 106-0032, Japan

**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.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation

Software Interoperability Coordinator, Department YBWA  
3605 Highway 52 N  
Rochester, MN 55901  
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, IBM License Agreement for Machine Code, or any equivalent agreement between us.

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

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

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

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

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

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

#### COPYRIGHT LICENSE:

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

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (IBM) (2006). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. 2006. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

---

## Programming Interface Information

This Linux in a Logical Partition documents intended Programming Interfaces that allow the customer to write programs to obtain the services of Linux.

---

## Trademarks

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

e(logo)server  
eServer  
IBM  
i5/OS  
iSeries  
Linux  
Power PC  
System i5  
System i

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

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

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

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

---

## Terms and conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

**Personal Use:** You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

**Commercial Use:** You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

---

## **Code license and disclaimer information**

IBM grants you a nonexclusive copyright license to use all programming code examples from which you can generate similar function tailored to your own specific needs.

SUBJECT TO ANY STATUTORY WARRANTIES WHICH CANNOT BE EXCLUDED, IBM, ITS PROGRAM DEVELOPERS AND SUPPLIERS MAKE NO WARRANTIES OR CONDITIONS EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT, REGARDING THE PROGRAM OR TECHNICAL SUPPORT, IF ANY.

UNDER NO CIRCUMSTANCES IS IBM, ITS PROGRAM DEVELOPERS OR SUPPLIERS LIABLE FOR ANY OF THE FOLLOWING, EVEN IF INFORMED OF THEIR POSSIBILITY:

1. LOSS OF, OR DAMAGE TO, DATA;
2. DIRECT, SPECIAL, INCIDENTAL, OR INDIRECT DAMAGES, OR FOR ANY ECONOMIC CONSEQUENTIAL DAMAGES; OR
3. LOST PROFITS, BUSINESS, REVENUE, GOODWILL, OR ANTICIPATED SAVINGS.

SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF DIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, SO SOME OR ALL OF THE ABOVE LIMITATIONS OR EXCLUSIONS MAY NOT APPLY TO YOU.





Printed in USA