

IBM

@server

iSeries

Linux in a guest partition







@server

iSeries

Linux in a guest partition



---

# Contents

<b>Linux in a guest partition</b> . . . . .	1
Plan to run Linux in a guest partition . . . . .	1
Linux support on iSeries servers . . . . .	2
Hosted versus nonhosted guest partition running Linux . . . . .	5
Virtual I/O in a guest partition running Linux . . . . .	5
Directly attached I/O in a guest partition running Linux . . . . .	6
iSeries I/O adapters (IOAs) supported by Linux . . . . .	6
Obtain Linux for iSeries . . . . .	10
Directly attached SCSI support for a guest partition . . . . .	10
Device parity protection . . . . .	11
Start device parity protection . . . . .	11
Display device parity status . . . . .	12
DASD concurrent maintenance . . . . .	12
Add disk units using concurrent maintenance . . . . .	13
Disk hardware migration . . . . .	15
Order a new server or upgrade an existing server to run a guest partition . . . . .	16
Create a guest partition to run Linux . . . . .	16
Configure a guest partition . . . . .	17
Assign I/O adapters (IOAs) to a guest partition. . . . .	17
Create a network server description and a network server storage space . . . . .	21
Configure the LAN console for a guest partition . . . . .	23
Connect the virtual console . . . . .	23
Install Linux in a guest partition . . . . .	24
Install from CD-ROM images in the integrated file system . . . . .	24
Start and stop the NWSD . . . . .	25
Manage Linux in a guest partition . . . . .	26
Logical partition authority needed for a guest partition . . . . .	26
Display and change hardware configuration . . . . .	27
Display the console log for the guest partition . . . . .	27
Display the operating environment of secondary partitions . . . . .	27
Display the reference code history for secondary partitions . . . . .	28
Change the guest partition host information . . . . .	28
Connect your guest partition to the LAN . . . . .	28
Directly attached LAN adapters . . . . .	28
Virtual Ethernet . . . . .	29
Enable virtual Ethernet . . . . .	29
Create an Ethernet line description for virtual Ethernet . . . . .	29
Virtual Ethernet connectivity methods . . . . .	30
Print system configuration for logical partitions . . . . .	35
Dynamic movement of shared processor units . . . . .	36
Dynamically adding disks to a Linux partition . . . . .	37
Network server description . . . . .	40
Link NWSSTG among multiple NWSDs . . . . .	40
Delete NWSDs for a guest partition . . . . .	40
Delete disk drives for a guest partition . . . . .	41
Application support . . . . .	41
Access iSeries data using the Linux ODBC driver . . . . .	41
OS/400 NetServer support for Samba . . . . .	41
How to use SAMBA and NFS to access files in the Integrated File System . . . . .	42
Back up and recover a guest partition . . . . .	45
Backup options for virtual and directly attached disk. . . . .	45
Utilities that run within Linux . . . . .	45
OS/400 commands . . . . .	46

How to use the shared read-only storage space . . . . .	46
Sharing access to disks in read-write mode . . . . .	48
How to clone Linux servers in iSeries partitions quickly . . . . .	49
Save and restore hosted Linux files to and from the host's shared tape drive . . . . .	49
Save and restore hosted Linux files by using Linux tar . . . . .	50
Save and restore hosted Linux data by using OS/400 SAV and RST . . . . .	50
Troubleshooting Linux virtual tape . . . . .	51
Save guest partition configuration data . . . . .	51
Back up the NWSD and disk drives associated with a guest partition . . . . .	52
Build a rescue image on an NWSSTG . . . . .	52
Use a recovery system from an NWSSTG . . . . .	53
Back up NWSDs for a guest partition . . . . .	54
Restore NWSDs for a guest partition . . . . .	54
Unlink disk drives from a guest partition . . . . .	54
Objects to save and their location on OS/400 . . . . .	55
IPL a guest partition . . . . .	55
Decide what IPL type to use when running Linux . . . . .	55
Boot from the A and B IPL source . . . . .	56
Recovering directly attached disks . . . . .	57
Troubleshoot Linux running in a guest partition. . . . .	58
Debug NWSD error messages . . . . .	58
Debug the processor multitasking error . . . . .	59
System reference codes (SRCs) for Linux . . . . .	60
Recover from a virtual Ethernet network failure . . . . .	61

---

## Linux in a guest partition

IBM and a variety of Linux distributors have partnered to integrate the Linux operating system with the reliability of the iSeries server. Linux brings a new generation of web-based applications to the iSeries. IBM has modified the Linux PowerPC kernel to run in a secondary logical partition and contributed the kernel back to the Linux community. This topic contains information on how to plan a strategy that you can use to ensure Linux is successfully installed on your system, how to configure a guest partition to run Linux, and how to manage and troubleshoot a guest partition running Linux.

### **“Plan to run Linux in a guest partition”**

Find out what you need to do before you install Linux on the iSeries. Understand what software and hardware requirements are needed to support Linux. Find out what configuration options are available and which options fit your company's needs.

### **“Create a guest partition to run Linux” on page 16**

Understand how to configure a guest partition using system service tools (SST) and how to install Linux on the iSeries. You can also find information about supported I/O adapters (IOAs) and how to configure a network server description (NWSA).

### **“Manage Linux in a guest partition” on page 26**

Find the information you need to manage a guest partition running Linux. Use the information to understand what IPL types you can use and how Linux partitions can communicate with other partitions on the server.

### **“Troubleshoot Linux running in a guest partition” on page 58**

Use this information to analyze and resolve system reference codes (SRCs) that are specific to Linux. Find out if the system you own requires you to disable the multitasking function of your server's processor.

---

## Plan to run Linux in a guest partition

Before you “Create a guest partition to run Linux” on page 16 your iSeries for Linux, you need to carefully plan and evaluate your software and hardware resources. This section guides you through the planning process before you create a guest partition running Linux.

### **“Linux support on iSeries servers” on page 2**

Evaluate each iSeries model to determine if your hardware will support Linux.

### **“Hosted versus nonhosted guest partition running Linux” on page 5**

Understand the difference between a hosted and nonhosted partition running Linux.

### **“Virtual I/O in a guest partition running Linux” on page 5**

Learn how you can create a Linux partition using virtual I/O resources.

### **“Directly attached I/O in a guest partition running Linux” on page 6**

Learn how you can create a Linux partition using directly attached I/O resources.

### **“iSeries I/O adapters (IOAs) supported by Linux” on page 6**

Find a list of IOAs and device drivers that are supported by Linux in a guest partition.

### **“Directly attached SCSI support for a guest partition” on page 10**

Understand how the ibmsis device driver can be used to support directly attached SCSI devices.


### **“Obtain Linux for iSeries” on page 10**

Find out who partnered with IBM to provide the Linux PowerPC kernel for the iSeries.


**“Order a new server or upgrade an existing server to run a guest partition” on page 16**  
Contact IBM or an IBM Business Partner to order a new server or upgrade an existing server to run a guest partition.



## Linux support on iSeries servers



To successfully partition an iSeries to run Linux, the server requires specific hardware and software. The primary partition must run OS/400 V5R1 or V5R2 and be updated with the latest PTFs. The latest

Linux-related OS/400 PTFs can be found on Linux for iSeries.  Linux is not supported on a primary partition. Additional information on software function can be found at OS/400 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 OS/400 and Linux partitions. Other models require the use of dedicated processors for a Linux partition. Those same models also require you to “Debug the processor multitasking error” on page 59 for the whole system, including the primary partition.

The Linux operating system supports single processors or multiple processors. This choice is made at the time the guest 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 guest partition.

The following tables help you determine if your system supports Linux in a guest partition.



## Model 270

Feature code	Number of processors	Supports 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
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

## Model 820

Feature code	Number of processors	Supports LPAR	Supports Linux	Linux shared processor	Disable processor multitasking
2395	1	Yes	No	No	No
2396	1	Yes	No	No	No
2397	2	Yes	Yes	No	Yes
2398	4	Yes	Yes	No	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

### 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

### 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

### Model 890

Feature code	Number of processors	Supports LPAR	Supports Linux	Linux shared processor	Disable processor multitasking
0197	24	Yes	Yes	Yes	No
0198	32	Yes	Yes	Yes	No
2487	16/24	Yes	Yes	Yes	No
2488	24/32	Yes	Yes	Yes	No


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

Refer to Linux for iSeries  for the latest updates to Linux running on an iSeries.



## Hosted versus nonhosted guest partition running Linux

A hosted guest partition is dependent on a hosting OS/400 partition for I/O resources. The hosting OS/400 partition can either be a primary or a secondary partition. The I/O resources a guest partition can use from a hosting partition include disk, CD, and tape devices.

The hosted guest partition must be started from the hosting OS/400 partition using a network server description (NWSD). The NWSD is used to control the guest partition. The guest partition can only be active when the hosting partition is both active and out of restricted state. When OS/400 is in restricted state, all NWSDs are automatically varied off. An OS/400 partition can host multiple guest partitions. You need to perform proper capacity planning  to ensure that the hosting partition can support guest partitions.

You should IPL a hosted guest partition by varying on an NWSD object. You should not power on a hosted guest partition using the Work with Partitions Status display. If you power on the guest partition running Linux use the Work with Partitions Status display, all of the virtual I/O devices will be unavailable.

A nonhosted guest partition is not dependent on a hosting OS/400 partition for any I/O resources. The guest partition has its own disk units or the partition makes use of networking support to do a network boot. A nonhosted guest partition can be started even if the primary partition is not fully active. A nonhosted guest partition can be started from the Work with Partitions Status display.

## Virtual I/O in a guest partition running Linux

Virtual I/O resources are devices owned by the “Hosted versus nonhosted guest partition running Linux” OS/400 partition that provide I/O function to the guest partition. The iSeries Linux kernel and OS/400 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 guest partition through an OS/400 partition. The virtual console can be established to the hosting partition 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 system errors.

**Virtual disk unit** can only be provided by a hosting partition. Virtual DASD provides access to NWSSTG virtual disks from 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 hosting partition. By default, a Linux partition can see all CD 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 “Create a network server description and a network server storage space” on page 21 for more information on changing NWSD attributes.

**Virtual tape** provides access to the OS/400 tape drives from a guest partition. By default, a guest 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 “Create a network server description and a network server storage space” on page 21 for more information on changing NWSD attributes.

**Virtual Ethernet** provides the same function as using a 1 Gigabyte (GB) Ethernet adapter. A guest partition can use virtual Ethernet to establish multiple high speed inter-partition connections. OS/400 and Linux partitions can communicate with each other using TCP/IP over the virtual Ethernet communication ports. Refer to Use virtual Ethernet in a guest partition for more information on virtual Ethernet.

For a better understanding of how your company might use guest partition with virtual I/O resources, read Logical and guest partition scenarios.

## Directly attached I/O in a guest partition running Linux

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. You can allocate disk units, tape devices, optical devices, and LAN adapters to a guest partition running Linux.

You must have an NWSA to install Linux in a guest 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 guest partition.

For a better understanding of how your company might use a guest 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/Os in a guest partition. [➤](#)

## iSeries I/O adapters (IOAs) supported by Linux

IOAs can be assigned to a guest partition. The iSeries supports the following adapters in a guest partition running Linux.

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

Adapters(CCIN)	Feature codes	Description	Linux device driver	SubSystem vendor/ SubSystem type	PCI vendor ID/PCI ID
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
2766	0612	Shortwave fibre channel adapter for point-to-point or arbitrated loop topologies; supporting tape and DASD attachment in OS/400; supporting DASD 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

Adapters(CCIN)	Feature codes	Description	Linux device driver	SubSystem vendor/ SubSystem type	PCI vendor ID/PCI ID
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
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
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 Impedance 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 Impedance Matching) and CIM packages.	icom	1014/0252	1014/0219
2838	0607	100/10-Mbps Ethernet	pcnet32	1014/0133	1022/2000


Adapters(CCIN)	Feature codes	Description	Linux device driver	SubSystem vendor/ SubSystem type	PCI vendor ID/PCI ID
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)	ibmsis	1014/0266	1069/B166

The Linux acenic, olympic, and pcnet32 device drivers are available with Linux. 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 tape, optical, and disk unit devices. This device driver is available from your preferred Linux distributor as a kernel module.

The most current version of this device driver can be found at the Linux for iSeries  website.

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 system 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-systems across fibre channel. The most current version of this device driver can be found at the Linux for iSeries  website. 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> website.



This list may change between OS/400 release boundaries. Refer to the Linux for iSeries  website.



## Obtain Linux for iSeries

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

Currently, IBM is working with the following major distributors of Linux: Red Hat Linux, SuSe Linux, and Turbolinux. All Linux distributions share a similar Linux kernel and development library. Linux distributors provide custom components that ease the installation and maintenance of Linux systems. Before installing another distributor's version of Linux, verify that the kernel has been compiled for the Power PC and the iSeries hardware. Otherwise, your system might be misconfigured and will not run Linux in a guest partition.

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

## Directly attached SCSI support for a guest partition

The ibmsis device driver is available to enable support for directly attached SCSI disk, tape, and optical devices. Disk compression is not supported under Linux. The ibmsis device driver, once installed, has two utilities that are used for "iSeries I/O adapters (IOAs) supported by Linux" on page 6. The sisconfig utility is used to configure directly attached disk and to perform error recovery. You can use this utility to display disk hardware status, start and stop device parity protection, initialize and format a disk, and perform other service actions. The sisupdate utility is used to update firmware on supported adapters. These utilities are run only in a guest partition running Linux. Root user authority is required to use these commands.

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 guest partition, contact your next level of support.

Directly attached SCSI drives in a guest 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	ibmsis
SCSI disk driver in base Linux kernel	sd.o
SCSI tape driver in base Linux kernel	st.o
SCSI CD-ROM driver in base Linux kernel	sr.o
SCSI driver in base Linux kernel	scsi_mod.o



## Device parity protection

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. To protect data, the disk input/output adapter (IOA) calculates and saves a parity value for each bit of data. Conceptually, the IOA computes the parity value from the data at the same location on each of the other disk units in the device parity set. When a disk failure occurs, the data can be reconstructed by using the parity value and the values of the bits in the same locations on the other disks. The system continues to run while the data is being reconstructed. The overall goal of device parity protection is to provide high availability and to protect data as inexpensively as possible. For more information on device parity protection as it relates to backup and recovery, refer to Device parity protection.

Guest partitions running Linux on the iSeries also support device parity protection. You can stripe data on parity protected disk units to improve performance using Logical Volume Manager (LVM) or Linux Software RAID 0 using the md driver. LVM controls disk resources by mapping data between a simple and flexible logical view of storage space and the actual physical disks. For more information on Logical Volume Manager or Linux software RAID, refer to the appropriate Linux HOWTO documentation. 

**Start device parity protection:** Usually, you start device parity protection when you attach a new disk subsystem to your system. You must start device parity protection 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, boot the guest partition into rescue mode prior to the installation, and follow the steps outlined below.

To start device parity protection, follow these steps:

**Attention:** When you work with device parity protection, the displays you see may differ slightly from the displays shown because of differences in disk subsystem requirements. The best way to ensure that you are running the correct step in each procedure is to refer to the titles of the displays.

1. From the guest partition, start `sisconfig`.
2. On the Work with Disk Units menu, select option 2 (Work with device parity protection).
3. On the Work with Device Parity Protection display, select option 2 (Start device parity protection) and press the Enter key. The Start Device Parity Protection display appears. It lists all the disk unit subsystems for which you can start device parity protection.
4. Type a 1 in the Option column next to the disk unit subsystems for which you want to start device parity protection. Press the Enter key.

**Attention:** You can select disks that are attached to different IOAs at the same time.

5. If the system detects a configuration problem, a warning report appears. If no devices are listed to start parity protection and you believe you have a valid configuration, look in `/var/log/messages` for possible error logs. You can do this manually or through `sisconfig`, using option 4 (Analyze log).
6. If you choose to continue, the Confirm Starting Device Parity Protection display appears. The display shows all the disk unit subsystems that you have selected and the individual disk units that are eligible to be started.

7. Press Enter to confirm your selection.
8. Press c to confirm the format and start parity protection operation. All data on the selected devices is lost.

**Attention:** At this point, the Enter key initiates the procedure for starting device parity protection. Once begun, this procedure continues to run until it is complete.

9. If the subsystems you have selected are correct, press the Enter key to continue.



## Display device parity status

The hardware device parity display is organized by device parity set. It includes all device parity sets on the system and all the disk units that are currently configured to use device parity protection. The possible values for the *Status* column are the following:

Status	Description
Active	This unit is part of a disk unit subsystem that has device parity protection. This unit is fully operational.
Failed	This unit is part of a disk unit subsystem that has device parity protection. This unit has failed. If another unit in this subsystem fails, data could be lost.
Synched	The parity for the parity set is being built from the data with the disk unit subsystem.
Rebuilt	This unit is part of a disk unit subsystem that has device parity protection. The data on this unit is being rebuilt from other units in the disk unit subsystem.
Unprotected	This unit is part of a disk unit subsystem that has device parity protection. This unit is operational. However, another unit in the disk unit subsystem has failed or is being rebuilt. If another unit in the disk unit subsystem fails, data could be lost.



## DASD concurrent maintenance

DASD concurrent maintenance allows disks to be replaced without powering down the server.

Concurrent maintenance support is a function of system unit hardware packaging. Not all servers support concurrent maintenance. On servers without device parity protection or software RAID, the server is not available when a disk-related hardware failure occurs. The server remains unavailable until the failed hardware is repaired or replaced. However, with device parity protection or software RAID, the failing hardware can often be repaired or replaced while the system is being used. If the server is not available when a disk-related hardware failure occurs, it may be possible to boot the Linux partition into rescue mode to perform the concurrent maintenance action, which would allow the failing hardware to be replaced without powering down the primary partition

When a disk failure occurs, you must know the exact location of the failed disk. Removing the wrong disk drive may cause results that cannot be predicted in your data files, and the next restart will take longer to complete.


The disk unit status display shows the status of individual disk unit. The possible values are:

Status	Description
Operational	The disk unit is operational and ready to accept input or output operations.
Not operational	This device cannot communicate with the IOA. Verify that the unit is powered on.
Not ready	The device cannot perform media-related functions, but it can still communicate with the IOA.

Status	Description
Read/write protected	The device cannot process either a read or a write operation. A device may be in this state because of a cache problem, a device configuration problem, or other types of problems that could cause a data integrity exposure.
DPY/Failed	This unit is part of a disk unit subsystem that has device parity protection. The disk unit failed within its device parity set, causing the loss of data protection for the device parity set.
DPY/Active	This unit is part of a disk unit subsystem that has device parity protection. The unit is operational and ready to accept input or output operations.
Performance degraded	The device is functional, but performance may be affected because of other hardware problems (such as a problem with the IOA cache.)
DPY/Unprotected	This unit is part of a disk unit subsystem that has device parity protection. Data protection is no longer in effect because of a failure in another resource.
DPY/Rebuilding	This unit is part of a disk unit subsystem that has device parity protection. Data protection is being rebuilt.
Format Required	The disk unit must be formatted to be made usable on this IOA.



**Add disk units using concurrent maintenance:** These devices are customer installable. If you elect not to do this installation yourself, contact IBM or an authorized dealer to make arrangements for them to perform the installation for a fee.

- Your unit can either be powered off (nonconcurrent) or powered on (concurrent) when you install your disk unit.
  - Be sure you understand the complexity of the powered-on (concurrent maintenance) procedures before you begin. If you do not feel comfortable using the concurrent maintenance procedure, go to [Install iSeries features](#) for detailed information on your iSeries server and how to install devices when your server is powered off.
  - Decide which way you want to install your disk unit. You can use the concurrent maintenance procedures, or you can refer to [Install iSeries features](#) for detailed information on your iSeries server and how to install devices when your server is powered off.
- For more information on a disk unit placement tool, refer to [Physical Device Placement Assistant](#).  After you have used the tool to determine the disk unit placement, return to these instructions to install the disk unit.

### Your system is powered on during installation

**Attention:** Disk units are fragile. Handle with care.

- Remove the unit front cover.
- Attach the disposable wrist strap to prevent electrostatic discharge from damaging a device.

#### Notes:

- Follow the same precautions you would use without the wrist strap. The 2209 Disposable Wrist Strap is for static control. It does not increase nor decrease your risk of receiving electric shock when using or working on electrical equipment.
- If the disposable wrist strap appears damaged or is cut off, contact your dealer for a replacement. Do not continue until a new wrist strap is available. Discard the damaged wrist strap.
- Remove the liner from the copper foil at the end, when you unroll the wrist strap.

4. Attach the copper foil to an exposed, unpainted metal surface on the frame of the system unit (electrical ground).

\_\_ 3. Find the disk unit locations in your system unit or your expansion unit. Refer to Install iSeries features for detailed information on your expansion unit.

\_\_ 4. Remove the disk unit covers. Squeeze the latch and pull out.

\_\_ 5. Look at the front of the system unit or the expansion unit and determine the next available position. When you have a system unit with an expansion unit, fill the slot positions in the base unit first.

\_\_ 6. Write down the position where the new disk will be installed \_\_\_\_\_ (for example; D03).

\_\_ 7. Write down the Frame ID where you will be installing the disk unit here \_\_\_\_\_. You will need it later. The base system Frame ID is 1. The Frame ID for each external expansion unit is displayed on the front display (for example: 02)

\_\_ 8. Be sure you are signed on the Linux partition with root authority.

\_\_ 9. Type `sisconfig` on the command line of the iSeries Linux session and press Enter.

\_\_ 10. Select Work with disk unit recovery from the Work with Disk Unit display. Press Enter.

\_\_ 11. Select Device Concurrent Maintenance from the display. Press Enter.

\_\_ 12. Type the Physical Location (Frame ID and position) where you will be installing the new disk unit. Refer to step 7 (See 14) and step 3 (See 14).

In this example, Physical Location: Frame ID 1 Position D06 , you would install a disk unit in frame 1 (base system), disk unit position 6.

\_\_ 13. Select option 2 (Install device for the action to be performed.)

\_\_ 14. Set the time delay for one minute—01.

\_\_ 15. Find the package that contains the new disk unit assembly and remove it from the static-protective package. Write the serial number (use the last four digits) of the disk unit here \_\_\_\_\_ before you install it.

\_\_ 16. Unlock the handle by pulling it out toward you and down before you install the disk unit. If the handle is not all the way down, the disk unit will not slide into the system unit or the expansion unit.

\_\_ 17. Support the disk unit assembly as you slide it half way into the system unit or the expansion unit. Do not push it any further.

\_\_ 18. Locate the light in the numbered arrow by the disk unit you are installing. When you press Enter in the next step, the light comes on and begins to blink rapidly.

\_\_ 19. Press Enter on the console.

\_\_ 20. When the light starts blinking, slide the disk unit in and lock it in place by closing the latch.

\_\_ 21. The light stops blinking and remains solid when you have installed the disk unit properly.

\_\_ 22. Return to the console and wait until the Concurrent Maintenance Results display appears.

Press Enter.

\_\_23. Do you have another disk unit to install?

\_ **No:** Continue with the next step.

\_ **Yes:** Look at the front of the system unit or the expansion unit and determine the next available position. Repeat steps 11 (See 14) through 22 (See 14).

\_\_24. When the Work with Disk Unit Recovery display appears, press q.

\_\_25. Select Display disk hardware status on the Work with Disk Units display. Press Enter.

\_\_26. The disk unit that you installed should be listed. The serial number should match the serial number you wrote in step 15 (See 14).

**Note:** If the disk unit is not listed, ensure that the disk unit was properly installed. Repeat steps Your system is powered on during installation (See 13) through 24 (See 14).

\_\_27. Press q once.

\_\_28. The installation of the disk unit is complete. You can now partition the disk and create a file system or add the device to an existing LVM volume group if you used LVM.

\_\_29. Select Work with device parity protection on the Work with Disk Unit display.

\_\_30. Select Include unit in device parity protection on the Work with Device Parity Protection display.

\_\_31. Did the Include Device Parity Protection Failed display appear?

\_\_32. The following conditions must be met before you can start device parity protection. . . . .

\_\_33. When the above conditions are met, select Enter to return to the Work with Device Parity Protection display.

\_\_34. Select Include unit device parity protection on the Work with Device Parity Protection display.

\_\_35. Select the disk units to include in device parity protection.

\_\_36. Press Enter to confirm your choice.

\_\_37. Press c to confirm format and include operations.

\_\_38. When the device parity protection is complete, the message "Selected unit(s) have been included successfully" appears.

The installation of the disk unit is complete.




**Disk hardware migration:** Normally, migration of disk units from one type of I/O adapter to another is a simple task. However, some migration paths are more complicated than others. The table below describes the various migration scenarios and the action required.

Old Adapter	New Adapter	Description
2763, 2748, 2778	2763, 2748, 2778	You can move disk units freely among these adapters with no special action.
2782, 2757	2782, 2757	You can move disk units freely among these adapters with no special action.
2763, 2748, 2778	2782, 2757	You can move disk units freely in this path. Parity-protected devices are transparently converted from RAID 5* to RAID 5** parity layout.
2782, 2757	2763, 2748, 2778	You can move unprotected disk units freely in this path. Moving parity-protected disk units requires a save and restore operation of the associated disk units.
2763, 2748, 2778, 2782, 2757	5702	Moving disk units requires a save and restore operation. After you attach the disk units to the new adapter, you must initialize and format them.
5702	2763, 2748, 2778, 2782, 2757	Moving disk units requires a save and restore operation. After you attach the disk units to the new adapter, you must initialize and format them.



## Order a new server or upgrade an existing server to run a guest partition

The LPAR Validation tool emulates an LPAR configuration and validates that the planned partitions are valid. In addition, the LPAR Validation Tool allows you to test the placement of OS/400 and Linux hardware within the system to ensure that the placement is valid. Refer to Logical Partitions  for information on the LPAR Validation Tool (LVT).

Contact your IBM marketing representative or IBM Business Partner to enter the order. You can enter the order by using the iSeries 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 guest partition, specify feature code 0142.

---

## Create a guest partition to run Linux

Before you start to configure a partition on your iSeries to run Linux, see “Plan to run Linux in a guest partition” on page 1 for assistance. You should also be familiar with basic logical partitions concepts before you start to create a guest partition. This topic guides you through the necessary steps for installing Linux on the iSeries.

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

### “Configure a guest partition” on page 17

Use iSeries Navigator to create a guest partition.

### “Assign I/O adapters (IOAs) to a guest partition” on page 17

Understand how to remove and assign adapters to a guest partition running Linux.

**“Create a network server description and a network server storage space” on page 21**

Create an NWSD to install Linux, provide virtual I/O resources to the guest partition, and understand how to control the starting and stopping of a Linux partition. Also, create an NWSSTG as a container for the Linux code.

**“Configure the LAN console for a guest partition” on page 23**

Understand how using Operations Console on the iSeries can establish a connection to the Linux console.

**“Connect the virtual console” on page 23**

Understand how to access a partition to install Linux and receive diagnostic messages using the virtual console.

**“Install Linux in a guest partition” on page 24**

Refer to your preferred Linux distributor for installation documentation.

**“Start and stop the NWSD” on page 25**

Learn how to vary on or vary off your network server description.

## Configure a guest partition

Before creating new partitions on the server, refer to [Back up your server](#) for instructions on performing a full system back up.

A guest partition on an iSeries server can support Linux, a non-OS/400 operating system. To run Linux in a guest partition, you must have V5R1 or V5R2 running in the primary partition. For more information on software release function, refer to [OS/400 logical partition function by release](#).

You can use a graphical wizard in iSeries Navigator to help you create guest partitions. iSeries Navigator comes packaged with iSeries Access for Windows and does not require an iSeries Access license to use.

In order to access logical partition functions, you must first configure the service tools server. For more information on the service tools server, refer to [Configure the service tools server](#).

To create a guest partition, follow these steps:

1. In iSeries Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the system.
3. Expand **Configuration and Service** and select **Logical Partitions**.
4. Right-click the physical system 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.

You will find detailed instructions and explanations in iSeries Navigator help topics.

## Assign I/O adapters (IOAs) to a guest partition

Before changing your hardware configuration, use the LPAR Validation Tool (LVT). This tool allows you to test the placement of OS/400 and guest partition hardware within the system to ensure that the placement is valid. The output generated from the LVT will validate the placement of hardware and ensure that your

guest partition is configured properly. Refer to [Logical Partition !\[\]\(465772ce2fc0e39b7001e2580b915cc2\_img.jpg\)](#) for information on the LPAR Validation Tool (LVT).

**Attention:**

Incorrect placement of adapters for guest partitions could result in unpredictable behavior in the guest partition.

If you want to use real hardware devices like communication lines, disk units, internal tape and CD devices in a guest partition, you must assign the controlling IOAs explicitly to the guest partition. These instructions do not apply to virtual devices like virtual Ethernet, virtual disk, virtual CD, or virtual tape provided by a hosting OS/400 partition.

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

- **Empty position:** IOA not yet installed.
- **Occupied position:** IOA installed, but not yet assigned to a guest partition or not supported by the guest partition to which it is assigned.
- **Communication IOA:** resource is a communication adapter and supported by the guest partition to which it is assigned.
- **Multiple Function IOA:** resource is a multiple function adapter and supported by the guest partition to which it is assigned.
- **Disk Controller:** resource is a disk controller and supported by the guest partition to which it is assigned.
- **Generic adapter:** IOA installed and may be supported by the guest partition to which it is assigned.
- **Storage IOA:** resource is a storage adapter and supported by the guest 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 guest partition. The logical and physical address is the only way to assure you are using the correct position in Add/Remove I/O resources during LPAR configuration.

A guest partition must have access to the bus on which its IOAs reside. The only restriction is that if the bus is shared with an OS/400 partition, the OS/400 partition must own it, shared (own shared) and the guest partition must use it shared (use shared).

### **Change bus ownership type (for shared bus ownership)**

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

1. In iSeries Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the system.
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**.

### **Move an unassigned IOA to a guest partition**

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

1. In iSeries Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the system.
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 guest partition to receive the IOA in **Move to — Logical partition**.
7. Click **OK** to move the IOA you specified.

### Move an unassigned IOA to an OS/400 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 misconfiguration of devices.

**Attention:**

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

To move an unassigned IOA to an OS/400 partition, follow these steps:

1. In iSeries Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the system.
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 OS/400 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. Select option 5 (Display Detail) next to the desired IOA and press Enter. Note the Resource name, Frame ID, and Card position.
14. 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.
15. From the Hardware Resource Concurrent Maintenance display, select option 9 (Work with controlling resource) next to the IOA selected and press Enter.
16. 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.
17. From the Confirm Add I/O resources display, verify the desired IOA has been added by verifying the logical address and pressing Enter.
18. Select **View to Refresh** the iSeries Navigator interface.
19. Right-click **Physical System** and select **Recovery** and **Clear non-reporting partitioning resources**.

### Move an IOA assigned to an IOP to a guest 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 guest 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 iSeries Navigator, expand **My Connections** or your active environment.
3. Select the primary partition of the system.
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 guest 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.

### **Assign an IOA assigned to an IOP to a guest partition**

This action should be performed when the IOP cannot be removed from the logical 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 misconfiguration of devices.

1. At the OS/400 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. 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. Select option 5 (Display Detail) next to the desired communication port and press Enter. Note the Resource name, Frame ID, and Card position.
7. 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.
8. From the Hardware Resource Concurrent Maintenance display, select option 9 (Work with controlling resource) next to the IOA selected and press Enter.
9. From the Work with Controlling Resource display, select option F7 (Relinquish from) next to the IOP displayed and press Enter.
10. Confirm the desired IOP/IOA connection is being relinquished by verifying the resource name and pressing 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.

11. At the OS/400 command line of the primary partition, type STRSST and press Enter.
12. From the Start Service Tools (STRSST) Sign On display, key a Service tools user ID and password and press Enter. LPAR Administration authority is needed to change the LPAR configuration.
13. From the System Service Tools (SST) display, select option 5 (Work with system partitions) and press Enter.
14. From the Work with System Partitions display, select option 1 (Display partition information).
15. From the Display Partition Information, select option 5 (Display system I/O resources).
16. From the Display System I/O Resources display, press F10 twice to see physical addresses (Frame ID and Card positions). Locate the desired Occupied position using the Frame ID and card position noted in step 6. 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 pressing Enter.
21. Return to the Work with Partition Configuration display, select option 3 (Add I/O resources) next to the guest 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 guest partition. If the guest 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 pressing Enter.
24. From the System Service Tools (SST) display, select option 5 (Work with system partitions) 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)

Refer to Linux for iSeries  for the latest updates to Linux running on an iSeries.

## Create a network server description and a network server storage space

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.

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

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

1. At the OS/400 command line, type CRTNWSD and press F4 for prompts.

2. From the Create Network Server Description display, provide the following information:

NWSD (Provide a name for the NWSD).  
RSRCNAME (\*NONE).  
TYPE(\*GUEST).  
ONLINE (\*NO or \*YES).  
PARTITION (Provide name of the partition).  
CODEPAGE (437).  
TCPPOPTCFG (\*NONE).  
RSTDDEVRSC (for virtual CD and tape devices) (\*NONE).  
SYNCTIME (\*TYPE).  
IPLSRC (\*NWSSTG).

**Attention:**

To boot an NWSD 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 boot from there. The NWSSTG must have a partition type "PReP Boot" (type0x41) marked as bootable.

IPLSTMF (\*NONE).

**Attention:**

This is the file to load from if you specify that the kernel boot 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).

**Attention:**

After your install, if your root file system (/) is not installed on the first partition of the first disk you will need to set a root parameter.

3. From the Work with Network Server Storage Spaces (WRKNWSSTG) select option 1 (Create Network Server Storage Space), and provide the following information:

NWSSTG (Name).  
NWSSIZE (Refer to the installation documentation of your preferred Linux distributor).  
FROMNWSSTG (\*NONE).  
FORMAT (\*OPEN).  
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.

4. Select option 10 (Add link) and press Enter, Add Server Storage Link (ADDNWSSTGL), and provide the following information:

NWSSTG (Name).  
NWSD (Name).  
DYNAMIC (\*YES).  
DRVSEQNBR (\*CALC).

**Attention:**

You cannot specify the order of disk for a guest partition. Each disk will be shown to Linux in the order they were added. This creates an issue if you need to delete a disk that is connected to a partition. For example, if there are ten disks connected to a partition and you need to replace the fifth disk, you will need to remove the last five disks, and add them again, in order, before you can replace the disk you want to change.

## Configure the LAN console for a guest partition

Operations Console is an installable component of iSeries Access for Windows. It allows you to use your PC as the system console, perform control panel functions, or both. With Operations Console, you can administer iSeries servers that are on your local area network (LAN). You can also remotely administer several iSeries servers at different locations, including other LANs. All traffic between the console and the iSeries system is encrypted to improve the security of the system.

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 iSeries system and the console.

To use the Operations Console LAN connection to provide the Linux console, use the same telnet commands as described in "Connect 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 system. 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.

If the Operations Console program is connected to more than one iSeries system, it is currently not possible to specify which system 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 iSeries system 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 iSeries system.

## Connect the virtual console

The virtual console provides the console function for the Linux system. It 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 hosting partition. TCP/IP must be configured and running on at least one OS/400 partition. The example below uses the IBM Personal Communications client.

To connect to a virtual console, 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 hosting partition, enter port number 2301 of the primary or hosting partition and press OK.
4. From the OS/400 Guest Partition Console display, select the partition to which you want to connect as console.

5. Enter the OS/400 service tools ID and password to connect to the Linux partition.

Another example of how to telnet from a dos command to the Linux console.

To connect to the virtual console, follow these steps:


1. From a MS DOS Command Prompt, telnet to your system and port 2301 (telnet xxxxxx 2301).
2. Select the partition to which you want to connect as console.
3. Enter the OS/400 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.

## Install Linux in a guest partition

Before you install Linux on a partition, you need to configure your iSeries to run Linux. Refer to “Configure a guest partition” on page 17 for more information on configuring Linux in a guest partition.

IBM is currently working with partners who will support Linux on an iSeries. Your preferred Linux distributor will provide specific documentation on how to install Linux on the iSeries. Refer to Linux for iSeries  for links to your preferred distributor.

## Install from CD-ROM images in the integrated file system

Most Linux distributors provide ISO CD-ROM images that can be downloaded from their website. 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 iSeries.

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 OS/400 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:
9. Select Mount (option 6) on the first CD image.
10. Vary on the device and press Enter. At the command line, enter the command VRYCFG and press Enter.
11. From the Vary on the device display, provide the following information:  
Virtual Optical Device Description  
Config Type of \*DEV  
Turn the status \*ON
12. At the command line, enter the command LODIMGCLG and press Enter.
13. From the Load or Unload Image Catalog display, provide the following information:  
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.) and press Enter.
14. At the command line, enter the command WRKOPTVOL and press Enter.
15. 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.

16. At the command line, enter the command WRKCFGSTS and press Enter.
17. From The Work with Configuration Status display, provide the following information:  
WRKCFGSTS \*NWS  
8 Network Server Description  
2 Update  
IPL source = \*STMF  
IPL stream file = /qopt/path to boot 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.

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

## Start and stop the NWSD

You will have to vary on or vary off the NWSD to IPL a guest partition running Linux. There are two ways to vary on or off the NWSD.

To vary on or off the network server description, follow these steps:

- From an OS/400 command line, type VRYCFG (Vary Configuration) and press Enter.
- From the Vary Configuration display, specify the name of the NWS (Configuration object), Type \*NWS, and Status \*ON or \*OFF.

or

- From an OS/400 command line, type WRKCFGSTS \*NWS (Work with Configuration Status) and press Enter.
- From the Work with Configuration Status display, select option 1 to Vary on or option 2 to Vary off the NWSD.

---

## Manage Linux in a guest partition

This topic contains information about managing a partition running Linux. You will learn how to IPL a partition running Linux and how to communicate and share information between a guest partition and an OS/400 partition. Additional information on managing a partition can be found at Manage logical partitions. Refer to your Linux distributor for Linux specific task management.

### **“Logical partition authority needed for a guest partition”**

Determine the authority needed to perform task on a guest partition.

### **“Display and change hardware configuration” on page 27**

View the console log and system reference code history for each guest partition running on your system. Learn how to determine which OS/400 partition is hosting a guest partition.

### **“Connect your guest partition to the LAN” on page 28**

Learn how to connect your guest partition to the LAN. Understand how Proxy ARP, Network Address Translation (NAT), and TCP/IP Routing will help you configure your partition to the LAN.

### **“Print system configuration for logical partitions” on page 35**

Use dedicated service tools (DST) or system service tools (SST) to print copies of your system hardware configuration.

### **“Dynamic movement of shared processor units” on page 36**

Understand how the ability to move shared processing units dynamically becomes important when you need to adjust to changing workloads.

### **“Network server description” on page 40**

Learn more about linking and unlinking storage spaces among network server descriptions. Understand what OS/400 command to use to delete disk drives in a guest partition.

### **“Application support” on page 41**

Learn more about the iSeries ODBC Driver for Linux and the Linux client running Samba on the iSeries.

### **“Back up and recover a guest partition” on page 45**

Learn how to back up the network server description and the disk drives that are associated with a guest partition.

### **“IPL a guest partition” on page 55**

Understand the advantages and disadvantages of booting from the A or B IPL source and how you can safely IPL a partition running Linux.

## Logical partition authority needed for a guest partition

The authorities that you grant to service tool users determines what logical partition information they can access and what tasks they can perform. Take care in assigning service tool user profile privileges to manage partition security.

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:

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

## Display and change hardware configuration

Use the following topics to learn how to display and change your servers hardware configuration.

- “Display the console log for the guest partition”  
View the console log for each guest partition running on your system.
- “Display the operating environment of secondary partitions”  
Find out what version of the operating system is running in secondary partitions.
- “Display the reference code history for secondary partitions” on page 28  
View system reference code history for secondary partitions.
- “Change the guest partition host information” on page 28  
Learn how to determine which OS/400 partition is hosting a guest partition.

### Display the console log for the guest partition

You can use the display console log for the guest partition to see console information for the guest partition running Linux.

You can perform this procedure from system 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 system partitions); from DST select option 11 (Work with system partitions) and press Enter.
3. Select option 1 (Display partition information).
4. Select option 10 (Display guest environment console log).

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

### Display 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 iSeries Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the system.
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.
6. Click **Help** if you need more information on this field.
7. Click **OK**.

### **Display the reference code history for secondary partitions**

A reference code indicates a status or an error condition. The system records the reference code history (the last 200 reference codes) for secondary partitions.

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

1. In iSeries Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the system.
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 system reference codes.
6. Click **Help** if you need more information on this field.
7. Click **OK**.

### **Change the guest partition host information**

A hosted guest partition is dependent on an OS/400 partition for some or all I/O resources. The hosting OS/400 partition can either be a primary or a secondary partition.

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

1. In iSeries Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the system.
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 OS/400 partition that will act as the hosting partition.
6. Click **Help** if you need more information on this field.
7. Click **OK**.

## **Connect your guest partition to the LAN**

You can connect your guest partition to the 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 hosting partition.

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

- “Directly attached LAN adapters”
- “Virtual Ethernet” on page 29

### **Directly attached LAN adapters**

A guest partition running Linux can own its own LAN adapters. Once you allocate the supported LAN adapters to the guest partition, OS/400 does not realize the hardware exist and cannot use the I/O resource. For more information on allocating an IOA to a guest partition, refer to “Assign I/O adapters (IOAs) to a guest partition” on page 17.

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

## Virtual Ethernet

A guest 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. OS/400 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:

- “Enable virtual Ethernet”
- “Create an Ethernet line description for virtual Ethernet”
- “Virtual Ethernet connectivity methods” on page 30

**Enable 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 guest partition, a network device vethXX is created in the guest 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 iSeries Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the system.
3. Expand **Configuration and Service** and select **Logical Partitions**.
4. Right-click **Properties** and select the **Virtual Ethernet** page.
5. View the virtual Ethernet options for the logical partition. Click **Help** if you need more information on this field.
6. Click **OK**.

**Create an Ethernet line description for virtual Ethernet:** Creating an Ethernet line description is the first step in configuring OS/400 to use virtual Ethernet. This configuration allows a guest partition to communicate to an OS/400 partition using virtual Ethernet.

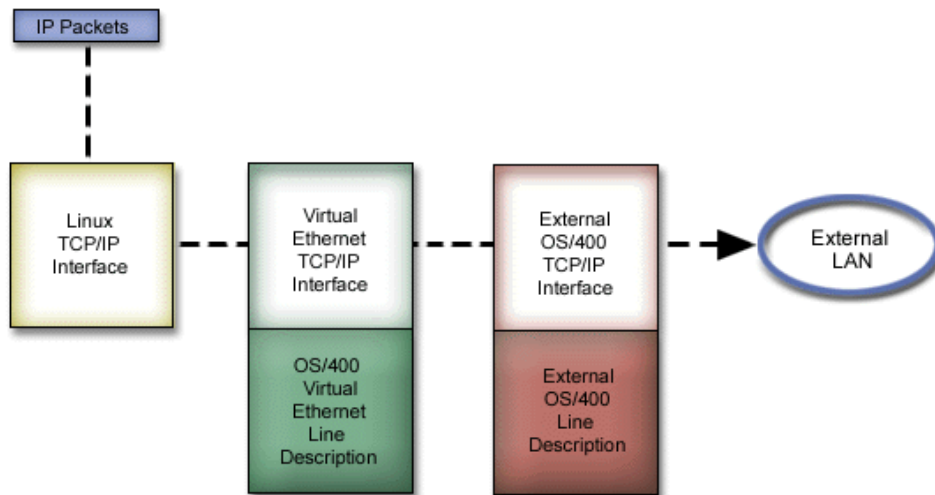
To configure a new Ethernet line description to support virtual Ethernet, complete the following steps:

1. At the OS/400 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:  
RSRCNAME  
LINESPEED (1G)  
DUPLEX (\*FULL)  
Press Enter.  
Maximum frame size (8996). By changing the frame size to 8996, the transfer of data across the virtual Ethernet is improved.  
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 iSeries system. This LAN segment is separate from any real LAN that the system might have a connection with. A virtual Ethernet is made up of a virtual line description and OS/400 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 guest partition is connected only to a virtual Ethernet segment, and you want your guest partition to communicate with systems on an external LAN, you will need to bridge TCP/IP traffic between the OS/400 external LAN and the virtual OS/400 LAN segment. A logical flow of the IP packets would look something like this:



IP traffic initiated by the guest partition goes from the Linux network interface to the virtual OS/400 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 (See 30)**

**Network Address Translation (NAT) (See 32)**

**TCP/IP Routing (See 34)**


### Proxy ARP

The Proxy ARP methods 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: More Cool Things Than Ever   
This redbook provides sample scenarios that demonstrate common solutions with example configurations. It also helps you plan, install, tailor, configure, and troubleshoot TCP/IP on your iSeries system.

- 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 guest partition as described in the following 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 guest 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 guest 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 OS/400 TCP/IP interface for your hosting partition. 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 29. Note the port of the hardware resource, because you will need it when you set up TCP/IP in your guest 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) RSRCTYPE(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.

OS/400 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 OS/400 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 guest 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 "Create an Ethernet line description for virtual Ethernet" on page 29. 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.


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 guest partition. Then from OS/400 ping the virtual Ethernet interface and the Linux interface.

### Network Address Translation (NAT)

NAT can route traffic between your guest 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 guest partition. Other forms of NAT like masquerade NAT would also work if your guest 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 OS/400 network connection. Since you will be using IP Packet rules, you must use iSeries Navigator to create and apply your rules.

The following illustration is an example of using NAT to connect your guest 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 example, any existing TCP/IP traffic for the iSeries 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 guest partition uses the virtual interface (192.168.1.1) and its own 192.168.1.5 interface to communicate with the hosting partition.

To make static NAT work, you need to first set up your OS/400 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 guest partition, make sure you selected to create virtual Ethernet. This is described in “Configure a guest partition” on page 17.
2. “Create an Ethernet line description for virtual Ethernet” on page 29 for your virtual Ethernet. Note the port of the hardware resource, because you will need it when you set up TCP/IP in your guest 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) RSRCTYPE(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 guest 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 “Create an Ethernet line description for virtual Ethernet” on page 29. 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.

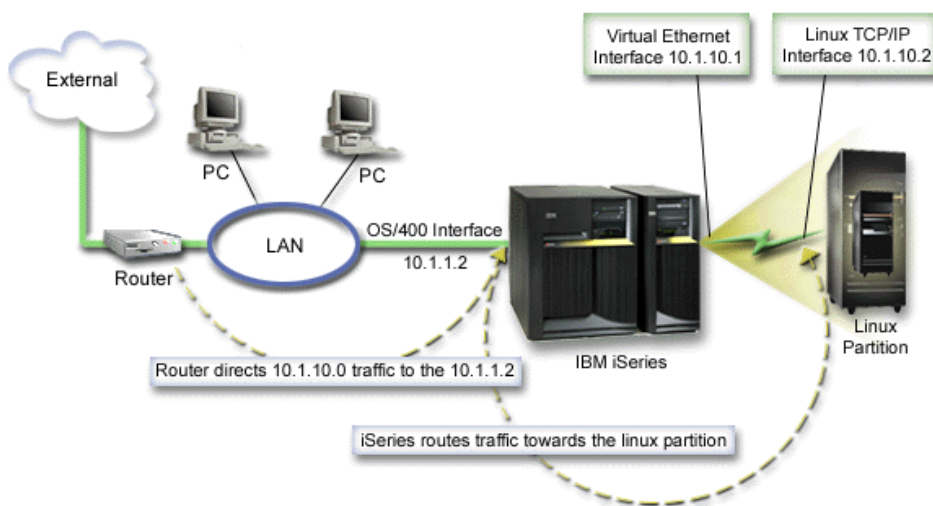
In the example:

- the interface IP address is 192.168.1.5.
  - the subnet mask is 255.255.255.0.
  - the gateway IP address is 10.1.1.3.
  - 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 guest partition, ping the Linux gateway and from OS/400, ping the guest partition. If the ping is successful, continue.
  8. Turn on IP datagram forwarding. This allows the OS/400 TCP/IP interfaces to pass IP packets between each other. *CHGTCPA IPDTGFWD(\*YES)*
  9. Connect to the hosting partition with iSeries Navigator. You must not connect to the hosting partition 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 mouse click on **Defined Addresses** and select **New Defined Address**.
    - c. Enter an Address name, the IP address of the guest 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 mouse click on **Defined Addresses** and select **New Defined Address**.
  - e. Enter an Address name, the IP address of the guest 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 mouse 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 Journalling set to off. For the example, you would enter:
    - Mapped address name = LINUXPART
    - To Address name= SHELL
    - Line = ETHLINE
    - Journalling = 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.
11. To test outbound communications, ping an external host from your guest partition. Then from that external host, ping your guest partition to test inbound communications.

## TCP/IP Routing

You can also route traffic to your guest partitions through your iSeries server with various routing techniques. This solution is not difficult to configure on your iSeries but depending on the topology of your network, it may not be practical to implement. Consider the following 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 OS/400, if you turn on IP Datagram Forwarding, the iSeries will route the IP packets to and from the guest 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 iSeries. 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 iSeries) 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 iSeries 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 guest partition, make sure you selected to create virtual Ethernet. This is described in “Configure a guest partition” on page 17.
2. “Create an Ethernet line description for virtual Ethernet” on page 29 for your virtual Ethernet. Note the port of the hardware resource, because you will need it when you set up TCP/IP in your guest 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) RSRCTYPE(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 guest 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 “Create an Ethernet line description for virtual Ethernet” on page 29. 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.  
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 OS/400 TCP/IP interfaces to pass IP packets between each other. *CHGTCPA IPDTGFWD(\*YES)*

## Print system configuration for logical partitions

If you are running OS/400 V5R2 on the primary and all of the other partitions on the system are guest partitions, it is strongly recommended that you print the system 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 system in a disaster recovery scenario.

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

**Attention:**

Printing a rack configuration listing through Hardware Service Manager within SST will only provide 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 system. For this reason, you should use the steps outlined below using the primary partition.

Follow these steps to print the system configuration:

1. From the primary partition start SST or DST.
2. From SST select option 5 (Work with system partitions); from DST select option 11 (Work with system partitions) and press Enter.
3. Select option 1 (Display partition information).
4. Select option 5 (Display system 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 system 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 OS/400 command line and print these three spooled files.

## Dynamic movement of shared processor units

You can dynamically move shared processors units among guest 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 guest partition. For shared processors, you can specify a minimum value equal to the minimum amount of processing power needed to support the guest partition. The maximum value must be less than the amount of processing power available on the system. 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 guest partition to another, follow these steps:

1. In iSeries Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the system.
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 guest 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

Virtual devices simplify hardware configuration on the iSeries system because they do not require you to add additional physical devices to the system in order to run Linux. You can allocate up to 48 virtual disks to a Linux partition, depending on which distribution you use. With each disk supporting up to 64 GB of storage, it is possible to have up to 1.2 TB of disk available to each Linux partition. Each virtual disk appears to Linux as a single physical unit. However, the associated space in the OS/400 integrated file system (IFS) is spread across all of the disk arms available on the iSeries server. This provides the benefits of RAID without the overhead of RAID support in the Linux kernel.

Version 5 Release 2 provides the ability to dynamically add virtual disks to a Linux partition. You can allocate disk space in the IFS and make it available to Linux without restarting OS/400 or Linux. The Linux system administrator can also configure the newly allocated disk space and make it available, again without interrupting the system.

1. Use the Create Network Server Storage Space (CRTNWSSTG) command to add disks to a Linux partition. Alternatively, you could use the Work with Network Server Storage Space (WRKNWSSTG) command and select option 1=Create. Both of these alternatives create the virtual disk, allocate it to the partition, and bring up the Create NWS Storage Space (CRTNWSSTG) display. On that display, enter the following parameter values:
  - In the Network server storage space field, enter the name you want to give to the storage space. This name has no direct correlation in Linux.
  - In the Size field, enter the size in megabytes for the new storage space.
  - In the Format field, enter \*OPEN.
  - In the Text description field, enter a meaningful description for the storage space.
2. Enter the Display Network Server Storage Space (DSPNWSSTG) command to verify that the new storage space exists in the list of storage spaces. Alternatively, you could use the Work with Network Server Storage Spaces command and select option 5=Display.
3. Enter the Add Network Server Storage Link (ADDNWSSTGL) command to allocate the new virtual disk to the appropriate network server descriptor. Alternatively, you could use the Work with Network Server Storage Space command and select option 10=Add link for the virtual disk. Both of these alternatives bring up the Add Network Server Storage Link display.



On that display, enter the following parameter values:

- To allocate the virtual disk, specify the name of the network storage description (NWSD) in the Network server description field.
  - To make the disk dynamically available to the partition (that is, available without rebooting the Linux partition), specify \*YES in the Dynamic storage link field.
4. That's it! As far as OS/400 is concerned, the space is available to the Linux partition. Now we need to run through a couple of steps within Linux to partition, format, and access the disk. We need to determine the device name of the new disk. The name depends on two factors:
    - The distribution of Linux
    - The number of disks currently allocated

Let's assume that we are running the SuSE distribution and that only a single disk is allocated. The first disk device is known as /dev/hda. The second disk (the newly allocated virtual disk) is known as /dev/hdb.

5. To create partitions on the new disk, we can use the fdisk command (a standard Linux command, the same for all distributions). Continuing with our example, we would enter the command (in Linux) as follows:

```
fdisk /dev/hdb
```

Keep in mind that you must have superuser (that is, root) privileges to execute this command. The command responds with the following prompt:

```
Command (m for help):
```

6. Enter 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:

```
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

7. We don't want a FAT16 partition, so we need to delete the current partition and then create a new one. To delete the partition, enter `d` at the command prompt. `fdisk` responds with the following prompt:  
Partition number (1-4):

8. Enter the partition number (in this case, 1), and press the Enter key. `fdisk` returns the command prompt, indicating that the deletion was successful.

9. Now we need to create the new partition. Enter the command `n` for a new partition. The following prompt appears:

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

We want to create a primary partition, so enter `p` and press Enter. The following prompt appears:

Partition number (1-4):

10. Since this is the first partition on the disk, enter 1 and press Enter. The following prompt appears:

First cylinder (1-200, default 1):

11. Since we will use the entire disk for this partition, press Enter to take the default of 1 for the first disk cylinder. The following prompt appears:

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

12. Again, since we will use the entire disk for this partition, press Enter to take the default of 200 for the last disk cylinder. The Command 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.

13. Up until this point, we have made no changes to the disk structure. To commit the changes, enter the command `w` (for write) and press the Enter key. `fdisk` responds with the following diagnostic messages:

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

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

14. Once the operation is complete, `fdisk` exits back to a command prompt. Now that the disk is partitioned, we need to format the partition. To do this, we can use the `mkfs` command (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. Therefore, 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.

15. Finally, within the Linux file tree, we need to create a directory that we can use to access the new file system. We 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:

```
mkdir /mnt/data
mount /dev/hdb1 /mnt/data
```

16. You now have new virtual disk space available to the running Linux partition. To facilitate automatic mounting of the resource across Linux boots, 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.

- “Link NWSSTG among multiple NWSDs”  
Learn more about linking network server storage among network server descriptions.
- “Delete NWSDs for a guest partition”  
Follow detailed steps to unlink storage spaces before you delete NWSD for a guest partition.
- “Delete disk drives for a guest partition” on page 41  
Understand what OS/400 command to use to delete disk drives in a guest partition.

### Link NWSSTG among multiple NWSDs

You have the ability to link network server storage space to multiple network server descriptions. 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 NWSDs, follow these steps:

1. At the OS/400 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.

### Delete NWSDs for a guest partition

Before you delete a network server description (NWSD), you need to “Unlink disk drives from a guest partition” on page 54 storage spaces that are associated with that NWSD. Then you can delete the NWSD.

To unlink the storage space for the system drive for NWSDs do the following:

1. On the OS/400 command line, type `RMVNWSSTGL NWSSTG(nwsdname1) NWSD(nwsdname)`. Press Enter.

To unlink the storage space for the install source drive do the following

1. Type `RMVNWSSTGL 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 command as often as needed `RMVNWSSTGL NWSSTG(nwsstgname) NWSD(nwsdname)` and press Enter.

To delete the network server description (NWSD) for a guest partition, follow these steps:

1. On OS/400, type the command `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, go to the next step.
4. Press F3 to return to the previous screen.
5. Enter a 4 in the Opt field to the left of the Network Server and press Enter.
6. On the Confirm Delete of Network Server Descriptions display, press Enter.

### **Delete disk drives for a guest partition**

Before you can delete a disk drive, you must “Unlink disk drives from a guest partition” on page 54 it from the NWSD. Once you have unlinked it, you can delete it.

You can use the OS/400 command line interface to delete the disk drive.

To delete the disk drive by using the OS/400 command line, follow these steps to use the Delete Network Server Storage Space (`DLTNWSSTG`) command:

1. Enter the command `DLTNWSSTG` and press F4.
2. In the Network server storage space field, enter the name of the disk drive.
3. Press Enter.

## **Application support**

Use the following topics to learn how to effectively use the Linux ODBC driver and Samba in a guest partition.

- “Access iSeries data using the Linux ODBC driver”  
Learn how the iSeries ODBC driver for Linux allows you to access iSeries databases.
- “OS/400 NetServer support for Samba”  
Learn more about a Linux client running Samba to connect to an iSeries NetServer.


### **Access iSeries data using the Linux ODBC driver**

The iSeries ODBC Driver for Linux allows you to access the iSeries database data from Linux applications written to the ODBC API. It is based on the ODBC driver in the iSeries Access Express for Windows product.

Refer to Linux for iSeries  for more information on using the Linux ODBC driver.

### **OS/400 NetServer support for Samba**

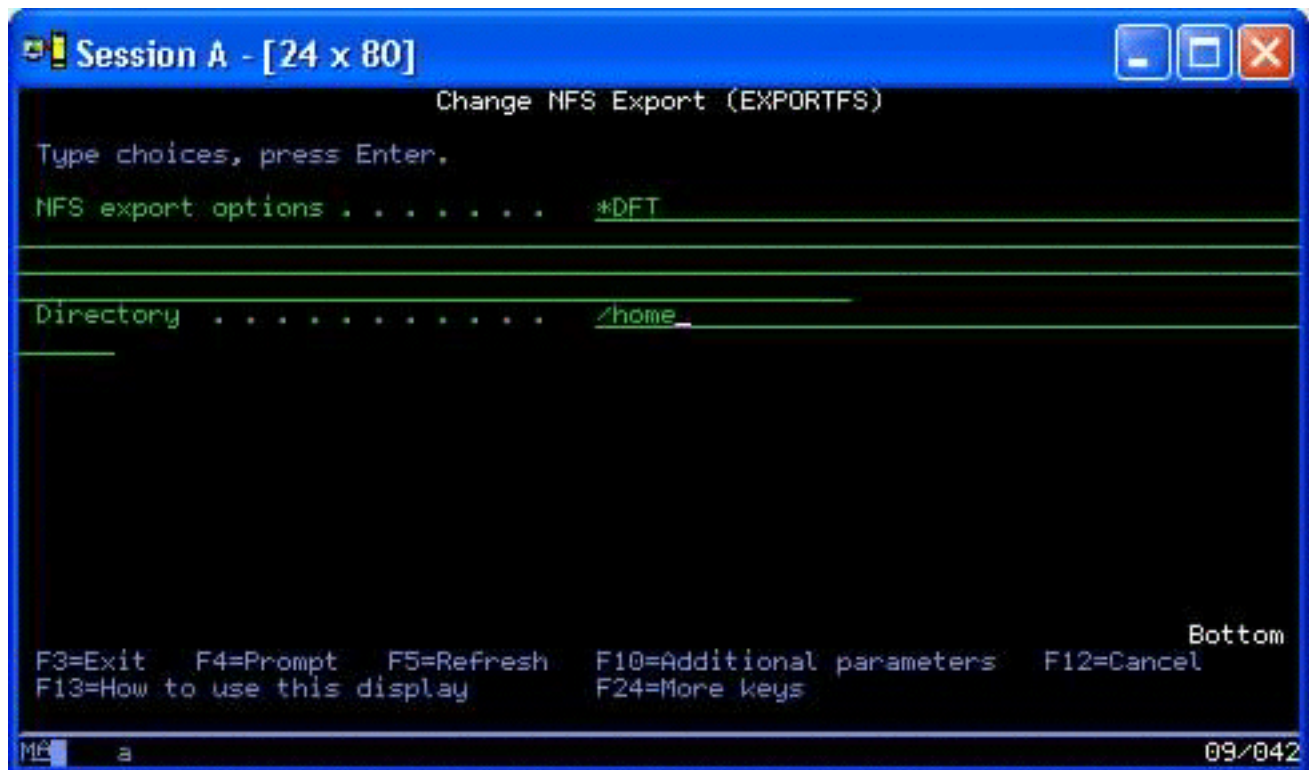
iSeries NetServer provides support for Linux. This support allows a Linux client running Samba to connect to an iSeries NetServer.

Samba is an Open Source client and file server compatible with Microsoft Networking that comes with many current distributions of Linux. For more information on OS/400 NetServer support for Samba, refer to Linux for iSeries  . >>

## 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 (IFS). Two of these tools are Network File System (NFS) and SAMBA. (SAMBA is Linux's implementation of Windows File Serving.)

1. To access IFS resources via NFS, the resource must first be available from OS/400. To make the resource available from OS/400, export the desired directory (that is, publish the directory so that others with NFS and the appropriate access rights can access the resource). You can export the directory either by editing the `/etc/exports` file within OS/400 or by using the Change NFS Exports (EXPORTFS) command. If you use the EXPORTFS command, the Change NFS Export (EXPORTSF) display appears.



In this case, the OS/400 is exporting the `/home` directory with default NFS options. If the NFS services are running within OS/400, the resource is now available for access from the Linux partition. Assuming that the NFS services are running within OS/400, the resource is now available for access from the Linux partition.

2. For the Linux partition to access the NFS resource, we must complete a number of steps:
  - a. Create an access point within the Linux directory structure that Linux will use to access the files. To create this access point, create a directory:

```
mkdir /mnt/ifsdata
```
  - b. Associate the network resource (in this case, NFS) with the newly created directory. You can do so either by using the `mount` command (for temporary access to the resource) or by adding an entry to the `/etc/fstab` file (for pervasive file access).
    - The `mount` command to access the resource would resemble the following:

```
mount -t nfs rchassam:/home /mnt/ifsdata
```



- The /etc/fstab file stores mount definitions that the mount command can use both during system initialization and once the system is running. An entry to access the NFS resource would resemble the following:

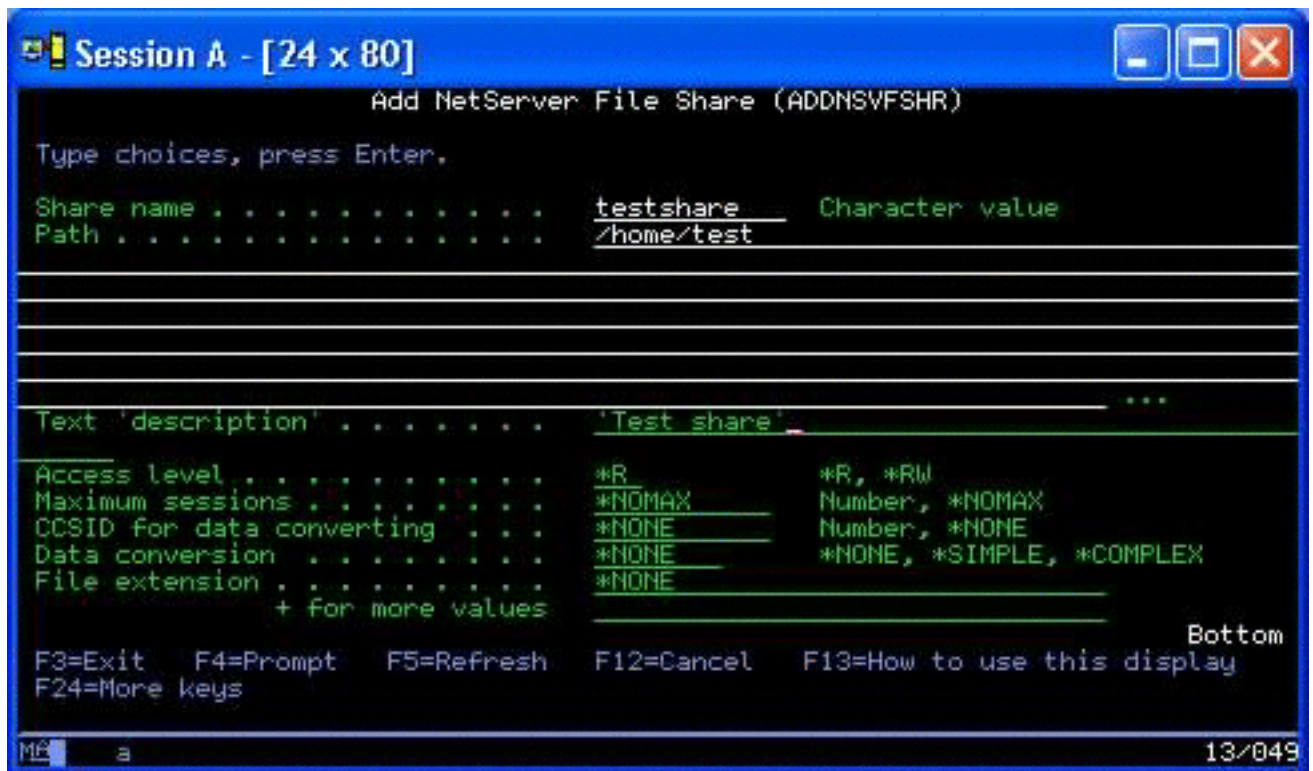
```
rchassam:/home /mnt/ifsdata nfs noauto 0 0
```

**Note:** The purpose of the noauto command in the entry above is so that Linux does not attempt to mount the NFS resource at IPL time. This is a good option to use if you do not know what the state of the NFS server will be at IPL.

3. In addition to NFS, Linux partitions can also use SAMBA to access file resources stored in the integrated file system (IFS).
  - a. File resources in the Integrated File Server that SAMBA will access must first be available (shared) via AS/400 NetServer. You can accomplish this with the NETS menu, which you access by entering the GO NETS command in OS/400.



- b. On the NETS Menu, select option 3, Add File Share. The Add NetServer File Share (ADDNSVFSHR) display appears.



- c. Enter the appropriate values in the following fields on that display:
- In the Share name field, enter the name that SAMBA will use in Linux to access the resource.
  - In the Path field, enter the location in IFS of the resource that will be shared.

At this point, assuming that NetServer is running on the OS/400 system, the resource is available for access from Linux via SAMBA.

4. For the Linux partition to access the SAMBA resource, you must complete several steps:
- Create an access point within the Linux directory structure that Linux will use to access the files. To accomplish this, create a directory:
 

```
mkdir /mnt/smbdata
```
  - Associate the network resource (in this case, SAMBA) with the newly created directory. You can do this either by using the mount command (for temporary access) or by adding an entry to the `/etc/fstab` file (for pervasive file access).
    - The mount command to access the resource would resemble the following:
 

```
mount -t smbfs -o username=tempuser //qrchassam/testshare /mnt/smbdata
```

The username option specifies the username that Linux should send to the SAMBA server for access. A prompt asks the user to enter his password. If the username and password match a known username-password combination within OS/400, Linux grants access and the files are available at `/mnt/smbdata`.

- Alternatively, you can add an entry to the `/etc/fstab` file. The `/etc/fstab` file stores mount definitions that the mount command can use both during system initialization and once the system is running. An entry to access the SAMBA resource would resemble the following:
 

```
//qrchassam/erwin /mnt/smbdata smbfs noauto,username=tempuser 0 0
```

**Note:** The purpose of the `noauto` command in the entry above is to prevent Linux from attempting to mount the NFS resource at IPL time. This is a good option to use if you do not know the state of the SAMBA server at IPL. Also, as with the mount command shown above, the username parameter indicates the username that Linux sends to the SAMBA

server. When the user tries to mount the resource, Linux prompts the user for his password. If the username-password combination matches that of a valid user on OS/400, the system grants access to the resource.





## Back up and recover a guest partition

Because iSeries integration of a guest partition combines two operating systems (OS/400 and the guest operating system), you may be able to use OS/400 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.

To back up a guest partition comprised of one or more virtual disk on iSeries, refer to the following articles:

- “Save guest partition configuration data” on page 51  
Understand the importance of saving configuration data for your server.
- “Back up the NWSD and disk drives associated with a guest partition” on page 52  
Understand how to save the network server description and storage spaces for guest partition.
- “Build a rescue image on an NWSSTG” on page 52  
Understand how to build a rescue image on a network storage space.
- “Use a recovery system from an NWSSTG” on page 53  
Follow the steps to use the rescue image created on the network storage space.
- “Back up NWSDs for a guest partition” on page 54  
Understand how to save the network server description and storage spaces for a guest partition.
- “Restore NWSDs for a guest partition” on page 54  
Understand when you would restore all the configuration objects for a guest partition.
- “Unlink disk drives from a guest partition” on page 54  
Learn how to unlink disk drives if you are uninstalling a guest partition.
- “Objects to save and their location on OS/400” on page 55  
Understand what objects to save and how to save objects on an OS/400 partition.

If you need to recover your entire system, refer to Backup and Recovery  . 

## Backup options for virtual and directly attached disk

There are two types of options for backing up and recovering data involving Linux partitions on the iSeries server:

- Utilities that run within Linux
- Commands that run within OS/400

**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 “Save and restore hosted Linux files to and from the host’s shared tape drive” on page 49.
- You can write the file to a file in the partition’s file system. Following is an example of this alternative.

```
tar -cvf /tmp/etc.tar /etc
```

where

c = create a tar file  
v = verbose (show the files that are being added to the tar file)  
f = name of the tar file follows  
/tmp/etc.tar = name of the tar file  
/etc = 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 system 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.

**OS/400 commands:** Users of virtual disks also have powerful tools to use in OS/400 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 boot 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 boot the Linux partition after a total iSeries reinstallation. You can also transport and restore saved virtual disks to other iSeries servers via FTP and tape.

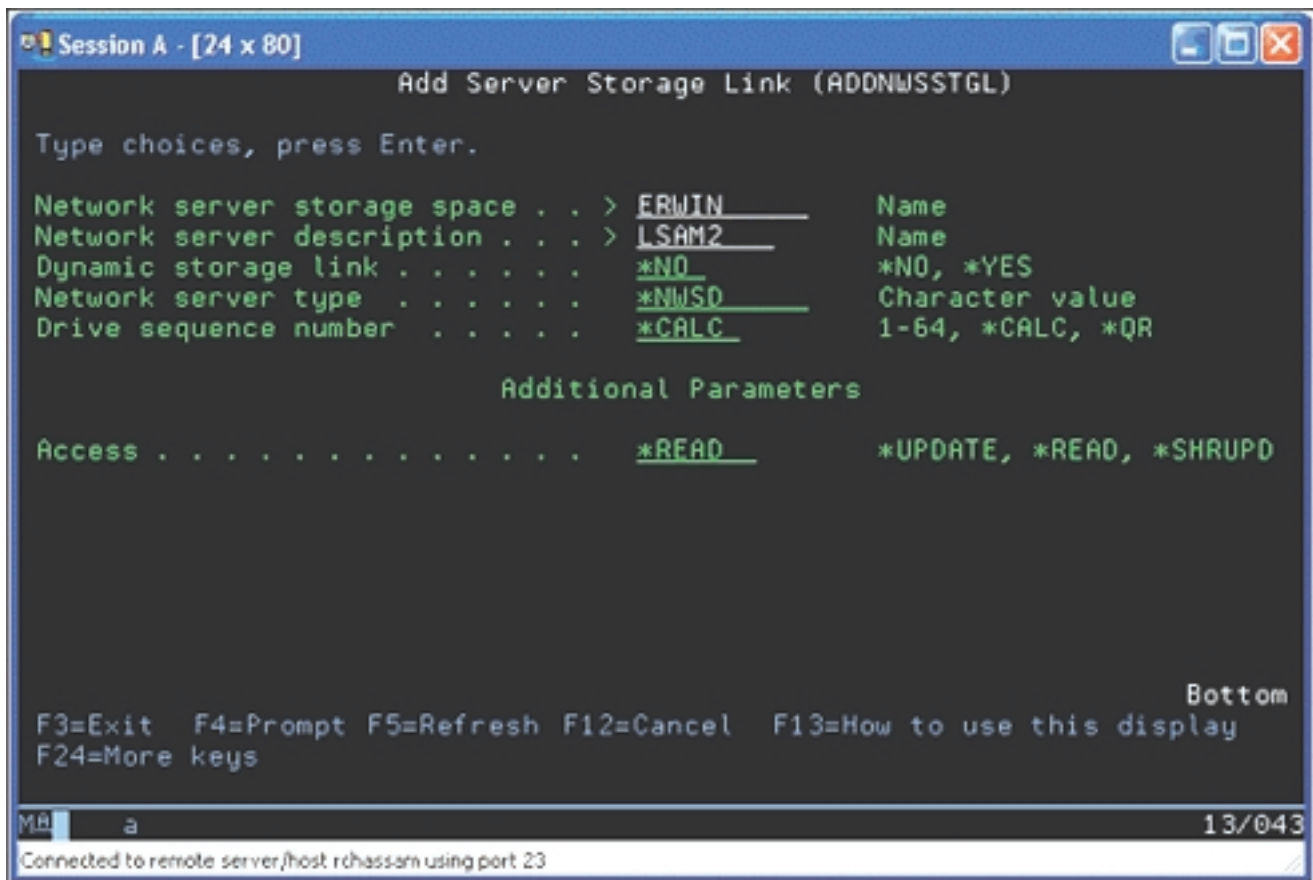


## How to use the shared read-only storage space

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.

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



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.

You can specify that Linux has read-only access to the disk in either of 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
```

To update a shared storage space, 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 of OS/400 is required, but it does not have to be in the primary partition.
- The maximum size supported for a storage space is 64 GB.
- The maximum number of storage spaces supported is 20, depending on the Linux distribution.



### Sharing access to disks in read-write mode

Although simultaneous access to a virtual disk by multiple logical partitions is limited to “How to use the shared read-only storage space” on page 46, 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.

To enable multiple partitions to access a virtual disk in read-write mode, you must set the access type to shared update (\*SHRUPD). You specify this access type on the Access parameter of the Add Network Server Storage Link (ADDNWSSTGL) command. This command links the network storage space to the network storage descriptor.



On the Add Network Server Storage Link (ADDNWSSTGL) display, you can specify either \*YES or \*NO in the Dynamic storage link field.

To display the Access field, you must press the F9 key.

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.



## How to clone Linux servers in iSeries partitions quickly

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 iSeries system and restore it to the same or another iSeries system 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 (CRTNWSD) 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, description, and access mode for the new virtual disk and press the Enter key. At this point, the system 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 VRYLPAR command.

A clever way to save additional disk storage is to split a virtual disk into two disks: a larger one for “How to use the shared read-only storage space” on page 46 and a smaller one for “Sharing access to disks in read-write mode” on page 48. That way, you only have to copy the read/write disk. You can share the read-only disk.



## Save and restore hosted Linux files to and from the host's shared tape drive

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 OS/400 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/iseri/vt0`, `vt1`, etc. TurboLinux has its own naming convention.
- To use the tape device from Linux, you must vary the tape off under OS/400. To use the OS/400 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 OS/400 data and tar data on the same tape volume.

**Save and restore 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 filesystem. For example, you cannot save a kernel residing in the PReP boot partition by using the tar command.

*Save 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.

*Restore 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.

**Save and restore hosted Linux data by using OS/400 SAV and RST:** Backups of hosted Linux data from OS/400 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 boot partition.

*Save Linux hosted data by using OS/400 SAV:* On OS/400, your data is in a network server storage space. The OS/400 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 OS/400 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.

*Restore Linux hosted data by using OS/400 RST:* The OS/400 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 OS/400 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 OS/400 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 OS/400.

**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.



**Save guest partition configuration data**

Logical partition configuration data is automatically maintained for the life of the physical system. Each OS/400 partition load source contains a copy of the configuration data.

A guest partition boot devices do not contain partition configuration data. Systems comprised solely of a single primary OS/400 partition and one or more guest partitions should have mirrored OS/400 load source devices in the primary partition. For more information on how to save logical partition data for one or more guest partitions, refer to Save all configuration data.

Only disaster recovery to a different physical system would require that you rebuild the configuration from the beginning. You should “Print system configuration for logical partitions” on page 35 when you make changes to your logical partition configuration. This printout will help you as you rebuild the configuration.

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.

**Attention:**

Logical partitions that you keep powered 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.

## **Back up the NWS and disk drives associated with a guest partition**

When you install the guest partitions with virtual disk on iSeries, the hosting OS/400 partition creates a network server description and disk drives for your guest operating system that you need to back up. Some of the disk drives are system-related (the installation and system drives); others are user-related. Because your guest operating system may consider them a unified system, you need to save all the disk drives and the network server description to restore properly.

The implementation of a guest partition for iSeries allows you to save and restore virtual disk as OS/400 network server storage space objects. These objects are saved as part of the OS/400 system when you perform a full OS/400 system backup. You can also specifically save the network server description and storage spaces that are associated with a guest partition on iSeries. Daily backup of the system drive is a good idea.

To back up the network server description and the disk drives that are associated with a guest partition on iSeries, see "Back up NWSs for a guest partition" on page 54.

## **Build a rescue image on an NWSSTG**

One rescue solution for a guest partition is to create a small network storage space (NWSSTG) that can be left on the IFS solely for the purpose of rescuing guest partitions. First, a quick explanation of what a Linux Rescue system normally entails.

Many distributors include a rescue image on their installation disks that will boot 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 guest partition. You can simplify this process by creating a network storage space with a rescue system in it at the time you install your guest partition.

Before creating a rescue image on network storage it is important to document the configuration information for each of your guest 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 system 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 boot 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 (NWS). You don't need to create a new NWS for this step. You could unlink an existing storage space and temporarily link your rescue storage space to any of your existing NWS's. Next, start your distributor's installation system as described in the documentation and follow the prompts. Make sure that you create a PrEP boot 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 boot up the system for you. At this time you should take some time to verify that the rescue system has all the utilities that you are going to need. For a guest partition, you are going to want to do an `"rpm -qa | grep ibmsis"` to make sure that the utilities to work with native DASD 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 guest partitions may require and ftp the files with the configuration information for your other guest partitions onto the rescue system 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 booting 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.

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

### **Use a recovery system from an NWSSTG**

To use the rescue image 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 boot 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 DASD 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 system) or `/dev/sda2` in the case of a partition on native DASD. It is important to note that the drive you are trying to rescue, when using virtual DASD, 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 system.) 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.

## Back up NWSDs for a guest partition

When you save the storage space objects that are associated with a guest partition virtual disks on iSeries, you also need to save the network server description (NWSD). Otherwise, a guest partition may not be able to re-establish items such as the partitions file system permissions. To save the NWSD, you use the Save Configuration (SAVCFG) command:

1. On the OS/400 command line, type SAVCFG.
2. Press Enter to have OS/400 save the NWSD configuration.

### Attention:

The save configuration command (SAVCFG) will save 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 will need to be added manually once the configuration and the dynamically linked storage spaces have been restored.

## Restore NWSDs for a guest partition

In a disaster recovery situation, you would restore all the configuration objects, which include the network server description (NWSD) for your guest partition. In some situations, for example when you migrate to new hardware, you need to specifically restore the NWSD. To have OS/400 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 OS/400 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 OS/400 restore the NWSD.
5. When you are done restoring the NWSD and all its associated storage spaces, “Start and stop the NWSD” on page 25 the guest partition.

## Unlink disk drives from a guest partition

Unlinking Linux virtual disk drives (network server storage spaces) disconnects them from the guest partition, making them inaccessible to users. You also need to unlink disk drives if you are uninstalling a guest partition.

You can use the OS/400 command line interface to unlink a disk drive.

To unlink a disk drive by using the OS/400 command line, follow these steps:

1. “Start and stop the NWSD” on page 25 for your guest partition.
2. On the OS/400 command line, enter the command RMVNWSSTGL and press F4.
3. In the Network server storage space field, enter the name of the storage space you want to unlink.
4. In the Network server description field, enter the name of the server from which you want to unlink the storage space.
5. If you are unlinking a statically linked disk drive that you plan to relink, specify \*NO 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.
6. Press Enter.
7. If you are uninstalling a guest partition, your next step is to “Delete disk drives for a guest partition” on page 41. Otherwise, “Start and stop the NWSD” on page 25 for your guest partition.

## Objects to save and their location on OS/400

Many objects are created as a result of installing a guest operating system in a partition and using virtual storage. Some of these objects are system-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 OS/400 GO SAVE command in the hosting partition. Option 21 saves the entire system. Option 22 saves system 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 OS/400 and the command to use. Manually saving parts of your system. has more information about using the save commands.

### Objects to save

Object content	Object name	Object location	Object type	Save command
<b>For guest partitions with virtual disk</b>				
Guest partition and virtual disk drive	stgspc	/QFPNWSSTG	User-defined network server storage spaces in system	GO SAVE, option 21 or 23 SAV OBJ('/QFPNWSSTG/stgspc') DEV('/QSYS.LIB/TAP01.DEVD')
			User-defined network server storage spaces in user ASP	SAV OBJ('/QFPNWSSTG/stgspc') ('/dev/QASPnn/stgspc.UDFS') DEV('/QSYS.LIB/TAP01.DEVD')
<b>For all guest partitions with a hosting partition</b>				
Messages from the guest partition	Various	Various	Server message queue	GO SAVE, option 21 or 23 SAVOBJ OBJ(msg) LIB(qlibrary) DEV(TAP01) OBJTYPE(*MSGQ)
				OS/400 config objects for guest partitions
Various	Various	QU SRSYS	Various	GO SAVE, option 21 or 23 SAVLIB LIB(*NONSYS) or LIB(*ALLUSR)

## IPL a guest partition

Use the following topics to learn how to IPL a guest partition.

- “Decide what IPL type to use when running Linux”  
Learn how you can safely IPL a partition running Linux.
- “Boot from the A and B IPL source” on page 56  
Understand the advantages and disadvantages of booting from the A or B IPL source and who has authority to perform the task.

## Decide what IPL type to use when running Linux

The IPL type parameter on the NWSD determines the source from which a guest operating system is loaded. The first thing to understand is that on an iSeries system 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 OS/400 using a network server description (NWSD.)

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

IPLSRC values	Description
A	Load the guest 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 guest 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 <b>reserves</b> this type of IPL for future use.
*Panel	The partition is started from the source indicated on the operator's panel.
*NWSSTG (network server storage space)	This IPL type is used to boot a partition from a virtual disk. OS/400 will find the kernel in the virtual disk and load it in location C. Then the partition will be set to IPL from C. OS/400 searches the first virtual disk connected to the NWSD for a partition marked bootable, and of type 0x41 (PReP boot.) 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 boot a partition from a kernel loaded in the OS/400 Integrated File System. Note that the Integrated File System includes files on the Optical (CD) drive on OS/400. OS/400 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 guest 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 system uses during the initial program load (IPL). These parameter can have the following values:

IPLSRC values	Description
A	Load the guest 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 guest 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 <b>*NWSSTG</b> for more information on how the C IPL type is used.
D	IBM support <b>reserves</b> this type of IPL for future use.

### Boot from the A and B IPL source

A guest partition running Linux can boot from either the A or B IPL source. However, Linux must be installed on the system 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.


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 boot source is that Linux boots faster. A disadvantage is that the boot 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 boot, you can modify the kernel command line parameters either through the Work with partitions service tool on your OS/400 primary partition or through the /proc file system under Linux.

Follow these steps to set the kernel command line:

1. At the OS/400 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.
3. From the System Service Tools (SST) display, select option 5 (Work with system partitions) 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 guest commands) next to the guest 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

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 boot from a kernel in the integrated file system (\*STMF). After you boot the partition, you can use Linux commands to rewrite the Linux kernel and enable the partition to be bootable.

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

1. Put a kernel in your IFS.
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. Enter 8=Display description in the Option column next to the Linux partition you want to boot. A description of that network server appears.
4. Enter 2=Edit in the Option column next to IPL source, IPL stream file, and IPL 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/kernels/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 (that is, enter telnet <your machine> 2301).
8. Vary on the partition.
9. Log in to the partition.
10. If your system does not have a copy of the kernel that you want to install, FTP it on to your system. You could even NFS share it off of the OS/400.
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 IPL 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 <many more spaces> > /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 IPL source = B and the IPL stream file = \*NONE.
18. Now you can vary on your partition and use the kernel installed in the iSeries slot B.



---

## Troubleshoot Linux running in a guest partition

Use this topic to understand how to analyze and resolve errors specific to Linux running on the iSeries. If your problem is specific to logical partitions, refer to Troubleshoot logical partitions for additional assistance. Linux specific troubleshooting problems require the assistance of your Linux distributor.

### “Debug NWSD error messages”

Find a list of error codes and messages related to problems with the network server description.

### “Debug the processor multitasking error” on page 59

Understand how to determine whether your server should have the processor multitasking function disabled.

### “System reference codes (SRCs) for Linux” on page 60

Find a list of system reference codes specific to Linux and suggested corrective actions for resolving errors.

### “Recover from a virtual Ethernet network failure” on page 61

Understand how to recover virtual Ethernet if you have upgraded your kernel or changed a network device.

## Debug NWSD error messages

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 guest partition running on the system. All error messages related to the NWSD should appear in QSYSOPR indicating a description of the problem and a resolution to the issue.

### NWSD error messages



Reason codes	Code explanations
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 NWS in the OS/400 partition that is active and using the partition specified in the PARTITION parameter of the NWS.
00000005	The partition specified in the PARTITION parameter of the NWS is powered on (perhaps through the LPAR configuration interface or from another OS/400 partition.)
00000006	The partition is set to boot 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 NWS is set to boot 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 bootable.
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.
00000009	The partition identified as the hosted partition is not configured. To select a host partition, use the Work with Partition Status display and press F11 (Work with partition configuration). From the Partition Configuration display, select option 13 next to the partition to assign the partition a host.
00000010 00000011 00000080	Contact your next level of support to find a proper solution to the problem.
00001088 00001089 0000108A	The kernel does not appear to be valid. This error is frequently caused if you do not FTP the kernel in binary mode.
0000108B 0000108C	The kernel does not appear to be compatible with the version of OS/400 in the primary partition.
000010A3 000010A9 000010AA	There is an insufficient amount of processors assigned to the partition or there are not enough shared processors available.
000010A4 000010A5	There is an insufficient amount of memory available for the partition.
000010AE	This error will occur on systems that only support dedicated processors when either you have specified a shared processor for a Linux partition or you have the QPRCMLTTSK system value set to 1.

## Debug the processor multitasking error

In order to run Linux on certain iSeries systems, processor multitasking must be disabled in OS/400. Processor multitasking causes the iSeries processor to cache information when switching between tasks. This function is not supported by Linux in certain iSeries systems.

An IPL of the system is required to activate the change. To disable processor multitasking from an OS/400 command line, change the QPRCMLTTSK system value to 0 (CHGSYSVAL QPRCMLTTSK '0') in the primary partition.

If a guest partition is started without disabling the multitasking function, the IPL of the guest partition will fail with system reference code (SRC) B2pp 8105 000010AE.

Before creating a guest partition, disable processor multitasking on systems with the following feature codes:

Servers	Feature codes requiring QPRCMLTTSK (0)
820	<ul style="list-style-type: none"> <li>• 2397</li> <li>• 2398</li> <li>• 2426</li> <li>• 2427</li> </ul>
830	<ul style="list-style-type: none"> <li>• 2400</li> <li>• 2402</li> <li>• 2403</li> <li>• 2351</li> </ul>
840	<ul style="list-style-type: none"> <li>• 2418</li> <li>• 2420</li> <li>• 2416</li> <li>• 2417</li> <li>• 2419</li> </ul>

## System reference codes (SRCs) for Linux

The following list contains SRCs specific to Linux and suggested corrective actions. To display SRCs, see “Display the reference code history for secondary partitions” on page 28 for a list of the last 200 reference codes for a partition.

Additional SRCs for logical partitions can be found at Resolve system reference codes for logical partitions. If an SRC is not listed, it may not be related to logical partitions. You should consult Troubleshooting and Service documentation or your next level of service.

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

### 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 guest partition will not IPL.

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

**Cause:** The size for 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 guest 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 system 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 guest partition.

**Recovery:** Collect the detailed hexadecimal data from the hosting partition 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 system 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 guest partition.

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

**Problem analysis procedure:** Review the SRC history of the hosting partition.

**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 guest partition.

**Recovery:** Collect the detailed hexadecimal data from the hosting partition 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 system 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. A 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 OS/400

## Recover from a virtual Ethernet network failure

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=c00000000ff2800)
```

In this case, the message is telling you that the OS/400 virtual line description veth0 now relates to a Linux net device of eth0, and veth=0 correlates to VLAN0 in OS/400.

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 OS/400 which is VLAN0.





Printed in U.S.A.