



IBM Systems - System i

Linux

Linux on an integrated xSeries solution

*Version 5 Release 4*







IBM Systems - System i

Linux

Linux on an integrated xSeries solution

*Version 5 Release 4*

**Note**

Before using this information and the product it supports, read the information in “Notices,” on page 95.

**First Edition (February 2006)**

This edition applies to version 5, release 4, modification 0 of i5/OS (product number 5722-SS1) and to all subsequent releases and modifications until otherwise indicated in new editions. This version does not run on all reduced instruction set computer (RISC) models nor does it run on CISC models.

© Copyright International Business Machines Corporation 2006.

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

---

# Contents

<b>Linux on an integrated xSeries solution</b>	<b>1</b>		Managing your integrated Linux server with the i5/OS command line . . . . .	62
What's New for V5R4 . . . . .	1		Using iSeries Navigator to work with integrated Linux server NWSDs . . . . .	63
Printable PDF . . . . .	2		Using the i5/OS command line to work with integrated Linux server NWSDs . . . . .	63
x86 Linux on iSeries . . . . .	2		Starting integrated Linux servers . . . . .	64
Which environment to use. . . . .	3		Shutting down integrated Linux servers . . . . .	66
Planning the installation . . . . .	3		Networking integrated servers . . . . .	69
Integrated Linux server overview . . . . .	3		Configure virtual Ethernet networks . . . . .	69
Linux operating systems tested . . . . .	4		Managing Linux drives . . . . .	70
Hardware support . . . . .	5		iSeries storage management overview . . . . .	70
Reviewing the planning documentation . . . . .	6		Introducing network server storage spaces . . . . .	71
Understanding virtual devices . . . . .	7		Understanding storage space linking . . . . .	72
Installing the hardware . . . . .	8		Manage Linux drives . . . . .	74
iSCSI Host Bus Adapter . . . . .	8		Creating Linux drives . . . . .	75
Integrated xSeries Adapter. . . . .	9		Linking Linux drives . . . . .	77
Integrated xSeries Server . . . . .	11		Unlinking Linux drives . . . . .	78
Completing the preinstallation tasks . . . . .	13		Deleting Linux drives . . . . .	79
Install the required iSeries licensed program for Linux integration . . . . .	13		Sharing iSeries tape and optical drives with Linux . . . . .	79
Installing the PTFs required for Linux integration support. . . . .	14		Restricting iSeries tape and optical drives . . . . .	81
Completing the iSCSI preinstallation tasks . . . . .	16		Deleting Linux servers. . . . .	82
Creating the driver diskette for Linux installs on the IXS and IXA . . . . .	16		Viewing messages for a Linux server . . . . .	82
Connecting a USB diskette drive to the IXS. . . . .	18		Linux-centric backup and recovery . . . . .	84
Connecting the integrated server into the Ethernet LAN . . . . .	18		Linux backup and recovery overview. . . . .	84
Starting the installation from i5/OS . . . . .	18		Linux recovery options . . . . .	85
INSLNXSVR command parameter descriptions . . . . .	23		Choosing a tape drive for use by your Linux backup application . . . . .	86
Examples: Running the Install Linux Server (INSLNXSVR) command . . . . .	49		Restricting iSeries tape drives that can be used by Linux . . . . .	86
Completing the installation from the Linux console . . . . .	50		Setting up an iSeries tape drive for use by Linux . . . . .	87
Completing the installation of RHEL 3 on the IXS or IXA . . . . .	50		Formatting tape media for use by Linux. . . . .	87
Completing the installation of RHEL 4 on the IXA . . . . .	51		Transferring control of a tape drive from i5/OS to Linux . . . . .	88
Completing the installation of SLES 8 on the IXA . . . . .	52		Transferring control of a tape drive from Linux to i5/OS . . . . .	90
Completing the installation of SLES 9 on the IXA . . . . .	54		Backing up files using Linux utilities and applications . . . . .	91
Completing a SLES 9 SP3 or later installation on an iSCSI attached BladeCenter or xSeries . . . . .	55		Troubleshooting integrated Linux servers . . . . .	91
Completing the postinstallation tasks . . . . .	57		Related information . . . . .	92
Run the ixsetup.sh script . . . . .	57		<b>Appendix. Notices . . . . .</b>	<b>95</b>
Run the ixsupdt utility . . . . .	58		Trademarks . . . . .	97
Maintaining the Linux integration code . . . . .	58		Terms and conditions . . . . .	98
i5/OS operations for integrated Linux servers . . . . .	58			
Overview . . . . .	58			
Setting up your operational interface . . . . .	60			
Managing your integrated Linux server with iSeries Navigator . . . . .	61			



---

## Linux on an integrated xSeries solution

This will help you learn about using x86-based Linux<sup>®</sup> solutions on the iSeries<sup>®</sup> integrated server platforms.

The x86-based integrated server platforms for iSeries servers consist of:

- Integrated xSeries Server (IXS)
- Integrated xSeries Adapter attached xSeries (IXA)
- iSCSI HBA attached BladeCenter and xSeries (iSCSI HBA)

There are numerous applications that are developed for x86 Linux. Take advantage of all of these applications by integrating Linux using the integrated server platforms.

Linux running on the integrated server platforms allows you to use the virtual disk, virtual optical, and virtual tape of the iSeries. There are several benefits to putting Linux on the iSeries. Linux is built on the same concepts that integrated xSeries servers use.

- Management of i5/OS<sup>®</sup> and Linux servers from a central location.
- Dynamically link and unlink additional virtual storage for backup or other maintenance.
- i5/OS backup provides for the Linux storage area to be backed up with the normal i5/OS backup.
- Logical servers allow testing with the production image and hardware.
- Easily switch the production environment to another server if a maintenance or a hardware error causes the server to go down.
- Reduced total cost of ownership with the iSeries warranty and maintenance coverage of the IXS.

This topic collection is the updated version of the Implementing Linux on Integrated xSeries<sup>®</sup> Solutions for iSeries Redbooks<sup>®</sup>. See “Related information” on page 92.

---


### What's New for V5R4

| This topic highlights changes made to the Linux on an integrated xSeries solution topic collection for V5R4.


### What's new as of 22 September 2006

| Support for Linux on iSCSI attached servers is added.


| The iSCSI HBA attached BladeCenter<sup>®</sup> and xSeries integrated server platforms on iSeries now support x86 64-bit Linux. The SUSE Linux Enterprise Server 9 is supported on iSCSI attached servers.

| To see the current list of Linux distributions supported on the integrated server platforms, see the Linux on integrated servers  page on the System i integration with BladeCenter and System x web site ([www.ibm.com/systems/i/bladecenter/linux/](http://www.ibm.com/systems/i/bladecenter/linux/))

| <http://www.ibm.com/systems/i/bladecenter/linux/> 



| For information about installing the Linux operating system on iSCSI attached hardware, see the iSCSI install readme first  page on the System i integration with BladeCenter and System x web site.

## | What's new

| This information includes and updates the information from the Implementing Linux on Integrated xSeries Solutions for iSeries  Redbook, SG24-6379-00.

## | How to see what's new or changed

| To help you see where technical changes have been made, this information uses:

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

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

---


## Printable PDF

Use this to view and print a PDF of this information.

To view or download the PDF version of this document, select Linux on an Integrated xSeries Solution (about 825 KB).

See "Related information" on page 92 for additional topics.

## Downloading Adobe Reader

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

### Related reference

"Related information" on page 92

---

## x86 Linux on iSeries

Learn how the x86 Linux on iSeries platform differs from other IBM® offerings.

Linux on x86-based platforms have a different implementation strategy than Linux on POWER™. IBM provides drivers to the distribution of Linux to function in this environment. The environment would look like a hosted partition in the iSeries. The same use of the virtual disk and I/O are used for the x86 version as used on POWER. The x86-based integrated server implementations give the ability to use PC adapters that are supported on Linux and allows for special application development.

*Table 1. Comparison between POWER hosted and x86 Linux environment.* This table compares the differences between the hosted POWER environment to the integrated server solutions for x86 Linux.

Resource	iSeries hosted POWER Linux	x86 Linux
Boot disk	*NWSSTG, STMF, via IOA, A or B slot	*NWSSTG
Other DASD	*NWSSTG or via IOA	*NWSSTG
Tape	Virtual or via IOA	Virtual
CD-ROM, DVD-RAM	Virtual or via IOA	Virtual
Ethernet LAN	Virtual or via IOA	Virtual or Via IOA
Console	Telnet to port 2301 on a host partition	Attached keyboard and display



## Which environment to use

There are two implementations of Linux for the iSeries platform. You can configure the Linux operating system to run on an iSeries logical partition or you can use an integrated x86 solution that includes IBM System x or BladeCenter hardware. It is important to understand which one meets your needs. The criteria for an application might dictate which to use.

Listed below are several points to consider for an x86 implementation:

- Does the application only run on x86 Linux?
- Is the application CPU intensive?
- Does the iSeries have enough CPW capacity to support a Linux on POWER implementation?
- Are special hardware devices required for an application that is not supported under the iSeries POWER Linux?
- Do you have an integrated server solution running Windows® today that you want to migrate to Linux?

Listed are things to consider for the POWER implementation.

- Do you need more control of resources on the iSeries?
- Do you want to use Dynamic processor movement?
- Do you need several small Linux partitions?

Other factors might dictate which method to use.

---

## Planning the installation

This topic covers planning for the installation of Linux on the integrated server platforms.

This provides the basic information needed to plan for the implementation of Linux on the integrated server platforms. The distributors provide special instructions for the installation of their Linux distributions on supported BladeCenter and xSeries models. You must follow these special instructions during the installation process.

The hardware installation and configuration is not covered in depth because it is identical to that which has been documented elsewhere for Windows integration. However, there are some minor differences in the supported hardware between integrated Windows and Linux servers, which are documented. The general considerations involved in installing the hardware for those who are unfamiliar with the integration product are summarized, and references are provided to more detailed information regarding supported hardware configurations. Postinstallation tasks you need to perform to complete the installation of the integrated Linux server are also covered.

## Integrated Linux server overview

This topic briefly describes the concepts of running the Linux operating system on an integrated server platform.

Linux implementations on an integrated server make use of virtual services provided by i5/OS and have some other differences compared with Linux running on a PC. The basic concepts for running the Linux or Windows operating systems on an integrated server platform are similar. See the Concepts topic in the Windows environment on iSeries topic collection for more information about the integrated server environment.


There is extensive documentation on the features and benefits of Windows and Linux integration on iSeries on the System i integration with BladeCenter and System x Web site:

<http://www.ibm.com/systems/i/bladecenter/> 

It is strongly recommended that you review this documentation if you are not already familiar with the product.

## Linux operating systems tested

Supported integrated server hardware and the Linux operating systems that have been tested with it are listed here.

The most recent version of this information is available on the System i integration with BladeCenter and System x Web site at <http://www.ibm.com/systems/i/bladecenter/linux/> .

The following abbreviations are used for distributions of the Linux operating system:

### RHEL3

Red Hat Enterprise Linux 3 Advanced Server (AS) or Enterprise Server (ES)

### RHEL4

Red Hat Enterprise Linux 4 Advanced Server (AS) or Enterprise Server (ES)

SLES8 SuSE LINUX Enterprise Server 8

SLES9 SuSE LINUX Enterprise Server 9

## Linux operating systems that have been tested with iSCSI attached BladeCenter and xSeries models

Only 64-bit, x86 versions of Linux have been tested on this hardware.

Table 2. Linux operating systems that have been tested with iSCSI attached BladeCenter and xSeries models

Integrated server hardware	Linux operating systems tested
x236	SLES9
x336	SLES9
x346	SLES9
x366	SLES9
x460	SLES9
x3850	SLES9
x3950	SLES9
HS20 (8843)*	SLES9
LS20 (8850)*	SLES9

**Note:** The blade models have only been tested with the BladeCenter 8867 model chassis.

## Linux operating systems that have been tested with IXS or IXA attached models

Only 32-bit, x86 versions of Linux have been tested on this hardware.

Table 3. Linux operating systems that have been tested with IXS and IXA-attached models


Integrated server hardware	Linux operating systems tested
2892	RHEL3
4812	RHEL3

Table 3. Linux operating systems that have been tested with IXS and IXA-attached models (continued)

Integrated server hardware	Linux operating systems tested
x235	RHEL3, SLES8, SLES9
x236	RHEL3, RHEL4, SLES8, SLES9
x255	RHEL3, RHEL4, SLES8, SLES9
x346	RHEL3, RHEL4, SLES8, SLES9
x360	RHEL3, SLES8, SLES9
x365	RHEL3, RHEL4, SLES8, SLES9
x440	RHEL3, SLES8
x445	RHEL3, RHEL4, SLES8, SLES9

**Note:** Although some Linux distributions support more than 8 processors, only 8 processors are supported on the integrated server platforms.

**Attention:** The Linux distributors may require special instructions for the installation on supported BladeCenter and xSeries models. You must follow these special instructions. Any special instructions can be found by checking the hardware certification list as described below.

- For Red Hat: <http://bugzilla.redhat.com/hwcert>  On this page you will see a **Quick Search** text field. Input your xSeries or BladeCenter model and click **Search**. You will see the associated Red Hat Enterprise Linux version listed. Click on your specific model to see details on the certification summary and details. In the details under Platform, **i386** means 32-bit, and **x86\_64** means 64-bit version of the operating system.

- For SUSE: <http://developer.novell.com/yessearch/Search.jsp>  On this page input the following and click **Search**.

**Keywords:** Input your BladeCenter or xSeries model

**Company:** IBM

**Product:** SUSE Linux Enterprise Server 8 for x86 or

SUSE Linux Enterprise Server 9 for x86 or

SUSE Linux Enterprise Server 9 for AMD64 & Intel EM64T

The **x86** means 32-bit and the **AMD64 & Intel EM64T** means the 64-bit version of the operating system.

Your model should be listed. Click on the bulletin number to see details on configurations and requirements for install.

## Hardware support

This topic describes the integrated server hardware platforms for Linux.

### iSCSI Host Bus Adapter support

The iSCSI Host Bus Adapter is hardware that allows you to attach a BladeCenter or xSeries system via an iSCSI network connection to iSeries. The i5/OS integrated server management and virtual device access is similar to the IXS and IXA, but the infrastructure is significantly different. To learn more about the iSCSI Host Bus Adapter support, see the following Web site:

<http://www.ibm.com/systems/i/bladecenter/iscsi/> 

## Integrated xSeries Adapter support

The Integrated xSeries Adapter is hardware that allows you to directly attach an xSeries system via a High Speed Link (HSL) connection to an iSeries. The i5/OS integrated server infrastructure for server management and virtual device access is the same as the Integrated xSeries Server.

| Linux on an integrated server has only been tested on some xSeries models. For details, see the *Linux distributions tested* section. For more information on supported xSeries servers, go to the following Web site:

<http://www.ibm.com/systems/i/bladecenter/ixa/servermodels/> 

**Note:** Pay particular attention to the footnoted material for your model.


You cannot connect an xSeries model that is not on this list to iSeries. The xSeries needs a CD-ROM drive and a diskette drive to install Linux. To determine the maximum number of Integrated xSeries Adapters you can connect to a particular model of the iSeries, go to the following Web site:

<http://www.ibm.com/systems/i/bladecenter/ixa/> 

## Integrated xSeries Server support

The Integrated xSeries Server is hardware that has been designed to allow you to simplify your server environment by reducing the hardware footprint of your x86-based servers, and by providing a degree of centralized management. The Integrated xSeries Server is installed inside the frame of the iSeries system, or an iSeries tower, and does not consume the space that a standalone server box might.

Linux on an integrated server has tested only on certain Integrated xSeries Server models.

For more information about configuring Integrated xSeries Servers, see <http://www.ibm.com/systems/i/bladecenter/ixs/> . This page provides general configuration information and will also help you determine the maximum number of Integrated xSeries Servers you can install on a particular model of the iSeries system.

A USB diskette drive is required to load the initial IBM-supplied drivers if you are installing Linux on an Integrated xSeries Server.

Note that you also need to connect an SVGA-compatible monitor, keyboard, and mouse to the Integrated xSeries Server to act as the Linux console.


## Reviewing the planning documentation

This topic covers the planning documentation for Integrated xSeries hardware.

The first task you must complete before you install Linux on the Integrated xSeries hardware is the planning documentation:

- | • **iSCSI Host Bus Adapter support**

| The iSCSI Host Bus Adapter is hardware that allows customers to attach a BladeCenter or xSeries  
| system via an iSCSI network connection to the iSeries. The i5/OS integrated server management and  
| virtual device access is similar to the IXS and IXA, but the infrastructure is significantly different. To  
| learn more about the iSCSI Host Bus Adapter support, see the following Web site:

| <http://www.ibm.com/systems/i/bladecenter/iscsi/readme/> 


- | • **Integrated xSeries Adapter**

If you are installing Linux on an Integrated xSeries Adapter-attached xSeries, work through the IXA install readme first .

**Note:** The IXA readme refers to a ServerGuide™ install. ServerGuide™ is xSeries software that simplifies the installation of a Windows server; however, it is not applicable to Linux.

Other documents that you should also review, especially the Integrated xSeries Adapter installation instructions:

- For pre-eServer i5 hardware, download:  
iSeries PCI Card and Integrated xSeries Adapter Card installation manual
- For eServer™ i5 hardware, download:

Integrated xSeries Adapter (IXA) installation manual 

These documents are particularly important when you come to physically install the Integrated xSeries Adapter card in the xSeries server, so be sure to use a current copy each time you do an install.

- **Integrated xSeries Server**

If you are installing Linux on an Integrated xSeries Server, work through the IXS install readme first .

## Understanding virtual devices

This topic describes different virtual devices that can be used with an integrated Linux server.

Integrated Linux and Windows servers feature virtual devices of various types. These devices can be confusing to a user who is new to the integrated server solutions products. Windows and Linux are essentially “guest” operating systems that support and manage the hardware that they interface to. As part of server consolidation, the integrated servers provide the capability to share hardware with i5/OS, thus reducing the amount of hardware required, and the extent of the hardware footprint that needs to be managed. The integrated server hardware utilizes virtual disk drives, virtual optical, virtual tape devices, and Virtual Ethernet LAN.

The integration software supports the following virtual devices:

### Virtual disk drives

A virtual disk is not optional for integrated servers. Although xSeries machines normally come with local hard drives, when an xSeries server is used in an integrated server environment, any installed hard drives or hard drive controllers must be removed.

A virtual disk is referred to as NWSSTG or a network server storage space. These objects are actually files in the integrated file system that appear to the integrated server operating system as hard drives. There can be a total of 32 of these virtual disk drive objects created and linked to each integrated server.

The installation command creates two virtual drives, which are referred to as the system drive and the install drive. We discuss these drives in more detail later.

All virtual drives for all the integrated servers on a particular iSeries are stored in the QFPNWSSTG directory, which is in the root directory of the Integrated File System. These devices implicitly utilize iSeries storage management, and thus data is spread across all the configured disks. Additionally, these drives are protected by iSeries RAID-5 if it has been configured, and therefore do not benefit from Windows or Linux software RAID.

### Virtual optical drives

The virtual optical function provides the capability for integrated Linux servers to utilize i5/OS optical drives. The drives need to be varied on i5/OS, and locked in Linux using the ixpdev command. You use an iSeries optical drive during the install of RHEL 3. Installation of all other

Linux distributions use the optical drive in the xSeries or BladeCenter systems. After Linux is installed, you can also use the iSeries virtual optical drive just like any other optical drive.

### Virtual tape drives

The virtual tape function provides the capability for integrated Linux servers to access i5/OS tape drives. The i5/OS tape drive must be varied off in i5/OS and locked in Linux using the `ixsdev` command. It is important to verify that the tape drive is supported for use with Linux integration support. A list of supported devices is maintained on the **System i integration with BladeCenter and System x** Web site:

[http://www.ibm.com/systems/i/bladecenter/product\\_info.html](http://www.ibm.com/systems/i/bladecenter/product_info.html)  .

### Virtual Ethernet

Virtual Ethernet LANs run across the same adapter as virtual disk, optical and tape devices. Therefore, no additional physical hardware is required. This type of connectivity can be used to communicate internally, as follows:

- Point-to-Point Virtual Ethernet LAN

As part of the installation, there is a line description and TCP/IP interface automatically created for the Point-to-Point Virtual Ethernet LAN. It is used to enable i5/OS to communicate with Linux to provide administrative control over the Linux server. The Point-to-Point Virtual Ethernet LAN is intended to be a secure connection that does not get routed to the external network.

- Other Virtual Ethernet connections

This function allows communication internally with other integrated servers or iSeries partitions (i5/OS, AIX® and POWER Linux) using virtual LAN connections. This technology could be used to offload the external network or to provide a fast, secure connection to a server or partition.

**Note:** Virtual Ethernet is only supported on IXA-attached xSeries servers. The Linux distributions that support virtual Ethernet include RHEL4 and SLES9. For the most current list of distributions that support virtual Ethernet, see the following Web site:

<http://www.ibm.com/systems/i/bladecenter/linux/>  .

---

## Installing the hardware

This topic describes how to install and configure the integrated server hardware.


This section is intended for those who are not familiar with the integrated server hardware. It provides an overview of the hardware, and points out some considerations for those who have not previously had experience with the integrated server platform.

### | iSCSI Host Bus Adapter

| This topic introduces the installation of the iSCSI HBA hardware.

### | iSCSI HBA Installation Overview

| The installation of the iSCSI HBA is a complex task that is much more involved than just plugging in the card. There are a number of preparation and planning steps that need to be followed. These steps are detailed in the *iSCSI install read me first* Web site at:

| <http://www.ibm.com/systems/i/bladecenter/iscsi/readme/>  .

- | At this stage you should follow the instructions from sections:
  - | • Prerequisites
  - | • Prepare for hardware installation
  - | • Hardware installation
- | **Note:** SLES9 on iSCSI only supports one iSCSI initiator to one iSCSI target configuration.

## Integrated xSeries Adapter

This topic discusses installation of the xSeries Adapter.

### IXA Installation Overview

This topic discusses an overview of the IXA installation.

Unlike the Integrated xSeries Server, the Integrated xSeries Adapter appears to the iSeries as a separate expansion tower. Most models of xSeries have their own diskette drive and optical drive, and can attach a tape drive. There are bays for disk drives but the server should be ordered without disk drives or disk controllers because these must be removed before connecting an xSeries to an iSeries using the Integrated xSeries Adapter. Integrated servers use virtual disk drives as well as virtual optical and tape drives. Most xSeries servers come with a diskette drive so there should be no need to connect a USB diskette drive to the server in order to install the Linux drivers.

Once an Integrated xSeries Adapter-attached xSeries is cabled into the iSeries HSL (high-speed link) loop and SPCN (system power control network) network, the xSeries should not be turned on or off using the power button on its front panel. In this case, power is controlled by iSeries, and the server can be turned on or off from iSeries Navigator or a green screen session.

In comparison with the Integrated xSeries Server, there is a larger degree of complexity in preparing an Integrated xSeries Adapter-attached xSeries server for installation of the Linux or Windows operating system because there are several components involved:

- xSeries BIOS
- RSA I or RSA II adapter and firmware (on certain models of the xSeries only)
- Advanced System Management software
- Firmware for the video adapter on the planar board or, in the case of the RSA II card, firmware for the video adapter on the RSA II card itself
- Integrated Ethernet port
- Integrated xSeries Adapter
- HSL loop and SPCN network

### xSeries hardware and firmware

This topic discusses the hardware and firmware needed for IXA.

The firmware should be updated as required by the particular xSeries model prior to installing the Integrated xSeries Adapter hardware, and before being cabled into the HSL loop or SPCN network.

Use the following link for downloading xSeries drivers, or for obtaining information related to xSeries firmware levels:

<http://www.ibm.com/pc/support/site.wss/> 

Alternatively, you can use UpdateXpress to update the xSeries firmware. UpdateXpress is a bootable CD that will automatically detect the level of firmware and present you with a list of suggested updates. Use the following web site to learn more about UpdateXpress. You can order the UpdateXpress CD or download the image from this website as well.



<http://www.ibm.com/pc/support/site.wss/document.do?lnocid=MIGR-53046> 

If you are still unsure about the correct firmware for your xSeries model you might need to contact xSeries support directly.

## Updating the xSeries BIOS for your integrated Linux server

This topic discusses xSeries BIOS, Advanced Systems Management, and how to maintain the most current level.

Verify the current level of BIOS for your xSeries model. This level can be obtained by pressing the F1 key during boot up and then displaying System Information. This allows you to check the BIOS and ASM levels installed on the xSeries.

If you need to upgrade the BIOS, see “Reviewing the planning documentation” on page 6 on page 5.

Before upgrading the BIOS be sure that the xSeries is not cabled into the HSL loop or SPCN network, and that it is turned on.

## RSA I or RSA II adapter and firmware

This topic describes the Remote Supervisor Adapter requirements for your integrated x86 Linux server.

Only certain xSeries models require the installation of a Remote Supervisor Adapter (RSA I or RSA II). This adapter is required to provide a management interface to the xSeries so that it can be controlled from i5/OS. For example, the RSA card is used to power up and power down the xSeries server. The RSA I or RSA II is only required for those xSeries servers that do not have the Remote Supervisor capability built in. You can determine whether your xSeries requires a Remote Supervisor Adapter by checking the following Web site:

<http://www.ibm.com/systems/i/bladecenter/ixa/servermodels/> 

Those xSeries servers requiring an RSA I or RSA II are noted.

If you are installing either of these cards make sure the xSeries is not turned on and is not cabled into the iSeries HSL and SPCN connections.

To install an RSA card refer to the iSeries PCI Card and Integrated xSeries Adapter (IXA) Card or Integrated xSeries Adapter (IXA) installation, see “Reviewing the planning documentation” on page 6.

**Important:** Read the RSA II documentation carefully when installing this card.

## Advanced System Management software

This topic discusses the advanced system management software.

It is recommended that you update this software to the latest level.

## Video card

This topic discusses video card requirements for your integrated Linux server hardware.

If you have an RSA II card in your xSeries it overrides the planar video adapter and you need to manually update the video driver as well as the RSA firmware. During the Power On Self Test, you see messages indicating that this is required. The update process is unique, see “Reviewing the planning documentation” on page 6 for more information on updating firmware.

**Note:** If your xSeries has an RSA II card, the display for your server needs to be connected to the video port on the RSA II itself.



## Integrated xSeries Adapter card

This topic discusses installation of the xSeries Adapter card.

If the RSA I or RSA II adapter is required for your xSeries server then it should be installed, and the firmware upgraded prior to installing the Integrated xSeries Adapter.

After the BIOS and firmware for the xSeries and the RSA adapter have been updated, you can install the Integrated xSeries Adapter into the xSeries Server. This requires that the xSeries be disconnected from the AC power.




Prior to installing the Integrated xSeries Adapter, a jumper needs to be set in many xSeries models to provide 3.3 volt standby power to the slot where the card is located. The jumper number is different for the various models, so verify the correct jumper settings in the iSeries PCI Card and Integrated xSeries Adapter Card or Integrated xSeries Adapter (IXA) installation manual.

## HSL loop and SPCN network

This topic discusses how to install HSL loop and SPCN network installation for an Integrated xSeries Adapter (IXA).

The Integrated xSeries Adapter (IXA) has connections for HSL, SPCN, and the xSeries service processor. The iSeries must be turned off before you cable up the Integrated xSeries Adapter installed in the xSeries. The i5 server can concurrently attach IXAs while the server is turned on.

HSL and SPCN configuration and connection are beyond the scope of this document, so refer to the following publications:

- Microsoft® Windows Server 2003 Integration with iSeries, SG24-6959, at the following Web site:  
<http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg246959.html> 
- Integrated xSeries Adapter and Migration Tower Coexistence in High-speed Link Loops, TIPS0082:  
<http://www.redbooks.ibm.com/abstracts/tips0082.html?Open> 
- Install High Speed Link (HSL) in the IBM Systems Hardware Information Center:  
<http://publib.boulder.ibm.com/series/v5r4/ic2924/index.htm?info/rzame/rzamehsl.htm> 

After the Integrated xSeries Adapter is installed, verify that it is operational by using the Work with Communications Resources command (WRKHDWRSC \*CMN). Check to see that the 2689 resources are present and operational.

If you have multiple Integrated xSeries Adapters, use the same procedures as described in “Verifying that an Integrated xSeries Server (IXS) is operational” on page 12 to correlate a resource name (LINxx) with a physical Integrated xSeries Adapter. Note that the Frame ID can be determined by removing the lid of the xSeries and reading the number on the LED located on one side of the Integrated xSeries Adapter (The frame ID is no longer used on i5 servers.). The LED displays the Frame ID in decimal, which you need to compare with the Frame ID next to the LINxx resource and press **Enter**.

## Integrated xSeries Server

This topic introduces the installation of the IXS hardware.

### IXS Installation overview

This topic discusses an overview of the IXS installation.

**Attention:** Only recent models of the Integrated xSeries Server can be used with the Linux operating system. Refer to the “Linux operating systems tested” on page 4 section to verify that your Integrated xSeries Server is valid.

The Integrated xSeries Server card installs into slots in the iSeries or one of its towers. It is basically an xSeries motherboard that can be inserted into an iSeries PCI slot. However, it lacks an integrated diskette drive, tape drive, optical drive, and hard drive. Virtual disk, tape, and optical drives can be made available to Linux running on the Integrated xSeries Server by the Linux integration support. Although there is no virtual diskette drive capability, you can connect a USB diskette drive to one of the USB ports on the Integrated xSeries Server.

It is important to understand that the Integrated xSeries Server must be installed in pre-defined PCI slots in the iSeries. Additionally, because it is a double-wide card, it takes up two slots, although it only plugs into one of them. To determine the possible slots that the Integrated xSeries Server can use, refer to one of the following documents:

- IBM eServer for i5, iSeries, and AS/400e™ System Builder, SG24-2155
- PCI Card Placement Rules for the IBM eServer iSeries Server, REDP-3638

You can download these publications from the following Web site:

<http://www.redbooks.ibm.com/> 

Each iSeries model has a maximum number of Integrated xSeries Server cards it supports based on the number and type of towers installed:

<http://www.ibm.com/systems/i/bladecenter/ixs/> 

Connectors for peripheral devices are on the card itself. You can connect the following devices:

- A 10/100/1000 Mbps Ethernet adapter is built into the card. Some models of the Integrated xSeries Server can also control up to three additional LAN adapters installed in adjacent slots. These can be 4/16/100 Mbps token ring, or 10/100/1000 Mbps Ethernet.
- An SVGA compatible monitor plugs directly into the card.
- A keyboard and mouse can be connected to the Integrated xSeries Server card using the supplied splitter cable, which plugs into a single port on the card. It is a relatively common misconception that you can only plug one of these devices directly into the port. In this case you would be unable to locate a port for the second device. You might conclude that you need to obtain a USB keyboard or mouse as the second device, but this is not the case.

The keyboard, mouse, and monitor can be connected to a KVM switch so they can be shared between multiple computers, but be aware that there can be compatibility problems with some switches.

- There are two USB ports that are available for you to connect peripheral devices.

**Important:** The installation of Linux on the Integrated xSeries Server requires a diskette drive. This is so that drivers can be loaded in order for Linux to be able to access iSeries virtual devices including the optical drive. Most xSeries servers have an onboard diskette drive, but the Integrated xSeries Server does not. Therefore a USB port must be used to plug in a USB diskette drive.

## Verifying that an Integrated xSeries Server (IXS) is operational

This task describes how to verify that an Integrated xSeries Server is operational.

### About this task

To verify that an Integrated xSeries Server is operational, follow these steps:

1. On an i5/OS command line, type `WRKHDWRSC *CMN` (Work with Communications Resources command) and press **Enter**.
2. Look under the Type column for a resource shown as 2892 or 4812. Verify that the Status column shows that all the 2892 or 4812 resources are Operational.

## Locating an Integrated xSeries Server (IXS)

This task describes how to locate the correct Integrated xSeries Server.

### About this task

If there is more than one set of 2892 or 4812 resources (LINxx) and you are unsure which resource corresponds to the physical Integrated xSeries Server you want to install the Linux operating system on, follow these steps:

1. On an i5/OS command line, type `WRKHDWRSC *CMN` (Work with Communications Resources command) and press **Enter**.
2. Look under the Type column for a resource shown as 2892 or 4812. Next to one of the LINxx resources, press **Enter**.
3. You can use the Frame ID and Card slot values to determine the physical location of the Integrated xSeries Server card (The frame ID is no longer used on i5 servers.) The Frame ID denotes the iSeries tower, and the Card position can be found by looking at the labels above or beside each PCI slot in the tower. You can also determine the Integrated xSeries Server model.

## Determining which Integrated xSeries Servers are in use

This task describes how to determine what IXS servers and resources are in use on your system.

### About this task

In a situation where you have multiple Integrated xSeries Servers and Linux is already installed on some of them, you might also need to determine which resources are already in use. You can use the Work with Network Server Descriptions command (`WRKNWSD`) to display a list of current Network Server Descriptions. To determine which Integrated xSeries Servers are in use, follow these steps:

1. On an i5/OS command line, type `WRKNWSD` and press **Enter**.
2. Type 5 next to each NWSD and press **Enter**. You will see a description of the NWSD. This display shows the Resource name used by the server and helps you identify which resources are already in use.
3. After displaying each NWSD and noting the associated resource name, compare these against the resources shown in the `WRKHDWRSC *CMN` display to determine which resources are not already configured.

You can create more than one integrated Linux server (with its corresponding NWSD) on an individual resource, but only one Linux server can be active on a particular resource at a time. Therefore, you might in fact find that you have more than one NWSD with the same resource name.

**Note:** If you are working on an iSeries with multiple i5/OS partitions you might also need to determine which partition the Integrated xSeries Server hardware is allocated to.

---

## Completing the preinstallation tasks

Before installing Linux on your integrated server, the following tasks need to be completed.

## Install the required iSeries licensed program for Linux integration

This topic discusses the i5/OS options and licensed programs that are required for integrated x86 Linux servers.

### About this task

Linux integration support requires the following iSeries licensed programs:

- 5722-SS1 opt. 29 Integrated Server Support
- 5722-LSV IBM i5/OS Integration for Linux on xSeries

- 5722-XE1 IBM iSeries Access for Windows which is needed for iSeries Navigator
- 5722-TC1 IBM TCP/IP Connectivity Utilities for i5/OS

l Additionally, the following are specific to iSCSI:

- 5722-SS1 opt. 03 Extended Base Directory Support
- 5722-SS1 opt. 30 Qshell
- 5722-SS1 opt. 34 Digital Certificate Manager
- 5722-JV1 Java Development Kit 1.4, Option 6.
- 5722-DG1 IBM HTTP Server for i5/OS which installed when you install 5733-VE2 IBM Virtualization Engine, SE version
- 5722-DR1 IBM Director Server

To verify that i5/OS licensed programs are installed, complete the following steps:


1. Type G0 LICPGM on a command line and press **Enter**. The Work with Licensed Programs menu is displayed.
2. Type 10 and press **Enter**. You see the Display Installed Licensed Programs display. Scroll down to verify that license programs are installed.
3. If the license programs are not installed you need to find the media that the system was installed from.
4. Insert the CD-ROM containing the license programs into the iSeries optical drive.
5. Return to the Work with Licensed Programs display and select option 11, Install licensed programs.
6. Page down the list of program products and type 1 next to the license program.
7. Verify that you selected the correct options on the confirmation display. Press **Enter**.
8. Type the device name for the optical drive being used. Press **Enter**.
9. The Linux integration software is now installed.

## Installing the PTFs required for Linux integration support

This topic discusses the required PTFs for installing integrated Linux support.

### Before you begin

Before you attempt to install Linux on an integrated server, you should ensure that all the PTFs listed on the following Web site are installed on your iSeries server:

<http://www.ibm.com/systems/i/bladecenter/ptfs.html> 

Most problems that arise during and after installation are due to the fact that required PTFs have not been installed.

PTFs fall into two categories:

- Cumulative and group PTFs
- Individual PTFs

### Cumulative PTFs

#### About this task

Cumulative PTF packages contain code updates for most of the iSeries Licensed Program Products including the license programs required for Linux integration support. They are made available on a regular basis. As the name implies, they are cumulative in that each package contains all PTFs supplied

with previous packages for a particular i5/OS release. Group PTFs are very similar to cumulative PTFs, but only contain PTFs for a particular program product or set of related products.

You should order the latest cumulative package well in advance so that you can install it at least a week before you plan to install Linux.

To determine the latest cumulative PTF level installed on your iSeries, complete the following steps:


Type `DSPPTF LICPGM(5722999)`, the Display PTF command. This command displays a list of PTFs for the i5/OS Systems Licensed Internal Code (SLIC). The cumulative PTF package numbers are at the top.

## Results

The top-most number (for example, TL03364) is the most current CUM, and 03364 is the date of release of the package, that is, day 364 of 2003 in the example.

Take note of the Status and IPL Action. The Status should be either Temporarily applied or Permanently applied. If the IPL Action is shown as anything other than None, enter 5 next to the PTF to determine if an IPL of the i5/OS partition or some other action is required. Note that cumulative PTFs require a restart (IPL) of the i5/OS partition.

You can determine how far behind you are with your cumulative updates by comparing the TLxxxxx number with the latest CUM listed on the following Web site:

<http://www.ibm.com/systems/i/bladecenter/ptfs.html> 

## Individual PTFs

### About this task

There might be PTFs available specifically for the iSeries Linux integration environment that have not yet been included in a cumulative PTF package. Linux integration support code fixes can be included in PTFs for iSeries Licensed Program products including:

- 5722-999: System Licensed Internal Code
- 5722-SS1: i5/OS
- 5722-LSV: IBM i5/OS Integration for Linux on xSeries
- 5722-XE1 IBM iSeries Access for Windows which includes iSeries Navigator.
- 5722-TC1 IBM TCP/IP Connectivity Utilities for i5/OS
- 5722-JV1 Java Development Kit 1.4, Option 6
- 5722-DR1 IBM Director Server

A list of all available PTFs not included in the latest cumulative PTF package can be found on the **System i integration with BladeCenter and System x** Web site at:

<http://www.ibm.com/systems/i/bladecenter/ptfs.html> 

You must install all individual PTFs listed for your release that are listed on this Web site. To determine if the listed PTFs are already installed on your system, use one of the following commands:

- For MFxxxxx (i5/OS SLIC)  
PTFs: `DSPPTF LICPGM(5722999) SELECT(MFxxxxx)`
- For SIxxxxx (5722-LSV)  
PTFs: `DSPPTF LICPGM(5722LSV) SELECT(SIxxxxx)`
- For SLxxxxx (5722-SS1)

| PTFs: DSPPTF LICPGM(5722SS1) SELECT(SIxxxxx)

## | **What to do next**

| MFxxxxx PTFs often require a restart (IPL) of the hosting i5/OS partition, while SIxxxxx PTFs rarely do.  
| If you need assistance applying PTFs, go to the information center at the following Web site:

| Install fixes 

| Updates to the integration code running on Linux are called service packs, and are provided in the form of PTFs to 5722-LSV. A Linux utility is used to download the integration service pack to the Linux server.

| **Note:** Do not confuse iSeries integration service packs with service packages for the Linux distributions.  
| The integration service packs only address issues in the i5/OS Linux support code, whereas  
| distribution service packs address issues in Linux.

## | **Completing the iSCSI preinstallation tasks**

| This topic discusses the iSCSI related tasks required immediately before Linux installation.

### | **About this task**

| For iSCSI attached xSeries or BladeCenter, there are number of tasks that are detailed in the *iSCSI install read me first* Web site at: <http://www.ibm.com/systems/i/bladecenter/iscsi/readme/>

### | **Example**

| At this stage you should follow the instructions from the **Prepare for operating system installation** section.

## **Creating the driver diskette for Linux installs on the IXS and IXA**

This topic discusses how to create a driver diskette that is required for Linux installation on the IXS and IXA.

### **About this task**

The driver diskette is required during installation so that the integration software needed to provide virtual devices for the IXS & IXA can be installed in Linux. To create the diskette, complete the following steps:

1. Locate the diskette images in the QIBM directory of the integrated file system.
2. From an i5/OS command line type the command `WRKLNK OBJ('/QIBM/ProdData/IntegratedOS/Linux/x86/Install/*')`
3. Look for an IMG file for either Red Hat Enterprise Linux or for SUSE Linux Enterprise Server, depending on what Linux distribution you have selected to use.

**Note:** Whatever versions of the distributions are supported when you perform your installation, there will be an IMG file for each one in the integrated file system directory `/QIBM/ProdData/IntegratedOS/Linux/x86/Install`.

4. On a PC running either Windows or Linux, make a directory on the local hard drive to store the diskette image:
  - For example, in Windows: MD C:\Redhat  
You can use any drive and directory name you wish.
  - For example, in Linux: `mkdir /home/redhat`



You can use any drive and directory name you wish.

5. Transfer the image file from the integrated file system directory `/QIBM/ProdData/IntegratedOS/Linux/x86/Install` to the directory you created in the previous step using one of the following methods:

- FTP the image file to the PC directory as a binary file.
- File sharing

Access the integrated file system directory using iSeries NetServer™ from Windows, or from Linux using Samba. Then copy the image file to the PC directory.

The `/QIBM` integrated file system directory is shared by default using iSeries NetServer. You should be able to access this directory and then drill down to `/QIBM/ProdData/IntegratedOS/Linux/x86/Install` without any setup.

If you have a problem accessing the `/QIBM` shared integrated file system directory you might need to review the setup of iSeries NetServer. In this case, go to the iSeries NetServer topic.

You can connect to the `/QIBM` integrated file system directory as follows:

#### Windows

- a. Using Windows, click **Start** → **Run** and enter the IP address of the i5/OS partition; for example, `\\161.50.4.3`
- b. Click **OK**. Drill down to the directory `/QIBM/ProdData/IntegratedOS/Linux/x86/Install` and copy the file to your PC.

#### Linux

- a. Using Linux, to access the integrated file system directory using Samba, you need to use the `SMBMOUNT` command to mount the `/QIBM` integrated file system directory in the Linux file system. The format of the command is: `smbmount //host-name/qibm /mnt/mydata -o username=user-name,password=password`
  - *host-name* is the iSeries TCP/IP host name or IP address.
  - *user-name* is a valid i5/OS profile name.
  - *password* is the password for user-name.

For more information on using the `smbmount` command to access an iSeries NetServer shared directory refer to the document at the following Web site:


[http://www.ibm.com/servers/eserver/iseries/netserver/pdf/LinuxNetServ\\_Usage.pdf](http://www.ibm.com/servers/eserver/iseries/netserver/pdf/LinuxNetServ_Usage.pdf) 

6. If you copied the diskette image to a Linux PC, create the diskette as follows:
  - a. Insert a formatted, blank diskette into the diskette drive.
  - b. Change the directory to the directory in which the image is stored.
  - c. Use the `dd` command to transfer the diskette image to the physical diskette. In the example below, the `dd` command is being issued from inside the `src` directory.

```
dd if=rhel3ixs.img of=/dev/fd0
```

The input file (`if`) is the image that was downloaded to the `src` directory. The output file (`of`) is the diskette drive.
  - d. To verify that the operation was successful you can mount the floppy device and list the files using the `ls -l` command.
7. If you copied the diskette image to a Windows PC, create the diskette as follows:
  - a. Insert a formatted, blank diskette into the PC's diskette drive.
  - b. For RHEL 3, SLES 8 or SLES 9, insert one of the following CD-ROMs into the PC's CD-ROM drive:
    - Red Hat Enterprise Linux 3 ES or AS: disc 1
    - SUSE Linux Enterprise Server 8: United Linux 1.0 disc 1
    - SUSE Linux Enterprise Server 9: disc 1

- c. Use Windows Explorer to navigate to the \dosutils\rawwritewin directory (RHEL 3) or \dosutils\rawwritewin directory (SUSE 8 or 9) on the optical drive and open the folder.
- d. For RHEL 4, the installation media does not contain any utilities for creating a driver diskette on Windows. However you can obtain the rawwritewin.exe utility from the above media or download versions 0.4 or above from the web site.

<http://uranus.it.swin.edu.au/~jn/linux/rawwrite.htm> 

**Note:** IBM does not endorse or accept any responsibility for the content, or the use, of this Web site. It is up to you to take precautions to ensure that whatever you select for your use is free of such items as viruses, worms, trojan horses and other items of a destructive nature.

- e. Double click **rawwritewin.exe**.
- f. Enter the path and name of the diskette image in the Image file text box.
- g. Click the **Write** button to write the image to diskette. When the diskette has been created, remove it and click **Exit**.
- h. You cannot read the diskette you have created from Windows. To verify the contents of the diskette you would need to use a Linux PC.

## Connecting a USB diskette drive to the IXS

This topic discusses USB diskette drives used for installation on the integrated xSeries server.

### About this task

If you are installing Linux on an Integrated xSeries Server you need to obtain a USB diskette drive. Plug the drive into one of the USB ports on the Integrated xSeries Server. Do not insert the driver diskette yet. If you have any of the USB ports already in use on the IXS, use the next available USB port (closest to the currently used one) for plugging the diskette drive.

## Connecting the integrated server into the Ethernet LAN

This topic discusses the external Ethernet connection which is required for all integrated server types (IXS, IXA, iSCSI).

### About this task

Make sure that you connect the external Ethernet connection to the network before you begin your Linux install so that the Linux install program can configure this connection.

### Example

**Note:** The external Ethernet connection can be used by the ixsetup.sh and ixsupdt programs which will be used during the postinstallation steps.

---

## Starting the installation from i5/OS

This topic describes the installation of Linux initiated from an i5/OS command line.

### About this task

The Install Linux Server (INSLNXSVR) command initiates the install of the Linux server operating system on an integrated server. INSLNXSVR also copies the integrated server support code on the install drive which will be installed during the post installation steps.

Linux server installation occurs in two steps. During the first step, the INSLNXSVR command creates all necessary objects to manage the server. This includes a network server description, message queue, line



descriptions, storage spaces and TCP/IP interfaces. For installs using the iSCSI HBA, network storage configuration objects for remote system, service processor and connection security are created by default unless the pre-existing objects are specified.

During the second step of the Linux server installation, the integrated server is varied on to start the Linux server installation.

Further Linux server installation is performed using the file server console and the normal Linux server install process.

When INSLNXSVR completes normally, the Linux server is left in a varied on state.

### **Restrictions:**

- You must have input/output system configuration (\*IOSYSCFG), all object (\*ALLOBJ) and job control (\*JOBCTL) special authorities to run this command.
- The file server must be varied off initially.
- The file server will be varied off and varied back on during the second step of the install as Linux server installs and requires the server to reboot.

### **Usage Notes:**

Any errors that occur during the first step of configuring the file server will result in the failure of this command.

After this command is run, if you need to manage the different resources created, use the following commands:

- To check out the status of the Linux server, use the Work with Configuration Status command; `WRKCFGSTS CFGTYPE(*NWS)`.
- To manage the server just installed, use the Work with Network Server Descriptions command; `WRKNWSD NWSD(network-server-name)`.
- To manage the line descriptions created by this command, use the Work with Line Descriptions command; `WRKLIND LIND(nwsdname*)`. The line descriptions are named using the network server name (NWSD parameter) specified on the INSLNXSVR command.
- To manage the TCP/IP interfaces created by this command, use the Work with TCP/IP Network Status (NETSTAT) command, option 1. Another option is to use the Configure TCP/IP (CFGTCP) command, option 1.
- To manage the network server configurations just created by this command, use the Work with NWS Configuration command; `WRKNWSCFG NWSCFG(nwsdname*)`. The network server configurations are named using the network server name (NWSD parameter) specified on the INSLNXSVR command.

If the installation fails at an early stage, all i5/OS objects that have been created are removed. However, if the installation fails after being transferred to the Linux console, the i5/OS objects remain. This is because you might be able to correct the error and recover the installation. In this case, restart the server and the installation code will try to resume from the point of failure. If you are unable to recover, or you decide to restart the installation from scratch, you can delete the Linux server instance using the DLTLNXSVR command.

To start the installation of Linux on the integrated server, complete the following steps:

1. If you are installing on an iSCSI attached BladeCenter or xSeries, complete the tasks listed in the section titled *Prepare for operating system installation* from the **iSCSI install read me first** Web site: <http://www.ibm.com/systems/i/bladecenter/iscsi/readme/>
2. If you are installing Linux on an IXA attached xSeries, locate the driver diskette but do not insert it into the diskette drive of the xSeries server yet.
3. If you are installing Linux on an IXS:
  - a. Verify that the USB diskette drive is plugged into a USB port on the server.
  - b. Verify that the iSeries optical drive is varied on and that it is not being used by another job.
  - c. Insert the Linux installation CD in the iSeries optical drive.
4. Verify that no other NWS D is varied on for the hardware resource you want to install on by running the CL command, and typing WRKCFGSTS \*NWS. Display each NWS D that is varied on and confirm that it is not using the resource you intend to use.
5. Sign on to an i5/OS green screen session with a profile that has the following authorities: \*IOSYSCFG, \*ALLOBJ, and \*JOBCTL or \*SECADM.
6. Change the QSYSOPR message queue to something other than \*BREAK so that the installation is not interrupted. For example, type the CL command CHGMSGQ MSGQ(QSYSOPR) DLVRY(\*NOTIFY).
7. On the i5/OS command line type **INSLNXSVR** and press **F4**.

This will display the parameters for the command. Below is a brief description of the required parameters. All other parameters are filled in with default values which can be modified as you prefer. For more details on the parameters, see “INSLNXSVR command parameter descriptions” on page 23 and “Examples: Running the Install Linux Server (INSLNXSVR) command” on page 49.

- a. Enter information the following parameters:

#### **Network Server Description**

Choose any name of up to eight characters that is meaningful to you. This does not have to be identical to the Linux server’s host name. It is used to track the i5/OS objects associated with the integrated server. This name is used as the prefix for the system and install drives, and the Point-to-Point Virtual Ethernet line description.

#### **Resource name**

For installs using an iSCSI HBA, enter **\*ISCSI**.

For installs on an IXS or IXA, enter in the associated hardware resource name (**LINxx**).

#### **Linux server distribution**

Select the appropriate distribution. At the time of writing, the valid options are \*RHEL3, \*RHEL4, \*SLES8 or \*SLES9. This might change as newer Linux distributions become available. **Note:** RHEL 3 is the only Linux distribution supported on the IXS.

- b. You might want to change the default values for the following optional parameters:

#### **Server storage space sizes**

**Install source size:** The default value is calculated based on the space required to hold the installation source data. Typically, a storage space of 200 megabytes is created for the install source disk. This virtual disk can be between 200 MB and 2047 MB. It does not need to be larger than the calculated value and it must be kept as \*FAT format. The Linux integration code installs some utilities and configuration files here for its own use. The install drive is not intended for public use.

**System size:** The default value is calculated based on the space required to hold the Linux installation. Typically, a storage space of 12000 megabytes is created for the system disk. The size of the system storage space can be anywhere between 1024 MB and 1024000 MB. Depending on how much data you want to store on this disk, you might need to make it considerably larger. The system calculates a default value based on the following information:

## Values used by the INSLNXSVR CL command to estimate the system size

### Linux system partition

The estimate for the root and boot (if created) partitions with room for updates is 8 GB.

### Swap partition

This needs to be at least as large as the amount of installed memory. We assume 4 GB.

Be careful about specifying a system drive size less than the calculated value. While this might work initially, over time as you install more distribution updates you might run out of space.

The amount of space you allocate for your storage spaces does not affect the performance of the Linux operating system when it is running on integrated server hardware. For Linux servers running on integrated server hardware, the size of the system drive managed by i5/OS has no effect on performance. The storage spaces for the Linux operating system are actually files residing on the iSeries disk storage, and the data is automatically spread across all physical disk drives in the specified auxiliary storage pool (ASP).

### Restricted device resources

You can use this parameter to list iSeries tape and optical devices that you do not want to make available to the Linux server as virtual devices.

For IXS or IXA solutions, the installation of RHEL3 requires at least one optical drive that is not restricted.

For iSCSI attached solutions, you must restrict any tape or optical device that is not available, such as:

- External devices that are not powered on
- Any device that has been physically removed but still has a device description

If you do not restrict these devices, other devices (including disk, optical, and tape devices) might fail to report on your Linux server.

8. If you are installing an iSCSI attached server, do the following steps:
  - a. Enter information for the following parameters:

#### Storage Path

This parameter specifies the storage path the storage spaces can use. The **Network server host adapter** field must have the name of the existing network server host adapter device.

#### Virtual Ethernet Path

This parameter has two fields that describe the virtual Ethernet point to point connection. The **Port** field is filled in with \*VRTETHPTP by default. The **Network server host adapter** field must have the name of the existing network server host adapter device.

- b. Use one of the following methods to enter information for the network server configuration (NWSCFG) objects that i5/OS will use to manage the remote system.
  - The iSCSI install readme first Web page instructs you to create the network server configuration (NWSCFG) objects for **remote system**, **service processor**, and **connection security** prior to executing the install command. You need to specify the names of these objects with the following parameters:

#### Remote system NWSCFG

Specifies the remote system network server configuration to use with this server.

### **Service processor NWSCFG**

Specifies the service processor network server configuration to use with this server.

### **Connection security NWSCFG**

Specifies the connection security network server configuration to use with this server.

### **IP security rule**

You must specify \*NONE. The IP security rule parameter is not currently supported.

- If you do not specify values for the remote system, service processor, and connection security network server configuration (NWSCFG) objects, the INSLNXSVR command will create them for you. It is recommended that you verify your network connections by creating these objects before running the command.

If you have not already created these objects, enter information for the following parameters:

### **Enable unicast**

This value should be \*YES. This will enable the unicast packet distribution method for transmission of packets to the specified service processor host name or internet address.

### **Service processor name or SP internet address**

You must specify the either the internet host name or address of the service processor for the remote system.

### **Initialize service processor**

Specifies how the remote system's service processor is secured. This value should be \*SYNC since a new service processor object will be created and needs to be synchronized with the remote system's service processor.

### **SP authentication**

Specifies the service processor user ID and password. You must specify the user ID and password since the Initialize service processor parameter will be set to the value \*SYNC .

### **Remote system identifier**

Specifies the identifying serial number, type and model of the remote system. You must specify these values for a blade system. You can use the default value of \*EID for an xSeries product can be used.

### **Remote interfaces**

This parameter is used to identify and configure the remote system's SCSI and LAN interfaces. Both the SCSI and LAN have fields for **Adapter address**, **Internet address**, **Subnet mask** , **Gateway address**. The SCSI interface requires a value for **iSCSI qualified name**.

If you configure the iSCSI initiator to use dynamic addressing on boot (as recommended) and the install command has use the default setting of \*DYNAMIC for the **Delivery** parameter, these values are provided to the iSCSI HBA during the boot process. If you use the default value of \*GEN the Install Linux Server (INSLNXSVR) command will automatically generate a value if it is not specified. This default is designed to work with dynamic boot option.

When the iSCSI initiator is configured for manual addressing on boot and the install command has the **Delivery** parameter set to \*MANUAL, these values must match

- what you have configured on the iSCSI HBA. In this case, you must specify a value other than \*GEN in the iSCSI qualified name field.
9. Press **Enter**. When the command runs, the following things occur:
    - The NWSD is created.
    - The system storage space is created. This might take a few minutes depending on the size.
    - The install storage space is created.
    - Files are copied to the install drive.
  10. If you are installing Linux (other than RHEL 3) on an externally attached BladeCenter or xSeries you will be prompted to input the installation media in the xSeries optical drive.
 

For example, a SLES 9 install on an IXA attached xSeries would prompt:

Please insert SLES9 disc 1 into the xSeries CD-ROM (C G).

Verify that the installation media is inserted in the xSeries CD-ROM drive. Type G then press **Enter**. (Do not insert the diskette at this time.)

As another example, a SLES 9 install on an iSCSI attached BladeCenter or xSeries would prompt:

Please insert SLES9 disc 1 into the xSeries CD-ROM.

The vary processing has already started by the time this message is displayed. The blade or xSeries system will be automatically powered on and starting the boot process.
  11. The configuration files are updated and then the installation is transferred to the Linux server's console. When this happens you see the following message at the bottom of the green screen session from which you ran the INSLNXSVR command:
 

Network server install completed for <network server name>.

The server boots after the i5/OS portion of the installation has completed. The Linux installation wizard is displayed on the server's console. You must now complete the installation from the Linux console.

## INSLNXSVR command parameter descriptions

This topic describes the command parameters associated with the Install Linux server (INSLNXSVR) CL command.

### Parameters for all integrated Linux servers

#### Network server description (NWSD)

The network server description is created using the values specified in this command. This name is also used as the prefix for the system and install drives, as well as the Point-to-Point Virtual Ethernet line description. By default, this is the name used as the computer name of the Linux server that is installed as well as the TCP host name for the Linux server.

This is a required parameter.

#### *name*

Specify the name of the network server. The network server name can be up to eight characters. The following characters are allowed in NWSD names:

- Alphabetical characters A through Z
- Digits 0 through 9

#### Resource name (RSRCNAME)

Specifies the resource name that identifies the hardware that the network server description uses. Use the Work with Hardware Resources (WRKHDWRSC) command with \*CMN specified for the TYPE parameter to help determine the resource name. Specify the resource name of the communications adapter. The resource name represents the File Server IOA.

This is a required parameter.

*name*

Specify the resource name. The maximum length of the resource name is 10 characters.

### **Linux server distribution (LNXSVRDST)**

Specifies the version of Linux server to install on this network server. At the time of writing, the valid options are \*RHEL3, \*RHEL4, \*SLES8 or \*SLES9. This might change as newer Linux distributions become available.

This is a required parameter.

### **character-value**

Specify the Linux server distribution version to install. You can use F4 while prompting the command to see the complete list of allowed values for this parameter.

### **Linux source directory (LNXSRCDIR)**

This parameter is only used to install RHEL 3. This parameter is not used with RHEL 4, SLES 8 or SLES 9 because these versions install from the media located in the optical drive of an IXA-attached xSeries server.

Specifies the directory associated with the Linux server CD-ROM image that is used as the source for the install. The directory name may reference an optical volume ('/QOPT/volume'), a folder ('/QDLS/folder'), or an integrated file system directory ('/dir1/dir2'). An example of a CD-ROM volume path name would be '/QOPT/RHEL\_3 I386/'.

To find out the name of a volume on an optical device, use the command: DSPOPT VOL(\*MOUNTED) DEV(device-name). If you do not know the name of the optical device, use the command: WRKCFGSTS CFGTYPE(\*DEV) CFGD(\*OPT)

To find out the name of a path in an integrated file system directory, use the Work with Object Links (WRKLNK) command. WRKLNK will show the directory object path names on the system.

### **\*DFT**

The default directory path name to use is determined by searching the QOPT file system (/QOPT directory in integrated file system). The path name for the first volume in QOPT that contains a DOSUTILS directory with a LOADLIN.EXE file in it is used as the default.

**Note:** Using \*DFT on a system with an optical media library is not recommended. Every volume located in the '/QOPT' path will be searched for a valid Linux install source and will cause each volume to be retrieved and mounted. A specific optical volume path should be specified.

### **path-name**

Specify the Linux server source directory path name to use for the install.

**Note:** The specified directory must contain a DOSUTILS directory with a LOADLIN.EXE file in it.

You can use F4 while prompting the command to see a list of path names for optical volumes that are allowed values for this parameter. Folder or integrated file system path names that are not optical volumes are also allowed, but are not listed when using F4.

### **TCP/IP port configuration (TCPPORTCFG)**

Specifies the Linux TCP/IP configuration values that are specific to a port on the network server. This information consists of four parts including the identification of the network server port, the internet address, the subnet mask and the default gateway assigned to the port.

### **Virtual Ethernet port (VRTETHPORT)**

Specify the TCP/IP configuration for the virtual Ethernet used by the file server.

#### **Single values**

##### **\*NONE**

Specifies that there is no Linux TCP/IP port configuration.

#### **Other values (up to 4 repetitions)**

##### **Element 1: Port**

Specifies the network server virtual Ethernet port number to be configured. Specify one of the following values:

##### **\*VRTETHn**

The network server virtual Ethernet port 'n' is configured, where 'n' can have a value of 0 through 9.

##### **Element 2: Linux internet address**

###### **character-value**

Specify the Linux internet address for the port in the form, nnn.nnn.nnn.nnn, where nnn is a are decimal numbers ranging from 0 through 255. The internet address selected must be unique across all NWSD objects and the system's TCP/IP configuration.

##### **Element 3: Linux subnet mask**

###### **character-value**

Specifies the subnet mask for the Linux internet address in the form, nnn.nnn.nnn.nnn, where nnn is a decimal number ranging from 0 through 255.

##### **Element 4: Associated port**

Specifies the resource name that describes the port that is used to establish a connection between a Linux network server and the network.

**Note:** Use the Work with Hardware Resources (WRKHDWRSC) command with \*CMN specified for the TYPE parameter to help determine the resource name. The resource name is on the port. For example, the resource name may be CMN01 on a Ethernet port.



### \*NONE

An associated port resource name is not associated with the line.

### *name*

Specify the resource name.

### TCP/IP local domain name (TCPDMNNAME)

Specifies the local domain name associated with the network server.

A domain name can be a text string having 2 to 255 characters. Domain names consist of one or more labels separated by periods. Each label can contain up to 63 characters. The following characters are allowed in domain names:

- Alphabetical characters A through Z
- Alphabetical characters a through z
- Digits 0 through 9
- Minus sign (-)
- Period (.). Periods are only allowed when they separate labels of domain style name (refer to RFC 1034).

Uppercase and lowercase characters are allowed, but no significance is attached to the case. The case is maintained as entered. The first and last character of the host name must be an alphabetic character or a digit.

### \*SYS

Specifies that the local domain name for the network server should be the same value as is configured for the system.

### *character-value*

Specify a TCP domain name to be associated with the network server.

### TCP/IP name server system (TCPNAMSVR)

Specifies the internet address of the name server system that is used by the network server. Typically, this is the same value as it is for the system.

### **Single values**

### \*SYS

The name server system used by the network server should be the same as for the system.

### **\*NONE**

There is no name server to be used by the network server.

### **Other values (up to 3 repetitions)**



### *character-value*

Specify an internet address for the name server system to be used by the network server. Up to three remote name server systems can be specified. The name server systems are used in the order they are specified.

### **Server storage space sizes (SVRSTGSIZE)**

Specifies the size of the server storage spaces, in megabytes.

#### **Element 1: Install source size**

Specifies the size of the storage space that holds the files used to install the Linux server.

#### **\*CALC**

Specifies that the size should be calculated based on the space required to hold the install source. Typically, a storage space of 200 megabytes is created for the install source disk.

#### **integer**

Specifies the install source size value in megabytes. The size must be at least 200 megabytes but no larger than 2047 megabytes.

#### **Note:**

The Integrated Server Support code is copied to the install source drive. The size specified for the install source drive must be large enough to hold this data.

#### **Element 2: System size**

Specifies the size, in megabytes, of the storage space that the Linux server operating system is installed on.

#### **\*CALC**

Specifies that the size should be calculated based on the space required to hold the Linux installation. Typically, a storage space of 12000 megabytes is created for the system disk.

#### **integer**

Specifies the size of the system disk value in megabytes. The size must be at least 1024 megabytes but no larger than 1024000 megabytes.

**Note:** Be careful about specifying a system drive size less than the calculated value. While this might work initially, over time as you install more distribution updates you might run out of space.

### **Storage space ASP (SVRSTGASP)**

Specifies the auxiliary storage pool (ASP) identifiers for the storage space that will contain the files used to install the Linux server and the storage space that will contain the Linux server operating system.

**Note:** You cannot specify values for both the SVRSTGASP and STGASPDEV parameters.

**Note:** Null (omitted) values are specified with the characters \*N, which mean that no value was specified. The value specified for the corresponding Server storage ASP device (STGASPDEV)

element will be used if specified. Otherwise, the default value will be used. \*N is needed only when another value following the omitted element is being specified.

#### **Element 1: Install source ASP**

Specifies the auxiliary storage pool for the storage space that holds the files used to install the Linux server.

1

The storage space is created in auxiliary storage pool 1, the system auxiliary storage pool.

*1-255*

Specify the ASP identifier of the ASP to hold install files. Valid values depend on how many ASPs are defined on the system.

#### **Element 2: System ASP**

Specifies the auxiliary storage pool for the storage space that holds the Linux server operating system.

1

The storage space is created in auxiliary storage pool 1, the system auxiliary storage pool.

*1-255*

Specify the ASP identifier of the ASP to hold the operating system. Valid values depend on how many ASPs are defined on the system.

#### **Server storage ASP device (STGASPDEV)**

Specifies the auxiliary storage pool (ASP) device for the storage space that will contain the files used to install the Linux server and the storage space that will contain the Linux server operating system.

**Note:** You cannot specify values for both the SVRSTGASP and STGASPDEV parameters.

The ASP must have been activated (by varying on the ASP device) and have a status of 'Available'.

#### **Element 1: Install source ASP device**

Specifies the independent auxiliary storage pool device for the storage space that holds the files used to install the Linux server.

*name*

Specify the device name of the ASP to use for the network server storage space.

#### **Element 2: System ASP device**

Specifies the independent auxiliary storage pool device name for the storage space that holds the Linux server operating system.

*name*

Specify the device name of the ASP to use for the network server storage space.

### **Language version (LNGVER)**

Specifies the installed language environment used to display Integrated Server Support text and messages. Note that a smaller set of languages are available for messages displayed during installation than for messages displayed after installation.

#### **\*PRIMARY**

The installed language environment for Integrated Server Support text and messages is based on the language feature of the system's primary language.

*integer*

### **Synchronize date and time (SYNCTIME)**

Specifies whether the system should synchronize the network server date and time with the host system date and time.

The QTIMZON system value must be set to the correct value for time synchronization to work correctly.

#### **\*YES**

The system synchronizes the file server date and time with the host system date and time at every vary on and at least every 30 minutes thereafter

#### **\*NO**

The system synchronizes the file server date and time with the host system date and time when the network server description is varied on, but will not keep the date and time synchronized while the network server description is in a varied on state.

### **Restricted device resources (RSTDDEVRSC)**

Specifies the system's device resource names that are restricted and cannot be used by the Linux server. The installation of RHEL 3 requires that at least one optical drive is not restricted. This is necessary because RHEL 3 will use that optical drive to install the OS.

**Note:** Only tape and optical device resources can be restricted.

#### **Single values**

#### **\*NONE**

No device resources are restricted from the network server. Therefore, any tape or optical device resources that exist on the system can be used.

#### **\*ALL**

All tape and optical resources are restricted from being used by the network server.

#### **Other values (up to 10 repetitions)**

#### **\*ALLOPT**

All optical resources are restricted from being used by the network server.

**Note:** This value can only be specified once.

#### **\*ALLTAPE**

All tape resources are restricted from being used by the network server.

**Note:** This value can only be specified once.

*name*

Specify the name of the restricted device resource that cannot be used by the network server.

#### **Shutdown timeout (SHUTDTIMO)**

Specifies the server shutdown time-out value, in minutes. This is used to limit the amount of time that the servers operating system is allowed to shutdown before the server is varied offline.

15

The network server default shutdown time-out value is used.

*2-45*

Specify the time (in minutes) to wait. The system waits until the network servers operating system has shutdown successfully, or until the specified time passes before varying the network server offline.

### **| Additional parameters for iSCSI attached servers**

#### **| Activation timer (ACTTMR)**

| Specifies the amount of time (in seconds) the system will wait for the connection to be established to the remote server's service processor and to power on the remote server.

| **Note: This parameter is only valid when \*ISCSI is specified for the Resource name (RSRCNAME) parameter.**

| The network server default shutdown time-out value is used.

| 120

| The activate time of 120 seconds is used.

| *30-1800*

| Specify, in seconds, a value ranging from 30 through 1800.

#### **| Communications message queue (CMNMSGQ)**

| Specifies the name of a message queue to receive communications status messages.

| **Note: This parameter is only valid when \*ISCSI is specified for the Resource name (RSRCNAME) parameter.**

#### **| Single Values**

#### **| \*SYSOPR**

| Causes messages to be placed in the system operator message queue (QSYSOPR message queue in library QSYS).

| **Qualifier 1: Communications message queue**

| *name*

| Specify the name of a message queue to receive communications status messages.

| **Qualifier 2: Library**

| \*LIBL

| All libraries in the job's library list are searched until the first match is found.

| \*CURLIB

| The current library for the job is searched. If no library is specified as the current library for the job, the QGPL library is used.

| *name*

| Specify the name of the library to be used.

| Storage path (STGPTH)

| Specifies the storage path the storage spaces can use.

| **Note: This parameter is only valid when \*ISCSI is specified for the Resource name (RSRCNAME) parameter.**

| **Element 1: Network server host adapter**

| *communications-name*

| Specify the name of an existing network server host adapter (NWSH) device.

| Virtual Ethernet path (VRTETHPTH)

| Specifies the virtual Ethernet paths the Ethernet line descriptions can use. This information consists of two parts including the virtual Ethernet port and the network server host adapter (NWSH) device. You must enter at least one virtual Ethernet path which is the path to be used by the \*VRTETHPTP line description.

| **Note: This parameter is only valid when \*ISCSI is specified for the Resource name (RSRCNAME) parameter.**

| You can specify 5 values for this parameter.

| **Element 1: Port**

| \*VRTETHPTP

| The network server virtual Ethernet point to point port is configured.

| \*VRTETH0

| Virtual Ethernet port 0 is configured.

| **\*VRTETH1**

| Virtual Ethernet port 1 is configured.

| **\*VRTETH2**

| Virtual Ethernet port 2 is configured.

| **\*VRTETH3**

| Virtual Ethernet port 3 is configured.

| **\*VRTETH4**

| Virtual Ethernet port 4 is configured.

| **\*VRTETH5**

| Virtual Ethernet port 5 is configured.

| **\*VRTETH6**

| Virtual Ethernet port 6 is configured.

| **\*VRTETH7**

| Virtual Ethernet port 7 is configured.

| **\*VRTETH8**

| Virtual Ethernet port 8 is configured.

| **\*VRTETH9**

| Virtual Ethernet port 9 is configured.

| **Element 2: Network server host adapter**

| *communications-name*

| Specify the name of an existing network server host adapter (NWSH) device. The network server host adapter name does not need to be unique for each VRTETHPTH parameter on this NWSD.

| **Shutdown TCP port (SHUTDPORT)**

| Specifies the TCP port to use for shutdown.

| **Note:** This parameter is only valid when \*ISCSI is specified for the Resource name (**RSRCNAME**) parameter.

| **8700**

| Use the TCP port number of 8700.

| *integer*

| Specify the port number identifying the port that is to be used for shutdown. Valid values range from 1024 through 65,535.

| **Virtual Ethernet control port (VRTETHCTLP)**

| Specifies the TCP port to use for virtual Ethernet control.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)** parameter.

| **8800**

| Use the TCP port number of 8800.

| *integer*

| Specifies the port number identifying the port that is to be used for virtual Ethernet control. Valid values range from 1024 through 65,535.

| **Remote system NWSCFG (RMTNWSCFG)**

| Specifies the remote system network server configuration to use with this server.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)** parameter.

| **\*DFT**

| Use the system generated default remote system network server configuration name of 'nwsdnameRM' where nwsdname is the name of the network server description.

| *name*

| Specify the name of an existing remote system network server configuration.

| **Service processor NWSCFG (SPNWSCFG)**

| Specifies the service processor network server configuration to use with this server.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)** parameter.

| **\*DFT**

| Use the system generated default service processor network server configuration name of 'nwsdnameSP' where nwsdname is the name of the network server description.

| *name*

| Specify the name of an existing service processor network server configuration.

| **Connection security NWSCFG (CNNWSCFG)**

| Specifies the connection security network server configuration to use with this server.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)** parameter.

| **\*DFT**

| Use the system generated default connection security network server configuration name of 'nwsdnameCN' where nwsdname is the name of the network server description.

| ***name***

| Specify the name of an existing connection security network server configuration.

| **Default IP security rule (DFTSECRULE)**

| Specifies the default IP Security (IPSec) rule used between the hosting and remote system.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)** parameter.

| **Note:** This parameter is ignored when the connection security network server configuration specified for the **Connection security NWSCFG (CNNWSCFG)** parameter already exists.

| A value other than \*DFTSECRULE must be specified on the IP security rule (IPSECRULE) parameter in this case.

| **\*NONE**

| IP Security (IPSec) protocol security settings are not configured.

| ***\*GEN***

| Generate a random pre-shared key.

| ***character-string***

| Specify the pre-shared key.

| A pre-shared key is a nontrivial string up to 32 characters long.

| Valid characters are upper case A through Z, lower case a through z, numbers 0 through 9, and the following special characters:

| • Plus sign

| • Equal sign

| • Percent

| • Ampersand

| • Left parenthesis

| • Right parenthesis

| • Comma



| • Underline

| • Minus sign

| • Period

| • Colon

| • Semicolon

| **IP security rule (IPSECRULE)**

| Specify the relative entry of an existing IP security rules (IPSECRULE) parameter, defined in the existing connection security network server configuration that will be used as the initial IP security setting between the hosting and remote system.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)** parameter.

| **\*DFTSECRULE**

| Use the value specified on the **Default IP security rule (DFTSECRULE)** parameter.

| **Note:** This value is not valid when the **Connection security NWSCFG** specified on the CNNWSCFG parameter already exists.

| **\*NONE**

| Remote interface will not use any security rule.

| **1-16**

| Remote interface will use security rule specified.

| **Initialize service processor (INZSP)**

| Specifies how the remote system's service processor is secured.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)** parameter.

| Use this only if the interconnecting network is physically secure.

| Some service processors do not support secure connections. Use \*NONE for these service processors. Additional information can be found at **System i integration with BladeCenter and System x** at <http://www.ibm.com/systems/i/bladecenter/>.

| **\*MANUAL**

| Security parameters are manually configured on remote system's service processor. \*MANUAL provides the highest security.

| To use this option, it is required that the remote system's service processor is preconfigured with a user name, password and certificate. Certificate management will be required. This method is appropriate when connecting to the service processor via public networks to protect the password.

| **Note:** This parameter is ignored when the service processor network server configuration specified for the  
| **Service processor NWSCFG (SPNWSCFG)** parameter already exists.

| **\*AUTO**

| Parameters are automatically configured on the remote system's service processor.

| \*AUTO provides security without requiring preconfiguration of the remote system's service processor.  
| The remote system's service processor will have certificates automatically regenerated when the  
| certificates are near expiration. This option is appropriate if the interconnecting network is physically  
| secure or is protected by a firewall.

| **Note:** An administrator will need to regenerate the certificate using the Initialize NWS Configuration  
| (INZNWSCFG) command when the service processor certificate has expired, or if a new certificate and  
| password are desired at any time before the certificate expires.

| **Note:** This parameter is ignored when the service processor network server configuration specified for the  
| **Service processor NWSCFG (SPNWSCFG)** parameter already exists.

| **\*SYNC**

| Synchronize the self-signed certificate from the service processor. This option is used if multiple service  
| processor network server configurations are used for the same system or the service processor network  
| server configuration has been restored from backup and the service processor's certificate must be  
| synchronized. The current user name and password for the service processor must be specified on the **SP**  
| **authentication (SPAUT)** parameter to perform this option.

| **\*NONE**

| Provides no security.

| Use this only if the interconnecting network is physically secure.

| **Note:** Some service processors do not support secure connections. Use \*NONE for these service  
| processors. Additional information can be found at **System i integration with BladeCenter and System x**  
| at <http://www.ibm.com/systems/i/bladecenter/>.

| **Note:** This parameter is ignored when the service processor network server configuration specified for the  
| **Service processor NWSCFG (SPNWSCFG)** parameter already exists.

| **Enable unicast (ENBUNICAST)**

| Specifies whether unicast packet distribution is to be used. Unicast is a transmission method where  
| packets are sent directly to the specified **Service processor name (SPNAME)** or **SP internet address**  
| (**SPINTNETA**) parameter.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)**  
| parameter.

| **Note:** This parameter is ignored when the service processor network server configuration specified for the  
| **Service processor NWSCFG (SPNWSCFG)** parameter already exists.

| **\*NO**

| Disable unicast.

| **\*YES**

| Enable unicast.

| **Enclosure identifier (EID)**

| Specifies the identifying serial number, type and model of the enclosure containing the service processor.

| When specified, they are used to locate the system on the network.

| Look for these values on the label of the system.

| **Note:** This parameter is only valid when **\*ISCSI** is specified for the **Resource name (RSRCNAME)** parameter.

| **Note:** This parameter is ignored when the service processor network server configuration specified for the **Service processor NWSCFG (SPNWSCFG)** parameter already exists.

| **Single Values**

| **\*AUTO**

| Automatically retrieve the identifier when ENBUNICAST(\*YES) is specified.

| **Element 1: Serial number**

| *character-string*

| Specify the machine serial number.

| **Element 2: Manufacturer type and model**

| *character-string*

| Specify the machine type and model.

| The value is entered in the form ttttmmm where tttt is the machine type and mmm is the machine model number.

| **Service processor name (SPNAME)**

| Specifies the remote system's service processor host name.

| **Note:** This parameter is required when ENBUNICAST(\*YES) is specified.

| **Note:** This parameter is only valid when **\*ISCSI** is specified for the **Resource name (RSRCNAME)** parameter.

| **Note:** This parameter is ignored when the service processor network server configuration specified for the **Service processor NWSCFG (SPNWSCFG)** parameter already exists.

| **\*SPINTNETA**

| The remote system is identified by the value specified for the **SP internet address (SPINTNETA)** parameter.

| *host-name*

| Specify the remote system's service processor host name.

| **SP internet address (SPINTNETA)**

| Specifies the remote system's service processor internet address.

| The value is entered in the decimal form nnn.nnn.nnn.nnn, where nnn is a decimal number ranging from 0 through 255.

| **Note:**This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)** parameter.

| **Note:** This parameter is ignored when the service processor network server configuration specified for the **Service processor NWSCFG (SPNWSCFG)** parameter already exists.

| **Notes:**

| 1. This parameter is ignored when ENBUNICAST(\*NO) is specified.

| 2. This parameter is required when SPNAME(\*SPINTNETA) is specified.

| *internet-address*

| Specify the internet address of the service processor.

| The value is entered in the decimal form nnn.nnn.nnn.nnn, where nnn is a decimal number ranging from 0 through 255.

| **SP authentication (SPAUT)**

| Specifies the service processor user name and password. This is used to authenticate and secure the service processor.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)** parameter.

| **Note:** This parameter is ignored when the service processor network server configuration specified for the **Service processor NWSCFG (SPNWSCFG)** parameter already exists.

| **Notes** This value is only valid when INZSP(\*NONE) is specified on this command or in the existing network server configuration.

| **Single values**

| \*DFT

| The default service processors userid and password are used.

| **Element 1: User name**

| *character-string*

| Specify a name that represents the host configuration that owns the service processor. It is suggested that the remote system network server configuration name be used. If multiple remote system network server configurations use the same service processor at different times, each configuration must contain the same user name and password.

| **User password**

| *character-string*

| Specify the service processor password. Password must be at least 5 characters in length and contain at least one alphabetic character and one numeric or symbolic character.

| **SP certificate identifier (SPCERTID)**

| The SP certificate identifier specifies one of three possible fields that identifies the service processor's certificate.

| This parameter is specified to provide additional validation that the certificate is from the service processor. The contents of the selected field must exactly match the value of the field that was entered when the certificate was generated or requested from a certificate authority.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)** parameter.

| **Note:** This parameter is ignored when the service processor network server configuration specified for the **Service processor NWSCFG (SPNWSCFG)** parameter already exists.

| This parameter is required when INZSP(\*MANUAL) is specified and cannot have the value \*NONE.

| **Single values**

| **\*NONE**

| Service processor certificate is not configured.

| **Element 1: Component**

| **\*COMMONNAME**

| Selects the certificate's common name specified when the certificate was generated or requested from a certificate authority. On the remote supervisor adapter II this correlates to the "ASM Domain Name" field used to generate a self-signed certificate or generate a certificate signing request.

| **\*EMAIL**

| Selects the certificate's e-mail address specified when the certificate was generated or requested from a certificate authority. On the remote supervisor adapter II this correlates to the "Email Address" field used to generate a self-signed certificate or generate a certificate signing request.

| **\*ORGUNIT**

| Selects the certificate's organizational unit specified when the certificate was generated or requested from a certificate authority. On the remote supervisor adapter II this correlates to the "Organizational Unit" field used to generate a self-signed certificate or generate a certificate signing request.

| **Element 2: Compare value**

| *character-string*

| Specify the certificates component compare value. Enter no more than 255 characters of text, enclosed in apostrophes.

| **Remote system identifier (RMTSYSID)**

| Specifies the identifying serial number, type and model of the remote system. When specified, they are used to locate the remote system on the network.

| Look for these values on the label of the system.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)** parameter.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Remote system NWSCFG (NWSCFG)** parameter already exists.

| **Single values**

| **\*EID**

| Use the service processor identifier.

| **Element 1: Serial number**

| *character-string*

| Specify the machine serial number.

| **Element 2: Manufacturer type and model**

| *character-string*

| Specify the machine type and model.

| **Delivery method (DELIVERY)**

| Specifies how the parameters necessary to configure the remote system are delivered.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)** parameter.

| **Note:** This parameter is ignored when the remote system network server configuration specified for the **Remote system NWSCFG (NWSCFG)** parameter already exists.

| **\*DYNAMIC**

| Parameters are dynamically delivered to the remote system using Dynamic Host Configuration Protocol (DHCP).

| **\*MANUAL**

| Parameters are manually configured on the remote system using the BIOS utilities (System BIOS or Adapter BIOS - CTRL-Q).

## | CHAP authentication (CHAPAUT)

| Specifies the Challenge Handshake Authentication Protocol (CHAP) for the host system iSCSI target to authenticate the remote system initiator node.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)** parameter.

| **Note:** This parameter is ignored when the remote system network server configuration specified for the **Remote system NWSCFG (NWSCFG)** parameter already exists.

### | **Single values**

#### | \*NONE

| CHAP authentication is not enabled.

#### | **Element 1: CHAP name**

#### | \*NWSCFG

| The system will automatically generate a name for CHAP using the Network server configuration name.

#### | *character-string*

| Specify the name you want to use for the Challenge Handshake Authentication Protocol.

| Valid characters are upper case A through Z, lower case a through z, numbers 0 through 9, and the following special characters:

- | • Plus sign
- | • Equal sign
- | • Percent
- | • Ampersand
- | • Left parenthesis
- | • Right parenthesis
- | • Comma
- | • Underline
- | • Minus sign
- | • Period
- | • Colon
- | • Semicolon

#### | **Element 2: CHAP secret**

| \*GEN

| The system will automatically generate a random CHAP secret.

| *character-string*

| Specify the secret you want to use for the Challenge Handshake Authentication Protocol.

| Valid characters are upper case A through Z, lower case a through z, numbers 0 through 9, and the following special characters:

| • Plus sign

| • Equal sign

| • Percent

| • Ampersand

| • Left parenthesis

| • Right parenthesis

| • Comma

| • Underline

| • Minus sign

| • Period

| • Colon

| • Semicolon

| Boot device ID (BOOTDEVID)

| Specifies the PCI Function Address (Bus/Device/Function) of the iSCSI adapter in the remote system that will be used to boot from.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)** parameter.

| **Note:** This parameter is ignored when the remote system network server configuration specified for the **Remote system NWSCFG (NWSCFG)** parameter already exists.

| **Single values**

| \*SINGLE

| The single iSCSI adapter is used on the remote system.

| **Element 1: Bus**

| *number*



| Specify the bus number of the remote system's iSCSI adapter that will be used to boot.

| Valid values range from 0 through 255.

| **Element 2: Device**

| *number*

| Specify the bus number of the remote system's iSCSI adapter that will be used to boot.

| Valid values range from 0 through 31.

| **Element 3: Function**

| *number*

| Specify the bus number of the remote system's iSCSI adapter that will be used to boot.

| Valid values range from 0 through 7.

| **Dynamic boot options (DYNBOOTOPT)**

| Specifies the internal Dynamic Host Configuration Protocol (DHCP) Server configuration.

| **Note:** This is an advanced configuration function.

| This parameter is used to configure the internal DHCP Server that is part of the iSCSI Target Host Bus Adapter firmware. It is used to provide IP address and diskless boot parameters for the remote iSCSI Initiator.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)** parameter.

| **Note:** This parameter is ignored when the remote system network server configuration specified for the **Remote system NWSCFG (NWSCFG)** parameter already exists.

| This parameter is only valid when DELIVERY(\*DYNAMIC) is specified.

| **Element 1: Vendor ID**

| The client and server are preconfigured to a default vendor ID. Network administrators can configure clients to define their own identifying values to convey hardware, operating system or other identifying information. DHCP option 60 described in the IETF RFC 2132 is used for this function.

| **\*DFT**

| The default Client ID consists of the adapter address for the remote system's iSCSI adapter. This value will be used to identify the remote system.

| *character-string*

| Specify the Client ID of the remote system's iSCSI adapter that will be used to boot.

| **Element 2: Alternate client ID**

| Used by clients to specify their unique identifier to the server. Each client's identifier must be unique among all other client identifiers used on the effective DHCP network to which the client is attached (that is, the client's local subnet and any remote subnets reachable using DHCP relay). Vendors and system administrators are responsible for choosing client identifiers that meet this requirement for uniqueness. DHCP option 61 described in the IETF RFC 2132 is used for this function.

| **\*ADPT**

| The default Client ID consists of the adapter address for the remote system's iSCSI adapter. This value will be used to identify the remote system.

| *character-string*

| Specify the Client ID of the remote system's iSCSI adapter that will be used to boot.

| **Remote interfaces (RMTIFC)**

| Specifies the remote system's interfaces. This information is used to identify and configure the remote system's interfaces. Each adapter has two functions to support a SCSI and a LAN interface.

| **Note:** This parameter is only valid when \*ISCSI is specified for the **Resource name (RSRCNAME)** parameter.

| **Note:** This parameter is ignored when the remote system network server configuration specified for the **Remote system NWSCFG (NWSCFG)** parameter already exists.

| **Element 1: SCSI interface**

| Specifies the remote system's SCSI interfaces.

| **Element 1: Adapter Address**

| **\*NONE**

| No SCSI interface is configured for this adapter.

| *adapter-address*

| Specify the 12-character hexadecimal adapter address for the remote system's iSCSI interface.

| **Element 2: Internet address**

| *internet-address*

| Specify the internet address for the remote system's SCSI interface.

| The value is entered in the decimal form nnn.nnn.nnn.nnn, where nnn is a decimal number ranging from 0 through 255.

| **Element 3: Subnet mask**

| *subnet-mask*

| Specify the subnet mask for the remote system's SCSI interface.

| The value is entered in the decimal form nnn.nnn.nnn.nnn, where nnn is a decimal number ranging from 0 through 255.

| **Element 4: Gateway address**

| *gateway-address*

| Specify the gateway address for the remote system's SCSI interface.

| The value is entered in the decimal form nnn.nnn.nnn.nnn, where nnn is a decimal number ranging from 0 through 255.

| **Element 5: iSCSI qualified name**

| \*GEN

| The system will automatically generate the iSCSI qualified name.

| *name*

| Specify the iSCSI qualified name for the remote system's SCSI interface.

| The following characters are allowed in iSCSI qualified names:

| • Alphabetical characters A through Z converted to lower case (refer to RFC 3722)

| • Alphabetical characters a through z

| • Digits 0 through 9

| • Period (.)

| • Dash (-)

| • colon (:)

| **Element 2: LAN interface**

| Specifies the remote system's LAN interfaces.

| **Element 1: Adapter Address**

| \*NONE

| No LAN interface is configured for this adapter.

| *adapter-address*

| Specify the 12-character hexadecimal adapter address for the remote system's LAN or TCP Offload Engine (TOE) interface.

| **Element 2: Internet address**

| *internet-address*

| Specify the internet address for the remote system's LAN interface.

| The value is entered in the decimal form nnn.nnn.nnn.nnn, where nnn is a decimal number ranging from  
| 0 through 255.

| **Element 3: Subnet mask**

| *subnet-mask*

| Specify the subnet mask for the remote system's LAN interface.

| The value is entered in the decimal form nnn.nnn.nnn.nnn, where nnn is a decimal number ranging from  
| 0 through 255.

| **Element 4: Gateway address**

| *gateway-address*

| Specify the gateway address for the remote system's LAN interface.

| The value is entered in the decimal form nnn.nnn.nnn.nnn, where nnn is a decimal number ranging from  
| 0 through 255.

| **Text 'description' (TEXT)**

| Specifies the text that briefly describes the objects created by this command.

| **\*BLANK**

| Text is not specified.

| *character-value*

| Specify no more than 50 characters of text, enclosed in apostrophes.

| **Configuration file (CFGFILE)**

| Specifies the source file that contains configuration data, to be used in activating or further defining the  
| server.

| **Single values**

| **\*NONE**

| No configuration file is specified.

| **Qualifier 1: Configuration file**

| *name*

| Specify the name of a source file containing the configuration data member(s) for the server. At the time  
| the server is activated, all members in the file will be processed.

| **Qualifier 2: Library**

| **\*LIBL**

| All libraries in the job's library list are searched until the first match is found.

| **\*CURLIB**

| The current library for the job is searched. If no library is specified as the current library for the job, the QGPL library is used.

| *name*

| Specify the name of the library to be used.

| **Virtual PTP Ethernet port (VRTPTPPORT)**

| Specifies the TCP/IP configuration for the virtual point-to-point Ethernet port.

| **Notes:**

- | 1. The VRTPTPPORT parameter must be used for Integrated xSeries servers running on the host system.
- | 2. The subnet mask that is used for both sides of the virtual point-to-point Ethernet port is 255.255.255.0 by default. Therefore, the internet addresses that are chosen for both sides of the virtual point-to-point Ethernet port must have the same values for the first three parts of the internet addresses.

| **Element 1: Internet address**

| Specifies the internet address for the host side of the virtual point-to-point Ethernet connection. Specify one of the following values:

| **\*GEN**

| Specify \*GEN to let the INSLNXSVR command configure a virtual point-to-point Ethernet port with a generated internet address.

| *character-value*

| Specify the host internet address for the virtual point-to-point Ethernet port in the form, xxx.yyy.zzz.nnn, where xxx, yyy, zzz, and nnn are decimal numbers ranging from 0 through 255. The internet address selected must be unique across all NWSD objects and the host TCP/IP configuration. The xxx.yyy.zzz value in the internet address must match the xxx.yyy.zzz value for the internet address of the Linux server side of the point-to-point Ethernet port.

| **Element 2: Linux internet address**

| Specifies the internet address for the Linux side of the virtual point-to-point Ethernet connection. Specify one of the following values:

| **\*GEN**

| Specify \*GEN to let the INSLNXSVR command configure a virtual point-to-point Ethernet port with a generated internet address.

| *character-value*

| Specify the Linux server internet address for the virtual point-to-point Ethernet port in the form, xxx.yyy.zzz.nnn, where xxx, yyy, zzz, and nnn are decimal numbers ranging from 0 through 255. The internet address selected must be unique across all NWSD objects and the host TCP/IP configuration. The xxx.yyy.zzz value in the internet address must match the xxx.yyy.zzz value for the internet address of the host side of the point-to-point Ethernet port.

## | Examples

### | Example 1: Install Linux Using Non-iSCSI Hardware

| INSLNXSVR NWSD(RHEL3SVR) RSRNAME(LIN03) LNXSVRDST(\*RHEL3)

| SVRSTGSIZE(200 12000) LNGVER(2920)

| TEXT('RedHat 3 server')

| This command will install a Linux server with a Red Hat 3 distribution on non-iSCSI hardware. It will contain a 12 gb system drive and exist on resource name LIN03. The name RHEL3SVR will be associated with the network server description. The language version used will be 2920.

### | Example 2: Install Linux Using iSCSI Hardware

| INSLNXSVR

| NWSD(SLES9SVR) RSRNAME(\*ISCSI) LNXSVRDST(\*SLES9)

| SVRSTGSIZE(200 12000) SVRSTGASP(2) LNGVER(2920)

| SHUTDTIMO(2) ACTTMR(300) ENBUNICAST(\*YES)

| SPINTNETA('9.5.2.32') SPAUT(JSMITH ())

| TEXT('SuS E9 server')

| This command will install a Linux server with a SUSE 9 distribution on iSCSI hardware. It will contain a 12 gb system drive which resides on user ASP 2. The name SLES9SVR will be associated with the network server description. The language version used will be 2920. The shutdown timeout will be set to 2 minutes. The system will wait 300 seconds for the connection to be made to the remote server's service processor. Unicast packet distribution will be used. The remote server's service processor internet address will be set to 9.5.2.32. service processor authentication will be made with username JSMITH.

## **Error messages**

### \*ESCAPE Messages

#### **NTA1007**

Network server &1 must be varied off.

#### **NTA1199**

Vary on or off of the Linux server not successful.

#### **NTA119A**

Linux server installation not successful.

#### **NTA1024**

Storage space assigned to server &1 missing, damaged or not valid.

#### **NTA119F**

Internal error occurred.

## **Examples: Running the Install Linux Server (INSLNXSVR) command**

You can customize the parameters that you use for the Install Linux Server (INSLNXSVR) command depending on the configuration of your integrated server.

### **Installing RHEL3 on an Integrated xSeries Server (IXS)**

Many of the parameters on the INSLNXSVR command are optional. Here is a sample command string for installing Linux on an IXS solution:

```
INSLNXSVR NWS(MYIXS) RSRNAME(LINXX) LNXSVRDST(*RHEL3)
```

### **Installing SLES9 on an IXA attached xSeries solution**

Many of the parameters on the INSLNXSVR command are optional. Here is a sample command string for installing Linux on an IXA solution:

```
INSLNXSVR NWS(MYIXA) RSRNAME(LINXX) LNXSVRDST(*SLES9)
```

### **Installing SLES9 on an iSCSI attached xSeries solution using existing configuration objects**

Fewer parameters are required for the INSLNXSVR command when you create the i5/OS configuration objects in advance. This example uses an existing remote system, connection security, and service processor network server configuration (NWSCFG) objects. The iSCSI initiator is configured for DHCP boot.

```
INSLNXSVR NWS(MYISCSI) RSRNAME(*ISCSI) LNXSVRDST(*SLES9) STGPTH(MYNWSH) VRTETHPTH((*VRTETHPTP MYNWSH)) RMTNWSCFG(MYRM) SPNWSCFG(MYSP) CNNNWSCFG(MYCN) IPSECRULE(*NONE)
```

### **Installing SLES9 or later on an iSCSI attached integrated server using the INSLNXSVR command to create configuration objects**

You can use the Install Linux Server (INSLNXSVR) command to create the configuration objects for your integrated server.

In this example, the INSLNXSVR command generates the remote system, service processor, and connection security network server configuration (NWSCFG) objects that are required for your integrated server. The iSCSI initiator is configured for DHCP boot.

```
INSLNXSVR NWS(MYISCSI) RSRNAME(*ISCSI) LNXSVRDST(*SLES9) STGPTH(MYNWSH) VRTETHPTH((*VRTETHPTP MYNWSH)) IPSECRULE(*NONE) INZSP(*SYNC) ENBUNICAST(*YES) SPINTNETA('9.5.X.X') SPAUT(MYSPUID MYSPPWD) RMTIFC((NNNNNNNNNNN '192.168.x.x' '255.X.X.X') (NNNNNNNNNNN '192.168.X.X' '255.X.X.X'))
```

### **Installing SLES9 on an iSCSI attached BladeCenter using the INSLNXSVR command to generate create objects**

You can use the Install Linux Server (INSLNXSVR) command to create the configuration objects for your integrated server.

In this example, the INSLNXSVR command generates the remote system, service processor, and connection security network server configuration (NWSCFG) objects that are required for your integrated server. The iSCSI initiator is configured for DHCP boot.

```


I INSLNXSVR NWS(MYISCSIB) RSRNAME(*ISCSI) LNXSVRST(*SLES9) STGPTH(MYNWSH) VRTETHPTH(((*VRTETHPTP
I MYNWSH)) IPSECRULE(*NONE) INZSP(*SYNC) ENBUNICAST(*YES) SPINTNETA('9.5.X.X') SPAUT(MYSPUID
I MYSPPWD) RMTSYSID(SERLNO TYPMOD) RMTIFC((NNNNNNNNNNNN '192.168.x.x' '255.X.X.X') (NNNNNNNNNNNN
I '192.168.X.X' '255.X.X.X'))


```

## Completing the installation from the Linux console

This topic introduces how to complete the integrated server installation with the Linux console. The task varies depending on the distribution and integrated server platform.

**Attention:** The Linux distributors may require special instructions for the installation on supported BladeCenter and xSeries models. You must follow these special instructions. Any special instructions can be found by checking the hardware certification list as described below.

- For Red Hat: <http://bugzilla.redhat.com/hwcert>  On this page you will see a **Quick Search** text field. Input your xSeries or BladeCenter model and click **Search**. You will see the associated Red Hat Enterprise Linux version listed. Click on your specific model to see details on the certification summary and details. In the details under Platform, **i386** means 32-bit, and **x86\_64** means 64-bit version of the operating system.

- For SUSE: <http://developer.novell.com/yesssearch/Search.jsp>  On this page input the following and click **Search**.

**Keywords:** Input your BladeCenter or xSeries model

**Company:** IBM

**Product:** SUSE Linux Enterprise Server 8 for x86 or

SUSE Linux Enterprise Server 9 for x86 or

SUSE Linux Enterprise Server 9 for AMD64 & Intel EM64T

The **x86** means 32-bit and the **AMD64 & Intel EM64T** means the 64-bit version of the operating system.

Your model should be listed. Click on the bulletin number to see details on configurations and requirements for install.

**Note: IXS/IXA** – All installations on the IXS or IXA hardware involve a driver diskette. Only input the driver diskette when prompted because it is not bootable. Otherwise you might see the message “602 Invalid diskette boot record”. To recover from this scenario, do the following: v Remove the diskette from the diskette drive. v Press Ctrl+Alt+Delete to reboot the server.

**Note: IXA** – Any time a restart of the Integrated xSeries Adapter occurs, the attached xSeries is rebooted twice, so you see the Power On Self Test (POST) sequence twice. This is normal.

**Note: iSCSI** – By the time the INSLNXSVR command prompts for the Linux media, the BladeCenter or xSeries will already be automatically powered on and going through the boot process. The installation requires the Linux media to be inserted into the BladeCenter or xSeries CD drive so that it can boot. If you did not enter the Linux media before the end of the Power On Self Test (POST), you will see a message display on the Linux console indicating that there is no boot device or a blank screen with a blinking cursor. At this point, you should enter the Linux media in the BladeCenter or xSeries CD drive and press CTRL-ALT-DEL at the Linux console. The system will reboot and the installation will start from the CD drive.

## Completing the installation of RHEL 3 on the IXS or IXA

This topic discusses the steps necessary to complete the Red Hat Enterprise Linux installation



## About this task

To complete the installation of a Red Hat Enterprise Linux 3 ES or AS (RHEL 3) server:

1. The server boots and dialog will display, **Do you have a driver disk?** Select **Yes** and press **Enter**.
2. If the installation detects multiple devices that could serve as sources for a driver disk, select **fd0** and press **Enter**.
3. You will be prompted to insert the driver diskette into the drive. After you insert the driver diskette into the drive, press **Enter**. There will be a message indicating that it is reading the driver diskette and then loading the `ibmixsvs` driver.
4. You will be prompted to load more driver disks. There is no need to do this. Select **No** and press **Enter**.
5. The Red Hat graphical interface will start. You can continue through the installation making selections as you prefer however specific selections are required on the **Disk Partitioning Setup** and **Boot Loader Configuration** screens.
6. When the installation displays the **Disk Partitioning Setup** screen, you can choose to automatically partition the disk or manually partition the disk using Disk Druid. Unless you are an experienced user, we recommend that you choose to automatically partition the disk.

**Important:** The install drive (`/dev/sdb`) *FAT* partition must not be modified. Otherwise the server will not boot.

To use automatic partitioning:

- a. Select **Automatically partition** and click **Next**.
  - b. Two disks are displayed, `/dev/sda` and `/dev/sdb`. Remove the check before `/dev/sdb` so that `/dev/sda` is the only drive checked and select **Remove all partitions on this system**. Click on **Next**, then click **Yes** on the warning message.
  - c. Review the partition setup on the **Disk Setup** display. Click **Next**.
7. When the installation displays the **Boot Loader Configuration** screen, you must make sure that the **GRUB** boot loader will be installed on `/dev/sda`. Click **Next**.
  8. Continue on with installation making selections as you prefer. When you are done making selections, the installation will prompt you for CD-ROMs. When this phase is finished, the server will reboot.
  9. After the server reboots, the Setup Agent will guide you through some basic configuration. You can continue through the setup as you prefer; however, specific selections are required on the **Red Hat Login** and **Why Register** screens. The reason is that Red Hat updates can not be applied to your server until the postinstallation instructions have been completed which are in the section titled *Completing the postinstallation tasks*.
  10. On the **Red Hat Login** screen, select **Tell me why I need to register and provide a Red Hat login** and click **Next**.
  11. On the **Why Register** screen, select **I can not complete registration at this time. Remind me later** and click **Next**.
  12. When the setup phases are complete, you will be prompted to sign on to the server. Please follow the instructions described in the section titled **Completing the postinstallation tasks** of the Redbook.
  13. Once the postinstallation instructions are complete, you can apply Red Hat updates to your server.  
**Important:** After applying Red Hat updates, you *must* run `ixsupdt` before you reboot your server. Details for `ixsupdt` are listed in the section titled *Maintaining the Linux integration code*.

## Completing the installation of RHEL 4 on the IXA

This topic discusses the steps necessary to complete the Red Hat Enterprise Linux installation

## About this task

To complete the installation of a Red Hat Enterprise Linux 4 ES or AS (RHEL 4) server:

1. The xSeries server boots from the Red Hat Enterprise Linux 4 disc 1 located in the xSeries CD-ROM drive. The Red Hat Enterprise Linux install prompt will appear.

**Important:** You must enter **linux dd** so that you will be prompted for a driver diskette. If you do not enter **linux dd** within 1 minute, the installation will continue without prompting you for a diskette and a warning message indicating that no hard drives have been found will appear. To recover, press CTRL-ALT-DELETE to reboot the xSeries. The RHEL 4 installation will begin again. You must enter **linux dd** at the install prompt within 1 minute.

2. A dialog will display, **Do you have a driver disk?** Select **Yes.** and press **Enter.**
3. If the installation detects multiple devices that could serve as sources for a driver disk, select **fd0** and press **Enter.**
4. You will be prompted to insert the driver diskette into the drive. After you insert the driver diskette into the drive, press **Enter.** There will be a message indicating that it is reading the driver diskette and then loading the **ibmixsvs** driver.
5. You will be prompted to load more driver disks. There is no need to do this. Select **No** and press **Enter.**
6. The Red Hat graphical interface will start. You can continue through the installation making selections as you prefer however specific selections are required on the **Disk Partitioning Setup** and **Boot Loader Configuration** screens.
7. When the installation displays the **Disk Partitioning Setup** screen, you can choose to automatically partition the disk or manually partition the disk using Disk Druid. Unless you are an experienced user, we recommend that you choose to automatically partition the disk.

**Important:** The install drive (*/dev/sdb*) FAT partition must not be modified. Otherwise the server will not boot.

**To use automatic partitioning:**

- a. Select **Automatically partition** and click **Next.**
  - b. Two disks are displayed, */dev/sda* and */dev/sdb*. Remove the check before */dev/sdb* so that */dev/sda* is the only drive checked and select **Remove all partitions on this system.** Click on **Next**, then click **Yes** on the warning message.
  - c. Review the partition setup on the **Disk Setup** display. Click **Next.**
8. When the installation displays the **Boot Loader Configuration** screen, you must make sure that the **GRUB** boot loader will be installed on */dev/sda*. Click **Next.**
  9. Continue on with installation making selections as you prefer. When you are done making selections, the installation will prompt you for CD-ROMs. When this phase is finished, the server will reboot.
  10. After the server reboots, the Setup Agent will guide you through some basic configuration. You can continue through the setup as you prefer; however, specific selections are required on the **Red Hat Login** and **Why Register** screens. The reason is that Red Hat updates can not be applied to your server until the postinstallation instructions have been completed which are in the section titled *Completing the postinstallation tasks.*
  11. On the **Red Hat Login** screen, select **Tell me why I need to register and provide a Red Hat login** and click **Next.**
  12. On the **Why Register** screen, select **I can not complete registration at this time, remind me later** and click **Next.**
  13. When the setup phases are complete, you will be prompted to sign on to the server. Please follow the instructions described in the section titled *Completing the postinstallation tasks.*
  14. Once the postinstallation instructions are complete, you can apply Red Hat updates to your server.  
**Important:** After applying Red Hat updates, you *must* run **ixsupdt** before you reboot your server. Details for **ixsupdt** are listed in the section titled *Maintaining the Linux integration code.*

## Completing the installation of SLES 8 on the IXA

This topic discusses the steps necessary to complete the SUSE Linux Enterprise Server installation.

## About this task

To complete the install of a SUSE Linux Enterprise Server 8 server:

1. The xSeries server boots from the SLES 8 disc 1 in the xSeries optical drive.

**Important:** As soon as you see the SUSE Linux ENTERPRISE SERVER GUI start on the Linux server's console, press the **Alt** key.

If you do not press the **Alt** key on the initial screen, you will not be prompted for the Driver Update diskette. At this point you need to abort the installation and reboot your server. You can do this by pressing the **Abort Installation** button and selecting the **Exit/Reboot** option in the following prompt. When the server boots up again, make sure you press the **Alt** key on the initial install screen.

2. The following message is displayed: Please get your Driver Update Floppy ready. Press **Enter**.
3. Insert the driver diskette when prompted. Press **Enter**.
4. The YaST installation wizard starts and the License agreement is displayed. Click **Accept if you agree to** the terms and continue with the installation.
5. The YaST installation wizard will display the phases of install on the right part of the screen. Specific settings are required for the **Installation Settings** phase. Settings in all other phases can be answered as you prefer.

**Important:** If you did not press **Alt** as described in step 1, you will see a message indicating that no hard drives were found. You must click on **Abort Installation** and then select **Exit/Reboot**. Your server will reboot and will start the installation over again. Make sure you press **Alt** as described in step 1.

6. During the **Installation Settings** phase, the categories listed below require the indicated settings. For any other categories, you can select the settings you prefer. To modify the settings in these categories, either click on the category heading or click on **Change...** and select the category from the list.

**a. Partitioning - The default settings are OK. If you are an experienced user, you can modify the partition setup.**

**Important:** If you choose to modify the partition setup, you **must not** modify the FAT partition on /dev/sdb1. Otherwise the server will not boot.

**b. Software - You must verify that the correct 32-bit kernel is being installed. You must also verify that the gettext package is being installed. To check this, do the following:**

- i. Click on **Software**.
- ii. On the Software Selection screen, select **Default System** and click on the **Detailed selections** button.
- iii. Click on the **Filter** drop down list and select **Package Groups**. This will show the RPM packages for the install.
- iv. Under System select **Kernel**.
- v. In the list of kernels in the upper-right panel, if **k\_deflt** is checked, you do not need to make any changes. If **k\_smp** is checked, you **must** uncheck it and check **k\_psmpl**.
- vi. Under Development select **Tools**.
- vii. In the list of tools in the upper-right panel, check **gettext**.
- viii. Click **Accept**.

**c. Booting - This must be set to Grub (MBR on /dev/sda) which is the default.**

7. After you have finished customizing the settings click **Accept** then **Yes, install** when the warning message appears. The install will prompt you for United Linux 1.0 CDs and will reboot when it is done copying data to start the Configuration phase of the install.
8. Continue on with the rest of the configuration phases of install.

9. When the configuration phases are complete, you will be prompted to sign on to the server. You must follow the instructions described in section titled *Completing the postinstallation tasks*.
10. Once the postinstallation instructions are complete, you can apply SLES updates to your server.  
**Important:** After applying SLES updates, you must run **ixsupdt** before you reboot your server. Details for **ixsupdt** are listed in the section titled *Maintaining the Linux integration code*.

## Completing the installation of SLES 9 on the IXA

This topic discusses the steps necessary to complete the SUSE Linux Enterprise Server installation.

### About this task

To complete the install of a SUSE Linux Enterprise Server 9 server:

1. The xSeries server boots from the SUSE Linux Enterprise Server 9 disc 1 located in the xSeries CD-ROM drive. A list of options will appear in a menu list for 20 seconds. You **must** do the following:  
Select **Installation**.  
Press **F6** for Driver Update.  
Press **Enter**.  
If you do not make any selections within the 20 seconds, the default option to Boot from Hard Disk will be used. Because there is no operating system on the hard disk yet, the message 'Failed to start from Harddisk' will be displayed. Press OK and make the selections as described above.  
**Important:** If you do not press F6 for Driver Update, you will not be prompted for the driver update floppy, which will cause the installation to fail to find hard drives when you get to step 8. At which point you will have to abort the installation and reboot your server, which will start the installation again as described in step 8.
2. When you press F6, the following message is displayed: **Please get your Driver Update Floppy ready**. Press **Enter**.
3. Insert the driver diskette when prompted. Press **Enter**.
4. A dialog will be presented with the message **Please choose the Driver Update Medium**. From the menu list, select **fd0 (floppy)**, select the **OK** button and press **Enter**.
5. The driver diskette will be scanned and you will see the message **Reading the Driver Update...** and then the message **Driver Update OK**. Press **Enter**.
6. The dialog in step 4 will reappear. Select **Back** and press **Enter**.
7. The YaST installation wizard starts and the License agreement is displayed. Click **I Agree** to accept the terms and continue with the installation.
8. The YaST installation wizard will display the phases of install on the right part of the screen. Specific settings are required for the **Base Installation - Installation Settings** phase and the **Configuration - Online update** phase.  
Settings in all other phases can be answered as you prefer. For more information about installing SLES with YaST, you can reference the *SUSE Linux Enterprise Server Installation and Administration* manual which you will find on the SUSE Linux Enterprise Server 9 disc 1 under `/doc/en/manual.pdf`.  
**Important:** If you did not press **F6** as described in step 1, you will see a message indicating that no hard drives were found. You must click on **Abort Installation** and then select **Exit/Reboot**. Your server will reboot and will start the installation over again. Make sure you press **F6** as described in step 1.
9. During the **Base Installation - Installation Settings** phase, the categories listed below require the indicated settings. For any other categories, you can select the settings you prefer. To modify the settings in these categories, either click on the category heading or click on **Change...** and select the category from the list.

- a. Partitioning - The default settings are OK. If you are an experienced user, you can modify the partition setup.  
**Important:** If you choose to modify the partition setup, you **must not** modify the FAT partition on /dev/sdb1. Otherwise the server will not boot.
  - b. Software - You must verify that the correct kernel is being installed. You must also verify that the gettext package is being installed. To check this, do the following:
    - i. Click on **Software**.
    - ii. On the Software Selection screen, select **Default System** and click on the **Detailed selections** button.
    - iii. Click on the **Filter** drop down list and select **Package Groups**. This will show the RPM packages for the install.
    - iv. Under System select **Kernel**.
    - v. In the list of kernels in the upper-right panel, if **kernel-default** is checked, you do not need to make any changes. If **kernel-bigsmpt** is checked, you **must** uncheck it and check **kernel-smp**.
    - vi. Under Development select **Tools**.
    - vii. In the list of tools in the upper-right panel, check **gettext**.
    - viii. Click **Accept**.
  - c. Booting - This must be set to Grub (MBR on /dev/sda) which is the default.
10. After you have finished customizing the settings click **Accept** then **Yes, install** when the warning message appears. The install will prompt you for various SUSE Linux Enterprise Server 9 CDs and will reboot when it is done copying data to start the Configuration phase of the install.
  11. During the **Configuration - Network** phase of the install, you will see a screen labeled **Test Internet Connection**. You **must** select **No, Skip this test**.  
**Attention:** If you do not select **No, Skip this test**, a SLES update will be performed at this time which can cause unrecoverable damage to your server. A SLES update should only be done after the server has completed the install **and** you have completed the steps as described in "Completing the postinstallation tasks" on page 57.
  12. Continue on with the rest of the configuration phases of install.
  13. When the configuration phases are complete, you will be prompted to sign on to the server. You must follow the instructions described in "Completing the postinstallation tasks" on page 57.
  14. Once the postinstallation instructions are complete, you can apply SLES updates to your server.  
**Important:** After applying SLES updates, you *must* run **ixsupdt** before you reboot your server. Details for ixsupdt are listed in "Maintaining the Linux integration code" on page 58.

## Completing a SLES 9 SP3 or later installation on an iSCSI attached BladeCenter or xSeries

This topic discusses the steps necessary to complete the SLES 9 SP3 installation for iSCSI attached servers.

### About this task

1. The xSeries server boots from the SUSE Linux Enterprise Server 9 Service Pack 3 (or later) disc 1 located in the xSeries CD-ROM drive. You **must** do the following:
  - a. Select **Installation**.

**Important:** If you do not make any selections within the 20 seconds, the default option to Boot from Hard Disk will be used. Because there is no operating system on the hard disk yet, the message "Failed to start from Harddisk" will be displayed. Press OK and then select Installation.



b. In the Boot Option text field type: `insmod=qla4xxx autoyast=device://sdb1/install/sles9ins.xml`

**Important:** If you do not enter the Boot Option text as described above, the installation may fail to find hard drives during the install interview or when the server reboots to complete the install.

c. Press **Enter**.

2. The install will display the message **'Make sure the CD number 1 is in your drive.'** The installation wants disc 1 from the base media, not disc 1 from the Service Pack.
3. The YaST installation wizard starts and the License agreement is displayed. Click **I Agree** to accept the terms and continue with the installation. Since we specified an autoyast file in the boot options in step 1, the default answer is **I Agree** and will be automatically chosen after 10 seconds.
4. The YaST installation wizard will display the phases of install on the right part of the screen. Specific settings are required for the **Base Installation - Installation Settings** phase and the **Configuration - Online update** phase. Settings in all other phases can be answered as you prefer. For more information about installing SLES with YaST, you can reference the *SUSE Linux Enterprise Server Installation and Administration* manual, which you will find on the SUSE Linux Enterprise Server 9 disc 1 under `/doc/en/manual.pdf`.

**Attention:** If you see a message indicating that no hard drives were found, it may be due to one of these reasons:

- The Boot Option text was not entered as described in step 1.
- The SUSE Linux Enterprise Server 9 disc 1 was used to start the installation. You must use SUSE Linux Enterprise Server 9 Service Pack 3 (or later) disc 1 to start the installation.
- There may be an error in the iSCSI initiator settings or Remote system NWSCFG object.

At this point, you must click on **Abort**, then **Abort Installation**, and then select **Exit/Reboot**. Your server will reboot and will start the installation over again. Make sure you follow the steps in Step 1.

5. During the **Base Installation - Installation Settings** phase, you can accept the default settings. To change the settings, either click on the category heading or click **Change...** and select the category from the list. Here is a list of the restrictions for each category:
  - a. **Partitioning:** If you choose to modify the partition setup, you **must not** modify the FAT partition on `/dev/sdb1`. Otherwise the server will not boot.
  - b. **Booting:** The Bootloader Type must be **GRUB** with a **Location** on `/dev/sda` (MBR).
  - c. **Software** You must verify that the linux32 package (32-bit emulation for x86-64) is installed. You must also verify that the gettext package is being installed. To check this, do the following steps:
    - 1) Click **Software**.
    - 2) On the Software Selection screen, select **Default System** and click on the **Detailed selections** button.
    - 3) Click on the **Filter** drop down list and select **Package Groups**. This will show the RPM packages for the installation.
    - 4) Under System select **Kernel**.
    - 5) In the list of kernels in the upper-right panel, you must make sure **linux32** is checked.
    - 6) Select **DevelopmentTools**.
    - 7) In the list of tools in the upper-right panel, check **gettext**.
    - 8) Click **Accept**.
6. After you have finished customizing the settings click **Accept**, then **Yes, install** when the warning message appears. The installation will prompt you for **both** base and Service Pack discs. If the prompt does not specifically ask for a Service Pack disc, then it wants a base disc. When the installation is done copying data it will reboot to start the Configuration phase of the install.
7. During the **Configuration - Network** phase of the install, you will see a screen labeled **Test Internet Connection**. You **must** select **No, Skip this test**. **Attention:** If you do not select **No, Skip this test**, a SLES update will be performed at this time which can cause unrecoverable damage to your server. A

SLES update should only be done after the server has completed the install **and** you have completed the steps as described in “Completing the postinstallation tasks.”

8. Continue on with the rest of the configuration phases of install.
9. When the configuration phases are complete, you will be prompted to sign on to the server. You must follow the instructions described in “Completing the postinstallation tasks.”
10. Once the postinstallation instructions are complete, you can apply SLES updates to your server.

**Important:** After applying SLES updates, you **must** run **ixsupdt** before you reboot your server. Details for **ixsupdt** are listed in the section “Maintaining the Linux integration code” on page 58.

---

## Completing the postinstallation tasks

You must perform some postinstallation tasks in order to complete the setup of the integrated Linux server.

### Run the **ixssetup.sh** script

This topic discusses the **ixssetup.sh** script, necessary to complete the Linux installation.

#### About this task

**Attention:** This step is not an optional step. This step ensures that the server will:

- Shut down cleanly
- Prepare for Linux distribution updates
- Lock/unlock virtual tape and optical devices
- Exchange administrative information with the hosting i5/OS partition

To run the **ixssetup.sh** script, complete the following steps:

1. Verify that the integrated hardware is connected to the hosting i5/OS partition using an external LAN connection.
2. Log into the Linux server as root.
3. Start a shell prompt:
  - For Red Hat Enterprise Linux, right-click the desktop and select **New Terminal**.
  - For SUSE Linux Enterprise Server, click the monitor and shell toolbar icon.
4. Enter the following command:

```
mkdir /mnt/sdb1
```
5. Enter the following command:

```
mount /dev/sdb1 /mnt/sdb1
```
6. Enter the following command:

```
/mnt/sdb1/install/ixssetup.sh
```
7. You are prompted to enter the IP address or DNS Name (host name) of the hosting i5/OS partition. Linux integration support needs this data in order to set up a connection to exchange administrative information with i5/OS. It is important that you specify a valid IP address, or host name that is registered in the local DNS.
8. The server reboots.

## Run the ixsupdt utility

This topic describes the ixsupdt utility to attain the latest IXS kernel and user-level code.

### About this task

After you have successfully run the ixsetup.sh script, you must also run the ixsupdt command. This is to ensure that you have the latest IXS kernel and user-level code.

The ixsupdt command must also be run immediately after applying a Linux distribution kernel update and before the server is rebooted. For more details on ixsupdt, see the section titled Maintaining the Linux integration code.

---

## Maintaining the Linux integration code

This topic discusses how to maintain and manage your integrated Linux server code.

### About this task

The Linux integration software has some components that run on i5/OS, and others that run on Linux. The Linux components are installed when you create the Linux server. The i5/OS-based code can be maintained using the normal PTF process; however, you need a way to maintain the integration code installed on Linux. You can do this from Linux using the ixsupdt command. This command updates the kernel code and user level code.

**Note:** Linux distributions have methods to update the software (that is, online update, update CD-ROMs, etc.). If a Linux distribution's update includes a kernel update, we recommend that you back up your server prior to applying the update. Immediately after applying the Linux distribution's update, you must run the ixsupdt command. This is to ensure that the latest iSeries Linux integration drivers are available for the updated kernel. To see a list of kernels supported by the Linux integration drivers, please check the following Web site:

<http://www.ibm.com/systems/i/bladecenter/linux/> 

To run the ixsupdt command, complete the following steps:

1. Log in to Linux as either the root user or a user with root authority, and start a terminal session.
2. At the command prompt enter the ixsupdt command and press **Enter**. The syntax is:

```
ixsupdt <userid> [<address>]
```

Where *userid* is an i5/OS user profile and *address* is the IP address or host name of the hosting i5/OS partition. The ixsupdt command allows you to change the IP address or host name listed as the managing iSeries for your server.

Linux integration support needs this data in order to set up a connection to exchange administrative information with i5/OS. It is important that you specify a valid IP address, or host name that is registered in the local DNS.

3. Linux issues a prompt for the i5/OS user's password. Type in the password and press **Enter**.
4. If the kernel has been updated, a reboot is required to run with the updated kernel.

---

## i5/OS operations for integrated Linux servers

This topic introduces using i5/OS to manage your integrated Linux servers.

### Overview

There are some operational tasks that must be performed from Linux, and others from i5/OS. This information describes how to operate an integrated Linux server from the i5/OS side.



You can operate the integrated Linux server from i5/OS using either iSeries Navigator or by entering CL commands on the i5/OS command line.

**Note:** All references to Linux apply equally to the supported Red Hat and SUSE distributions, except where noted. The term *storage space* means a network server storage space. A storage space is a chunk of iSeries single-level disk storage that Linux sees as a physical disk drive. The terms storage space and Linux drive are used synonymously. An integrated Linux server is an instance of Linux running on an Integrated xSeries Server, or an xSeries server attached via an Integrated xSeries Adapter. Each integrated Linux server has its own corresponding network server description (NWSD). The Integrated Linux servers are also referred to as simply *Linux servers*.

When an integrated Linux server is installed on iSeries, the following objects are created:

### Network server description (NWSD)

This is the main configuration object of an integrated server that i5/OS uses to control the server and describe its attributes. Each integrated Linux server instance has a corresponding NWSD.

The NWSD contains information that points the Linux server instance to the physical integrated server hardware. All other i5/OS objects needed to run a Linux server are linked to, and controlled by, the NWSD. The NWSD provides independence between instances of Linux and the x86 hardware that they run on. This unique architecture enables some powerful functions, such as the ability to:

- Create multiple Linux server instances that can be run on the same integrated hardware resource, although not at the same time. This provides a “dormant” server capability.
- Switch Linux instances between integrated hardware resources. This provides a “hot spare” backup server capability that can be quickly activated in the case of failure of an Integrated xSeries Server or Integrated xSeries Adapter.

**Note:** In V5R4, the NWSD Network server type information for \*ISCSI server connections incorrectly lists \*LINUX32 for the server operating system. Only 64-bit x86 Linux is supported on iSCSI attached servers.

### Server instances

Each integrated server or network server description represents a Linux or Windows server instance, not an Integrated xSeries Server, Integrated xSeries Adapter, or iSCSI HBA.

Each server instance has a corresponding network server description (NWSD) and vice versa. You can create multiple instances of Linux or Windows on one Integrated xSeries Server, Integrated xSeries Adapter, or iSCSI HBA, although only one instance can be active at a time. In other words, iSeries integration for Linux and Windows separates the hardware from the Linux and Windows instances that were created to run on that hardware. This is a unique feature of the iSeries Windows and Linux integration architecture, and provides major benefits in terms of availability and hardware utilization.

When performing operations on integrated Linux servers using iSeries Navigator, it is usually addressed in terms of servers. When using CL commands through the i5/OS command line it is usually addressed in terms of network server descriptions or NWSDs.

### Network server storage space

*Storage spaces* are chunks of iSeries disk storage that are created in an i5/OS auxiliary storage pool (ASP) and reserved for Linux or Windows. Because storage spaces are virtual disk drives created from i5/OS single-level storage, each storage space is scattered across all the physical disk drives in the ASP. They can reside in the system ASP, a user ASP, or an independent ASP (IASP), and appear as objects in the /QFPNWSSTG integrated file system directory.

From an i5/OS perspective, this portion of iSeries disk storage is called a *network server storage space*, or storage space for short. From the Linux (or Windows) perspective, this portion of iSeries

disk storage is seen as a physical disk drive. In fact these drives are completely virtual, so we often refer to storage spaces and disk drives as virtual drives or simply drives. When we are discussing integrated Linux servers we use these terms interchangeably. Each disk drive icon or storage space represents a Linux or Windows (virtual) disk drive.

To associate a storage space with a particular NWSD we *link* it.

The installation process automatically creates two storage spaces, called the *system drive* and *installation drive*, and links them to the NWSD. You can create more storage spaces for your server and link them to the NWSD after the installation has completed. Linux sees them as additional drives.

The administration of a Linux server's disk storage has both an i5/OS component and a Linux component. On the i5/OS side we create and link the storage space to the Linux server. On the Linux side we perform disk administration tasks such as partition management and formatting. These tasks are performed in exactly the same way on an integrated Linux server as they are on a standalone Linux server.

i5/OS storage management relieves Linux of the responsibility for providing physical device support. Therefore, common disk administration issues encountered in standalone Linux servers, such as disk drivers, addressing, configuration, and protection, are eliminated when you use integrated Linux servers.

### **Virtual Ethernet LAN line descriptions**

Indented under each NWSD is a communications line description that ends in PP. This line description, known as the *Point-to-Point Virtual Ethernet line*, is automatically configured when the integrated server is created. This communications line provides a virtual 1 Gbps Ethernet connection between the integrated server and its hosting i5/OS partition, and enables the server to exchange information with i5/OS. You need a separate line description for each Virtual Ethernet connection you set up between an integrated Linux server and other integrated servers or iSeries logical partitions.

### **TCP/IP interface for Point-to-Point Virtual Ethernet LAN line**

A TCP/IP interface is created specifically for the Point-to-Point Virtual Ethernet LAN line. By default, it has an address of the form 192.168.x.y, where x and y are automatically generated, and a subnet mask of 255.255.255.0. You can change the TCP/IP interface of Point-to-Point Virtual Ethernet LAN line after the server has been created. This interface should be an IP address and subnet that are not being used in your network. The Linux server has a corresponding TCP/IP address (192.168.x.y+1) configured on its end of the Point-to-Point Virtual Ethernet LAN.

## **Setting up your operational interface**

This topic discusses the iSeries Navigator and i5/OS command line interfaces that you can use to manage your integrated Linux server.

### **About this task**

i5/OS uses two operational interfaces:

#### **iSeries Navigator**

iSeries Navigator is a Windows-based, fully graphical interface to i5/OS operations. It is an easy way to perform i5/OS operational tasks in a well-designed GUI. iSeries Navigator has been designed from the ground up to be logical, consistent, and easy to use. Wizards are used extensively to make configuration easier. When using iSeries Navigator you do not need to remember commands, and you see much more information in the window than on a text-based display.

Most i5/OS operational tasks can be performed through iSeries Navigator, although there are some exceptions such as the command that powers down the iSeries. iSeries Navigator is the best choice if you are new to iSeries.

### **i5/OS command line**

The i5/OS command line is the traditional text-based interface where you can run Command Language (CL) commands directly. It is similar in function to the command-line interface in Unix, Linux, and Windows. CL is the scripting language used on iSeries, and you can group CL commands together to create a CL program.

The command-line interface is available through a 5250 or green screen session, where the term 5250 refers to the text-oriented datastream used, and green screen refers to the default color of the characters in the text-based session. The iSeries Access program product provides 5250 sessions in Windows and in supported browsers. It is also possible to access a 5250 session on iSeries using telnet.

The command-line interface can be preferred by experienced iSeries users who know the CL commands, because it is faster and more powerful than the graphical interface. It is, however, more difficult to use for an inexperienced iSeries user.

Use the iSeries Navigator graphical interface to operate your integrated Linux servers, except where noted. For managing integrated Linux servers you will find it much easier to learn and navigate your way around compared with the command line.

## **Managing your integrated Linux server with iSeries Navigator**

This topic discusses how to get started using the iSeries Navigator operational interface with your integrated Linux server.

### **About this task**

To operate integrated Linux servers on iSeries using a graphical interface, you need to have iSeries Navigator installed on your Windows PC. iSeries Navigator is a no-charge component of the iSeries Access program product, which is included as standard with the i5/OS installation CD-ROMs.

Before you can use iSeries Navigator you must install it. You can install it from the iSeries Access CD-ROM, or directly from the iSeries itself, providing that the 5722-XE1 licenced program product is installed. Note that you do not need to have 5722-XE1 installed in order to use iSeries Navigator on your Windows PC. For more information on installing iSeries Navigator, refer to the Installing iSeries Navigator.

Note that iSeries Navigator requires that 5722-SS1, option 12 (i5/OS - Host Servers) be installed on your iSeries. This option should be installed by default, but you can check by using the Go Licensed Programs command (GO LICPGM) and selecting option 10.

After you install iSeries Access on Windows, you must download and install the latest service pack from the following Web site:

<http://www.ibm.com/servers/eserver/series/access/casp.html> 

To access the Linux administration functions of iSeries Navigator, complete the following steps:

1. Launch iSeries Navigator from your Windows desktop.
2. Click **iSeries connection** → **Integrated Server Administration**.

The options for administering integrated Linux servers are included under the Integrated Server Administration option of iSeries Navigator. However, there are some Windows administration functions, such as Windows User Enrollment, which are not available for Linux servers.

There are two iSeries Navigator - Integrated Server Administration displays that you use regularly when managing integrated Linux servers. They are:

**Servers**

This option is used to manage integrated Linux and Windows server instances installed on iSeries, and is similar in function to the WRKCFGSTS \*NWS CL command that you enter on the i5/OS command line.

**All Virtual Disks**

This option is used to manage integrated Linux and Windows server disk drives, and is similar to the WRKNWSSTG CL command that you enter on the i5/OS command line.

**| Managing your integrated Linux server with the i5/OS command line**

| This topic discusses how to use the i5/OS command line operational interface to manage your integrated server and some of the CL commands that are available to you.

**| About this task**

| The i5/OS command line interface is preferred by many iSeries professionals because it is faster and more powerful than the iSeries Navigator interface. However, for inexperienced users it is more difficult to use because it requires a knowledge of the i5/OS CL commands.

| To use the i5/OS command-line interface, type a CL command on the command line and press **F4**. If you need additional lines to enter commands, type the CL command CALL QCMD on the command line and press **Enter**. Note that CL commands can be entered in upper or lower case.

| The list of commands that you can use for managing integrated Linux servers are:

Commands	Description
ADDNWSSTGL	Add Server Storage Link
CRTNWSSTG	Create NWS Storage Space
DLTNWSSTG	Delete NWS Storage Space
DLTLNXSVR	Delete Linux Server
DSPNWSSTG	Display NWS Storage Space
INSLNXSVR	Install Linux Server
RMVNWSSTGL	Remove Server Storage Link
VRYCFG	Vary Configuration
WRKNWSSTG	Work with NWS Storage Spaces
WRKNWSSTS	Work with NWS Status
WRKNWSCFG	Work with network server configuration objects
WRKDEVD *NWSH	Work with NWSH device descriptions

| To use the i5/OS command-line interface you need to either gain access to a twinax-attached terminal, or more commonly, set up a 5250 (green screen) emulator under Windows or a supported browser. Both the iSeries Access program product and IBM Personal Communications product come with a 5250 emulator, which is a chargeable option. In order to use 5250 sessions to communicate with i5/OS, you need to have ordered the Enterprise package for your iSeries. If you do not have the Enterprise package, your 5250 capability is limited to console functions only.

## Using iSeries Navigator to work with integrated Linux server NWSDs

This topic discusses how to work with the Network Server Description (NWSD) using iSeries Navigator.

### About this task

The Integrated Server Administration component of iSeries Navigator is a valuable tool to help display and manage NWSD information.

To work with an NWSD using iSeries Navigator, complete the following steps:

1. Click **Servers** under Integrated Server Administration.
2. Right-click the server instance (NWSD) that you want to work with.
3. From the context menu of the selected NWSD you can:
  - Start, shut down, or shut down and restart the integrated Linux server instance.
  - View the status of the integrated Linux server instance.
  - View the properties of the NWSD.

## Using the i5/OS command line to work with integrated Linux server NWSDs

This topic discusses how to use CL commands to work with the Network Server Description (NWSD) for your integrated Linux server.

There are two CL commands that you use regularly when managing integrated Linux servers.

- Work with Configuration Status (Network Server Descriptions) `WRKCFGSTS *NWS`
- Work with Network Server Descriptions `WRKNWSD`

These commands display a list of the Linux, Windows, and AIX network server descriptions (NWSDs), and are similar to the Integrated Server Administration - Servers option of iSeries Navigator. Note that you cannot change the way the data is arranged on these displays.

### WRKCFGSTS \*NWS

The `WRKCFGSTS *NWS` display shows the status of each NWSD and its associated objects.

Type `WRKCFGSTS *NWS` on the i5/OS command line and press **Enter**. You see a list of the Linux, Windows, and AIX NWSDs, and their statuses.

The common statuses for an NWSD are:

- **ACTIVE** status of an NWSD is equivalent to a Linux server being up and running.
- **VARIED OFF** status of an NWSD is equivalent to a Linux server being shut down.
- **VARY ON PENDING** status of an NWSD is equivalent to a Linux server in the process of starting up.
- **SHUTDOWN** status of an NWSD is equivalent to a Linux server having been shut down from the Linux console.

### WRKNWSD

The `WRKNWSD` command displays a list of NWSDs.

1. Type `WRKNWSD` on the i5/OS command line. Press **Enter**. You see a list of the Linux, Windows, and AIX NWSDs.
2. Type 2 next to the NWSD you want to edit (or 5 to view). Press **Enter**.

The following parameters are used:

- Network server description
- Resource name
- Vary on wait
- Shutdown timeout
- Code page
- Server message queue and library
- Configuration file and library
- TCP/IP port configuration and sub parameters
- Restrict device resources
- Synchronize date and time
- Text description

You can change the following parameters when the server is down:

- Resource name
- Vary on wait
- Shutdown timeout
- Code page
- Server message queue and library
- Configuration file and library
- TCP/IP port configuration and sub parameters v Restricted device resources
- Restrict device resources
- Synchronize date and time
- Text description

You can change the following parameters when the server is up:

- TCP/IP port configuration and sub parameters. The change does not take effect until you restart the server.
- Text description

## Starting integrated Linux servers

This topic discusses how to start single and multiple integrated Linux servers using the iSeries Navigator and CL commands. It also includes instructions on how to start integrated Linux servers automatically at iSeries startup.

### About this task

**Note:** You cannot start an integrated Linux server from the Linux console; however, you can restart a running server from the Linux console.

## Starting a single integrated Linux server using iSeries Navigator

### About this task

The **Integrated Server Administration** function of iSeries Navigator provides an interface for starting Linux servers individually or in multiples.

To start a single Linux server using iSeries Navigator, complete the following steps:

1. Click **Servers** under Integrated Server Administration.
2. Right-click the server you want to start.
3. Select **Start**. After a few moments, you should see the status change to Started.



## Starting multiple integrated Linux servers using iSeries Navigator

### About this task

To start all integrated servers (both Linux and Windows) select one of the following methods:

- Right-click **Servers** in the left pane and select **Start all**.
- Click **Start all integrated servers** in the Taskpad pane.

To start selected servers only, hold down the control key (Ctrl) and click each of the servers you want started. Then right-click one of the selected servers and click Start.

### Tips:

- If you have two or more server instances created for the same Integrated xSeries Server or IXA-attached xSeries server, you should not try to start all servers simultaneously. Remember, these are server instances, not physical servers.
- Because of interdependencies between Linux servers, you might wish to start the servers in a specific order.

## Starting integrated Linux servers using CL commands

### About this task

To start a Linux server from an i5/OS command line, select one of the following CL commands:

#### WRKCFGSTS

1. Type WRKCFGSTS \*NWS. Press **Enter**.
2. Type 1 in the Opt column next to the network server description (NWSD) that you want to vary on. Press **Enter**.

To start multiple Linux servers using the WRKCFGSTS command, simply type 1 in the Opt column next to all the network server descriptions (NWSD) that you want to vary on. Press **Enter**.

#### VRYCFG

1. Type VRYCFG. Press **F4**.
2. Enter the NWSD that you want to vary on. Press **Enter**. Here is an example:

```
VRYCFG CFGOBJ(nwsd-name) CFGTYPE(*NWS) STATUS(*ON)
```

To start multiple Linux servers using the VRYCFG command, simply list all the network server descriptions (NWSDs) that you want to vary on in the CFGOBJ parameter. Press **Enter**. For example:

```
VRYCFG CFGOBJ(nwsd-name1 nwsd-name2 nwsd-name3) CFGTYPE(*NWS) STATUS(*ON)
```

Using the VRYCFG command in this way could lock up your green screen session for several minutes.

## Starting integrated Linux servers automatically at iSeries startup

### About this task

If you only have a single integrated server, or you do not need to start your servers in a specific order, there is a simple way to start an integrated server when the iSeries is turned on:

1. On the i5/OS command line type CFGTCP. Press **Enter**.
2. On the Configure TCP/IP display type 1. Press **Enter**.
3. Locate the interface for the Point-to-Point Virtual Ethernet line that corresponds to the server you want to start automatically. It is of the form nwsd-namePP. Enter 2 next to the interface. Press **Enter**.
4. Change the Autostart parameter to \*YES. Press **Enter**.

## Results

When you next power up the iSeries, the integrated server starts automatically.

## Shutting down integrated Linux servers

This topic discusses the multiple methods to shut down integrated Linux servers.

### About this task

Just as starting Linux servers can be accomplished in multiple ways, so can stopping Linux servers. In this section we cover ways to shut down Linux servers, and unique situations that arise in an integrated environment.

You must shut down your Linux servers before performing i5/OS system tasks such as:

- Backing up statically linked storage spaces. Statically linked storage spaces can only be backed up when Linux servers are shut down.
- Putting an i5/OS partition into restricted state. When you end i5/OS to restricted state, the iSeries can only be accessed through the console. If you put your iSeries into restricted state by ending all subsystems immediately, your integrated Linux servers are not automatically shut down. Because you normally put i5/OS into restricted state to perform a backup, you need to shut down the Linux servers in order to back up their components. If a server is running, components such as storage spaces are locked, and excluded from the back up. Therefore, you should shut down your integrated servers in a controlled manner before putting i5/OS into restricted state.
- Performing a complete system backup. Performing a complete system save using option 21 on the Save menu is an easy way to ensure that you have a complete backup of your iSeries system, including your integrated Linux servers. Option 21 on the Save menu has an option that allows you to automatically shut down integrated servers before starting a backup. To set this option:
  1. Type go save on the i5/OS command line. Press **Enter**.
  2. Type 21, press **Enter**.
  3. You see the Save the Entire System display. Press **Enter**.
  4. On the Specify Command Defaults display there is an option called "Vary off network servers". This option defaults to *\*NONE*, which does not shut down integrated servers. To ensure that integrated Linux servers are shut down before the backup starts, specify *\*ALL*.
  5. Fill out the other options. Press **Enter** to start the backup.
- Powering down the iSeries. It is strongly recommend that you shut down your servers in a controlled manner before powering down the iSeries using the Power Down System (PWRDWN SYS) command.

**Important:** Always shut down your integrated Linux servers before you put i5/OS into restricted state, and prior to powering down the iSeries. Otherwise you might lose data or corrupt the Linux operating system.

## Shut down a single integrated Linux server using iSeries Navigator

### About this task

To shut down a single Linux server using iSeries Navigator:

1. Click **Servers** under **Integrated Server Administration**.
2. Right-click the server you wish to shut down.
3. Select **Shut Down** or **Shut Down and Restart**.
4. You see a shut down confirmation window. Click **Shut Down**.



## Results

After a few moments you should see the server status change to Shutting down. The status automatically changes to Shut down when the server is completely shut down. You do not need to refresh the window.

## Shut down multiple integrated Linux servers using iSeries Navigator

### About this task

To shut down all integrated servers simultaneously select one of the following methods:

- Right-click Servers in the left pane and select **Shut down all**.
- Click **Shut down all integrated servers** in the Taskpad pane.

To shut down selected servers simultaneously, hold down the control key (Ctrl) and click each of the servers you want shut down. Then right-click one of the selected servers and click **Shut down**.

Because of interdependencies between servers, you might wish to shut down the servers in a specific order. If you want to do this automatically, see “Shut down a single integrated Linux server using CL commands.”

## Shut down a single integrated Linux server using CL commands

### About this task

To shut down a single Linux server from an i5/OS command line, select one of the following CL commands.

#### WRKCFGSTS

1. Type WRKCFGSTS \*NWS. Press **Enter**.
2. Type 2 in the Opt column next to the network server description (NWSD) that you want to vary off. Press **Enter**.

To shut down multiple Linux servers using the WRKCFGSTS command, simply type 2 in the Opt column next to all the network server descriptions (NWSDs) that you want to vary off. Press **Enter**.

#### VRYCFG

1. Type VRYCFG. Press **F4**.
2. Enter the following parameters:
  - The NWSD that you want to vary off.
  - The status: \*OFF.
3. Press **Enter**.

Here is an example:

```
VRYCFG CFGOBJ(nwsd-name) CFGTYPE(*NWS) STATUS(*OFF)
```

To shut down multiple Linux servers using the VRYCFG command, simply list all the network server descriptions (NWSDs) that you want to vary off in the CFGOBJ parameter. Press **Enter**. For example:

```
VRYCFG CFGOBJ(nwsd-name1 nwsd-name2 nwsd-name3) CFGTYPE(*NWS) STATUS(*OFF)
```

## Shut down integrated Linux servers from the Linux console

### About this task

Every Linux server can be shut down from the Linux console by entering the command **shutdown -h now**. The Linux operating system shuts down but the NWSD is still considered to be varied on.

When Linux is shutdown from the console on an IXS or IXA attached server, the WRKCFGSTS command shows the NWSD with a status of SHUTDOWN, and the status of the communications objects under it are FAILED.

When Linux is shutdown from the console on an iSCSI attached server, the WRKCFGSTS command shows the NWSD with a status of VARIED ON.

On the WRKCFGSTS \*NWS display you can enter a 2 next to the NWSD to complete the vary off.

## **Shut down integrated xSeries servers from the Linux console**

### **About this task**

Shutting down a Linux server that is running on an Integrated xSeries Adapter-attached xSeries from the Linux console is a special case. In this case you should not shut Linux down through the Linux console. The reason for this is that when you shut down an integrated Linux server in this way, the xSeries server is turned off completely. This causes the Integrated xSeries Adapter to lose power and stop communicating over the HSL loop. The System Licensed Internal Code (SLIC) on the iSeries interprets this as one of its towers (the Integrated xSeries Adapter-attached xSeries in this case) having been turned off, and removed from the HSL loop. The iSeries goes into HSL recovery. The other iSeries towers in the loop should remain unaffected because data flows both ways around the loop, and the other towers still have a connection to the iSeries system unit. But if multiple Integrated xSeries Adapter-based integrated servers are shut down from the Linux console, towers that are cabled between the shut down servers on the HSL loop become isolated and no longer have a connection to the iSeries system unit. If the isolated towers contain disk units, the iSeries terminates abnormally. Also, if the loop contains a migration tower, shutting down a single Integrated xSeries Adapter-based server from the console can also bring down the iSeries because data can only flow one way around the HSL loop to a migration tower.

For integrated Linux servers running on Integrated xSeries Adapter-attached xSeries machines, the correct way to shut down is from iSeries Navigator or the i5/OS command line.

For more information on HSL configuration, search for redp-3652 at the following Web site:

<http://www.redbooks.ibm.com/> 

## **Shut down integrated Linux servers by forcing a shut down**

### **About this task**

When you issue a shut down command to an integrated Linux server from iSeries Navigator or an i5/OS command line, the shut down command is passed to the server to enable it to shut itself down cleanly. By default, i5/OS waits 15 minutes (900 seconds) before it performs a forced shut down of an integrated Linux server if the server has not already shut itself down. If you want to change this default value, do the following steps:

1. Shut down the Linux server.
2. On an i5/OS command line enter the Change Network Server Description (CHGNWSD) command and press **F4**.
3. Scroll down to the Shutdown time-out (SHUTDTIMO) parameter and enter the new default value. Press **Enter**.
4. Restart the server.

## **Shut down integrated Linux servers by powering down the iSeries**

### **About this task**

It is strongly recommended that you do not power down your i5/OS partition when Linux servers are active. When the Power Down System (PWRDWN SYS) \*IMMED command is issued, i5/OS does not try to shut down any active integrated Linux servers. When the iSeries powers down it cuts power to the

Linux servers. There is a possibility of data corruption if the servers suffer a forced shut down in this way. Always shut down your Linux servers cleanly prior to issuing the PWRDWNSYS command.

**Note:** The Power Down System (PWRDWNSYS) \*CNTRLD command will attempt to vary off any active servers prior to powering down the system.

**Important:** Never issue a PWRDWNSYS \*IMMED command without having your integrated Linux servers in shutdown status.

---

## Networking integrated servers

This topic introduces methods to network your integrated Linux servers.

### Configure virtual Ethernet networks

This topic describes how to configure a virtual Ethernet network between integrated Linux servers.

#### About this task

**Note:** If you are installing an integrated Linux server from scratch, the installation command (INSLNXSVR) can configure virtual Ethernet networks for you.

To configure a virtual Ethernet network, complete the following steps:

1. Configure an Ethernet line description for the integrated server.
  - a. On i5/OS, enter the Create Line Description (Ethernet) command, CRTLINETH, and press **Enter**.
  - b. In the Line Description field, enter the name of your network server description (NWSD), followed by V and a number (0 through 9) that corresponds to the virtual Ethernet network that you will use (for example, *NWSDnameV0*).
  - c. In the Resource Name field, enter \*NWSD.
  - d. In the Network server description field, enter the name of the NWSD that will use the adapter.
  - e. Press **Enter** twice.
  - f. In the Port Number field, enter the port number that corresponds to the virtual Ethernet network that you will use. Virtual Ethernet network port values are \*VRTETH0 through \*VRTETH9.
  - g. In the Local Adapter Address field, \*ADPT is required.
  - h. The Line Speed must be 1G and Duplex must be \*FULL.
  - i. Press **F10** to see more parameters. The Max frame size should be 8996 for virtual Ethernet.
  - j. Optional: Page down to the Text 'description' field and write a brief description.
  - k. Optional: Page down to the Link Speed field and specify \*MAX. Press **Enter**.
  - l. Press **Enter** to create the Ethernet line description.
  - m. Optional: If you want this integrated server to be connected to more than virtual Ethernet network, repeat all of the above steps to create a line description for each network, using different port values in step 1f.
2. Change the NWSD of the integrated server to use the Ethernet line descriptions.
  - a. Enter the i5/OS command CHGNWSD and press **Enter**.
  - b. In the Network server description field, enter the name of your NWSD and press **F4**.
  - c. Page down to TCP/IP port configuration.
  - d. In the blank to the right of the words + for more values, enter a + and press **Enter**.
  - e. In the Port field, enter the port number. (For example, \*VRTETH0).
  - f. In the Internet address field, enter the IP address that the integrated server will use.
  - g. In the Subnet mask field, enter the subnet mask that the integrated server will use.
  - h. In the Max Transmission Unit field, enter 8996.

- i. Press **Enter**.
3. Repeat the procedure for all the integrated servers you want to connect to the network, specifying the same virtual Ethernet port for each one.
4. Restart the integrated servers. A virtual Ethernet adapter device driver will be automatically installed and set to the Linux TCP/IP address that has been specified for it in the NWSD. However, an IP address entered at the integrated server console overrides the values that are set in the NWSD.
5. Test to see that the virtual Ethernet network is functioning, for example by pinging from one server to the IP addresses you specified for the other servers.

## What to do next

---

### Managing Linux drives

This topic describes information and advice on how to manage iSeries disk storage allocated to integrated Linux servers.

#### iSeries storage management overview

This topic discusses iSeries storage management methods.

When Linux is running on the integrated server hardware, it uses a portion of the iSeries disk capacity. For this reason, the administration of a Linux server's disk storage has both an i5/OS component and a Linux component. The i5/OS component is used to create and link a chunk of storage to the Linux server. The storage management function of i5/OS relieves the Linux server from the responsibility of providing physical device support. Therefore, many of the common disk administration tasks encountered in stand-alone PC servers (disk drivers, addressing, configuration and protection) are eliminated when you use an integrated Linux server. The Linux component of disk storage administration is used to perform the remainder of the disk administration tasks, including formatting and partitioning. These tasks are performed on an integrated Linux server in exactly the same way as they are on a stand-alone Linux server.

The key to understanding how disk storage is allocated to integrated Linux and Windows servers is an understanding of how i5/OS storage management works on the iSeries. The heart of storage management on the iSeries is a technology called single-level storage. Single-level storage is a revolutionary storage management architecture that not only gives the iSeries outstanding disk I/O performance, but greatly reduces the amount of administration required.

The major features of single-level storage are:

- Single storage pool

The management of physical disk drives is implemented in the Systems Licensed Internal Code (SLIC), which is similar in concept to the BIOS on a PC.

By default, the operating system and applications see only a single large pool of virtual storage (called the System Auxiliary Storage Pool or system ASP) rather than physical drives. Therefore, the management of physical storage is hidden from the user.

To increase the size of the pool, simply add disk drives to the iSeries system and they automatically become part of the system ASP. Note that under some circumstances you might wish to create additional storage pools that are called user ASPs and independent ASPs.

- Scattering of data

Instead of an object being stored on a single physical disk drive, single-level storage scatters objects across all physical drives, transparently to the user.

iSeries disk management supports fully parallel disk I/O, which provides outstanding disk I/O performance because each object on the system is accessible by multiple disk arms concurrently.

There is no need to be concerned about particular disk drives filling up, or moving data from one disk to another to improve performance because all data management is taken care of by SLIC. Therefore, iSeries does not require a Database Administrator. SLIC also ensures that there is no disk fragmentation.

- Single address space

Memory and disk on iSeries form a single 64-bit address space.

A single address space enables objects to be accessed by name rather than hardware address, which provides additional integrity and reliability.

## Introducing network server storage spaces

This topic discusses network server storage spaces and how to name them.

Operating systems such as Windows, Unix and Linux work with what they see as physical disk drives; there is little or no virtualization of storage at an operating system level. Because i5/OS virtualizes all disk storage, you can carve out chunks of disk space from the storage pool to form virtual disk drives, which we can then allocate to Linux and Windows. We call these virtual disk drives, storage spaces. Integrated Linux and Windows servers, as well as AIX 5L™ and Linux running in iSeries partitions, see these storage spaces as physical disk drives.

**Important:** Because drives, as seen by integrated Linux and Windows servers, are physically scattered over all disk drives in the ASP, you can size a Windows or Linux drive up to the available storage in the specified ASP.

The i5/OS object that is used to create a Windows or Linux virtual drive is called a Network Server Storage Space (NWSSTG), or storage space for short. These storage spaces are stored in the root of the i5/OS integrated file system in a directory called /QFPNWSSTG. You can use the Work with Links (WRKLNK) command from an i5/OS command line to view the contents of the /QFPNWSSTG directory. This same storage space architecture is used not only by integrated Linux servers, but also by integrated Windows servers, and Linux and AIX 5L running in iSeries logical partitions. Storage spaces can be interchanged between each of these different operating systems.

The amount of disk storage that you create for your servers is taken directly from the iSeries available storage, and each Windows or Linux virtual drive is physically scattered across all the real disk drives in the iSeries disk pool.

Storage spaces are different from other i5/OS file objects because the size that you specify for a storage space is completely allocated at the time it is created. This is because integrated Linux and Windows servers need to be able to connect to and format a drive of a fixed size.

The Linux server install command (INSLNXSVR) creates only two drives. These drives are used to store the Linux operating system (called the system drive) and some IBM-supplied drivers (called the installation drive).

It is a good idea to make a backup of the system drive before and after you make changes to the operating system. If something should happen, you can recover by restoring a backup of the system drive, rather than rebuilding the server from scratch. In order to recovery quickly from a system failure, you should not store user files on the system or installation drives. Files and data that change frequently should be stored on a different drive.

Before you start creating new drives for your server, take some time to calculate what the server needs now and in the future. After the server has been installed you can create additional drives for your Linux server at any time. These drives can be linked to the server while it is shut down (static linking) or started (dynamic linking). This means that you do not need to allocate large portions of your iSeries storage when the server is created; you can create additional drives of any size you wish (up to the limit) when they are needed.

Here is a summary of the operations that you can perform on Linux drives:

- Create a new Linux drive.
- Delete a Linux drive.
- Link a Linux drive.
- Unlink a Linux drive.
- Clone a Linux drive.
- Expand a Linux drive.

Linux disk drive operations can be performed in two ways:

- Using iSeries Navigator.
- Using CL commands.

## Naming storage spaces

It is important to decide on a naming convention for your storage spaces; otherwise you might have trouble correlating storage space names that you see from the i5/OS side with drives that you see from the Linux side. This can be especially difficult if you have both statically linked and dynamically linked drives.

When the Linux server is created, two drives are created by default: The system and installation drives. They pick up the name of the NWSD with a 1 appended for the system drive, and a 2 appended for the installation drive. For example, for an NWSD named REDHAT1, a system drive named REDHAT11 and an installation drive named REDHAT12 are created. It is not possible to rename these drives directly. You must copy them and then supply a new name to the copy. However, it is rarely necessary to rename these two drives, as it is obvious which NWSD they are associated with.

We do, however, recommend that you systematically name any additional drives that you create for a server. We suggest that you use the next available Linux device name for each new drive. By default, Linux allocates the SCSI device names /dev/sda and /dev/sdb to the system and installation drives, respectively. Therefore, we recommend that you name the first additional storage space you create for a server nwsd-namesdc, nwsd-namesdd, and so on. Using this naming convention for an NWSD named REDHAT1, additional storage spaces would be named REDHAT1SDC, REDHAT1SDD, and so on. Note that the maximum length of a storage space name is 10 characters.

## Understanding storage space linking

This topic discusses the benefits of storage space linking.

Another benefit of storage spaces is that they are independent. This means that when you create a new storage space, it is not tied to any particular Linux server. A storage space must be linked to a network server description (NWSD) object before the Linux server can “see” it. Even after you link a storage space you can unlink it and link it to another Linux server. Note that a storage space can only be linked to one server at a time. However, you could use this virtual connectivity to move a data drive among different Linux servers at the touch of a mouse.

- | The Network server type of the NWSD will determine how the network storage space can be linked.
- | NWSDs for iSCSI attached xSeries or BladeCenter solutions running Linux have a Server connection of \*iSCSI. The iSCSI Linux server will only allow storage spaces to be linked as dynamic.
- | NWSDs for IXS or IXA attached xSeries running Linux have a Server connection of \*IXSVR. NWSDs for IXS or IXA will allow storage spaces to be linked as either fixed (static) or dynamic.
- | • Fixed links (IXA/IXS)



| A storage space can only be linked as fixed when the integrated Linux server is shut down. Up to 16 storage spaces can be linked as fixed. You should always link storage spaces as fixed where possible, saving dynamic links for times when you need to add a drive while the server is active.

| The system and installation drives are linked as fixed when the server is created. Do not try to change these links to dynamic.

- | • Dynamic links (iSCSI, IXA, IXS)
  - | For iSCSI attached Linux servers, up to 64 storage spaces can be linked as dynamic.
  - | For IXS and IXA servers, up to 16 storage spaces can be linked as dynamic. A storage space can be linked as dynamic when the integrated Linux server is started or shut down. However, you would normally save your dynamic links for adding drives when the server is active. If you use up all your dynamic links, and you need to add more drives, the server must be shut down. Note that it is possible to change the link type from dynamic to fixed by varying off the NWSD, unlinking the storage space and relinking it as fixed.

After a storage space has been linked, you must format it under Linux before Linux can use it.

## Unlinking storage spaces

You might want to unlink a storage space in order to delete the storage space or to relink it to another Linux server. Storage spaces can be unlinked in two ways:

- Linux server is shut down
 

This is the method you would normally use to unlink a storage space from a Linux server. You can use this method regardless of whether the storage space was linked with a fixed or dynamic link. In this case, you shut down the Linux server and then perform the unlink operation.
- Linux server is up and running
 

At V5R3 or later, you can unlink a storage space while the Linux server is up and running. This is referred to as dynamic unlinking. Before you dynamically unlink a storage space you must ensure that:

  - No users are using the Linux drive to be unlinked.
  - The drive was linked dynamically.
  - If the disk is part of an active logical volume group, it cannot be unlinked.

| At V5R4 iSCSI attached Linux servers support dynamic unlinking. You can dynamically unlink any disk except for the storage spaces with sequence numbers 1 (system drive) and 2 (install drive).

## The link sequence

The link sequence is the order in which storage spaces are presented to the Linux server by i5/OS. Statically linked drives are always presented before dynamically linked drives in the link sequence. The link sequence corresponds to the order in which Linux sees its disk drive devices: /dev/sda, /dev/sdb, /dev/sdc, and so on.

The easiest way to determine the link type and sequence number for a storage space is to use the WRKNWSSTG command which shows a set of storage spaces linked to Linux server Redhat1. Some storage spaces are linked statically (\*FIX); others are linked dynamically (\*DYN). You can also view the link sequence graphically through iSeries Navigator.

| Notice that the link sequence for statically linked drives (\*FIX) starting at 3. The equivalent iSeries Navigator display shows the link sequence for statically linked drives (Fixed) starting at 1. Therefore, a statically linked drive that has a link sequence number of n on the WRKNWSSTG green screen display is shown with a link sequence number of n - 2 on the equivalent iSeries Navigator display.

| **Important:** You must retain the correct link sequence for the Linux system and installation drives. If the system or install drive is not linked with the correct link sequence, the server does not start. Therefore, you should never unlink the system or installation drives.

- For IXS and IXA platforms, the system drive must be statically linked with sequence number 3; the installation drive must be statically linked with sequence number 4.
- For iSCSI attached servers, the system drive must have link sequence number 1 and the install drive must have link sequence number 2.

Be careful about unlinking storage spaces from Linux servers that have an entry in `/etc/fstab`. If you unlink a drive from an integrated Windows server (other than the system or installation drives), Windows tolerates the missing drive and boots up normally, assuming that there are no application errors. Linux, however, detects the missing drive and halts the boot sequence. In this case, you need to sign on as root and go into maintenance mode to remove the file system entry for the drive you unlinked before the server can continue the boot process. If you relink the storage space, you need to add back the file system entry you previously deleted.

When you unlink a storage space, i5/OS optionally compresses the link sequence. For example, if you have four statically linked drives (link sequence numbers 3, 4, 5 and 6), and three dynamically linked drives (link sequence numbers 1, 2 and 3), when you boot the Linux server, these drives appear to Linux in the same order, but with the statically linked drives preceding the dynamically linked drives:

- `/dev/sda` = fixed link sequence number 3
- `/dev/sdb` = fixed link sequence number 4
- `/dev/sdc` = fixed link sequence number 5
- `/dev/sdd` = fixed link sequence number 6
- `/dev/sde` = dynamic link sequence number 1
- `/dev/sdf` = dynamic link sequence number 2
- `/dev/sdg` = dynamic link sequence number 3

If you unlink the storage space corresponding to the third statically linked drive (link sequence 5), the link sequence is compressed so that there are now three statically linked drives (3, 4 and 5) followed by the three dynamically linked drives (1, 2 and 3). When you restart the server, Linux issues an error for `/dev/sdg` because it allocates a device handle to the remaining drives in order. This can create a problem if you have applications accessing data through a specific mount point because, in the example given, after you unlink the storage space with fixed link sequence number 5, a different drive is mounted to each mount point (except for the system and installation drives).

It is possible to unlink a drive without compressing the link sequence. However, this has no effect because Linux still allocates a device handle to the drives it sees in order. That is, it does not leave a “hole” in the list of allocated devices. Therefore, in the previous example, Linux again issues an error for `/dev/sdg`.

- | It is recommended that, if you do need to unlink storage spaces, you compress the link sequence unless
- | you are planning to relink the drive later on, in which case the link sequence should not be compressed.
- | You should also remove the associated entry in `/etc/fstab` before you reboot or shut down your Linux
- | server.

## Manage Linux drives

This topic discusses how to manage your Linux drives using iSeries Navigator and CL commands.

### Managing Linux drives using iSeries Navigator

You can manage Linux disk drive operations using the iSeries Navigator - Integrated Server Administration - All Virtual Disks interface.



## Managing Linux drives using CL commands

You can manage Linux disk drive operations using the Work with Network Server Storage Spaces (WRKNWSSTG) CL command through the i5/OS command-line interface. The WRKNWSSTG CL command displays a list of the Linux and Windows storage spaces, and is similar to the Integrated Server Administration - All Virtual Disks option of iSeries Navigator.

In addition to listing the storage spaces, this display shows additional information for each storage space, including:

*Table 4. Storage space information*

Label	Description
% Used	The percentage of the storage space used
Size	The size of the storage space in megabytes
Server	The name of the NWSD that the storage space is currently linked to
Drive	The link sequence number of the drive in terms of fixed or dynamic links
Link Type	The link type (fixed or dynamic). A Linux server sees statically linked (*FIX) drives before dynamically linked (*DYN) drives. The Linux server sees statically linked drives in the order specified by the Drive parameter. The Linux server sees dynamically linked drives in the order specified by the Drive parameter.
Format	<p>The nominal format of the storage space A drive that has not been formatted by Linux shows up in green screen with the nominal format you specified when you created the drive (*FAT, *FAT32, *NTFS or *OPEN). It is a nominal format because, in reality, the drive must be formatted by Linux before it can be used. We recommend that you specify *OPEN as the nominal format.</p> <p>The actual Linux format is not reported. This is because the implementation of Linux integration support does not know about Linux formats such as ext2 and ext3. Given that we recommend that you always create storage spaces using the *OPEN format, all Linux drives should appear as *OPEN whether they are formatted by Linux or not. The exceptions are the system and installation drives that are always reported by the Linux integration code as *FAT32 and *FAT respectively.</p>
Access	The type of access for this storage space. Read/write (*UPDATE) is the only option supported for Linux.
ASP	The Auxiliary Storage Pool (ASP) in which the storage space was created. This can be the system ASP, a user ASP, or an independent ASP.

## Creating Linux drives

This topic discusses how to create drives for your integrated Linux server.

## Creating Linux drives using iSeries Navigator

Adding a drive to an integrated Linux server is a two-step process:

- Firstly, you need to create the drive.
- Secondly, you need to link the drive to the server's NWSD.

To create a Linux drive using iSeries Navigator, complete the following steps:

1. Select one of the following methods:
  - Right-click All Virtual Disks in the left pane and select **New Disk** and input your desired values in the dialog.
2. Fill in the parameters as described in the following list:

Table 5. Linux drive description

Label	Description
Disk drive name	The name of the Linux drive you are creating. This name is also used for the storage space object created in the /QFPNWSSTG directory of the integrated file system. You should consider using a naming scheme to allow easy identification of storage spaces and to allow using generics (*) on the save commands.
Description	Description of the drive.
Initialize disk	Checkbox to copy data to the new drive from another drive. This option is used when cloning a Linux drive.
Capacity	Enter a size for the new drive in megabytes (MB) or gigabytes (GB).
Disk pool	The iSeries disk pool where the drive is created. The default is Disk Pool 1 (the system ASP), but drives can be created in user ASPs and independent ASPs.
Planned file system	You can use the value that you prefer. The suggested value is *OPEN for Linux specific file systems such as ext2, ext3, etc.

3. Click **OK**. A progress window opens,

The process of creating a storage space can range from a few minutes to a few hours, depending on the size. When i5/OS finishes creating the storage space it is listed with the other storage spaces.

**Important:** You must now link the drive to the server's NWSD before Linux can "see" it.

**Note:** You can also link the disk to the server by checking the "Link disk to a server" option on the New Disk panel.

## Creating Linux drives using CL commands

To create a Linux drive using a CL command, complete the following steps:

1. Type WRKNWSSTG. Press **Enter**. The Work with Network Server Storage display. Type 1 in first row of the Opt column and type in the name of the storage space that you want to create. Press **Enter**.
2. Alternatively, you can type CRTNWSSTG on the command line. Press **Enter**. The Create Network Server Storage Space display appears. Type in the name of the storage space that you want to create.
3. Fill in the parameters as described in the following list. Press **Enter**.
4. Type in a description. Press **Enter**. You see a message at the bottom of the display confirming that the storage space has been created.

Table 6. Storage space description

Label	Description
Storage space name	The name of the storage space. This name is also used for the storage space object created in the /QFPNWSSTG directory of the integrated file system. You should consider using a naming scheme to allow easy identification of storage spaces and to allow using generics (*) on the save commands.
Size	Enter a size for the new storage space in megabytes (MB) or gigabytes (GB).
From storage space	The storage space to copy data from. This option is used when cloning a Linux drive.
Format	A nominal format for the new storage space. For Linux we recommend this parameter be set to *OPEN. The storage space is not actually formatted by this command. Only Linux can format it.
ASP ID	The iSeries auxiliary storage pool where the storage space is created. The default is the system ASP, but storage spaces can be created in user ASPs and independent ASPs.
Description	Descriptive name for the storage space.

**Important:** You must now link the drive to the server’s NWSD before Linux can “see” it.

## Linking Linux drives

This topic discusses how to link integrated Linux server drives.

For your Linux server to see a newly created storage space, it must be linked to the server’s NWSD.

### Linking Linux drives using iSeries Navigator

To link your Linux drive, complete the following steps.

1. Click **All Virtual Disks** under Integrated Server Administration.
2. Right-click the drive you want to link.
3. Select **Add Link** to open the Add Link to Server window.
4. For “Link to server” select the NWSD that corresponds to your Linux server from the drop-down list.
5. Choose a Link type, either fixed or dynamic.
6. Optionally, you can specify the link sequence position, which is the order in which the Linux server sees it drives. This parameter defaults to the next available sequence position. Click View Sequence for a listing of the current order in which the storage spaces are linked.
7. Access to disk drive: Exclusive - Update is the recommended choice for a Linux server.
8. Click **OK** to link the drive to the NWSD.
9. You see the message confirming that the disk was linked successfully to the NWSD, and reminding you that you need to format it.

**Important:** After you have created and linked a storage space for a Linux server, you must partition and format the drive using a Linux utility before it can be used.

### Linking Linux drives using CL commands

To link a Linux drive using a CL command, complete the following steps:

1. Type in WRKNWSSTG. Press **Enter**. The Work with Network Server Storage Spaces display appears. Type 10 in the Opt column next to the storage space that you want to link. Press **Enter**.
2. Optionally, you can type ADDNWSSTGL on the command line. Press **Enter**. The Add Server Storage Link display.
  - a. For “Network server storage space” enter the storage space name.
  - b. For “Network server description” enter the NWSD that corresponds to your Linux server.
3. For “Dynamic storage link” leave the default of \*NO if you want to link the storage space statically, or change it to \*YES for a dynamic link. Press **Enter**.
4. Optionally, you can specify the drive sequence number, which is the order in which the Linux server sees it drives. This parameter defaults to the next available sequence position.
5. Press **Enter**. You see a message at the bottom of the display confirming that the storage space was linked successfully to the NWSD.

**Note:** After you have created and linked a storage space for a Linux server, you must partition and format the drive using a Linux utility before it can be used.

## Unlinking Linux drives

This topic discusses how to unlink integrated Linux server drives.

You can unlink a Linux drive from a Linux server either when the server is shut down or when it is up and running. Unlinking a dynamic storage space from a running server is only possible at V5R3 or later, and before you unlink the drive you must ensure that:

- No users are using the Linux drive to be unlinked.
- The drive was linked dynamically in the first place.
- If the disk is part of an active logical volume group, it cannot be unlinked.

### Unlinking Linux drives using iSeries Navigator

To unlink a Linux drive using iSeries Navigator, complete the following steps:

1. If you want to unlink the Linux drive dynamically, unmount the drive from the server using the umount command from the Linux server’s console.
2. If you do not want to unlink the drive dynamically, shut down the Linux server to which the drive is linked.
3. Click **All Virtual Disks** under **Integrated Server Administration**.
4. Right-click the drive you want to unlink.
5. Select **Remove link** to open the Remove Link from Server window. The disk drive name and description that you specified when you created the storage space are displayed.
6. Click **Remove** to unlink the drive.

### Unlinking Linux drives using CL commands

To unlink a Linux drive using a CL command, complete the following steps:

1. If you want to unlink the Linux drive dynamically, unmount the drive from the server using the umount command from the Linux server’s console.
2. If you do not want to unlink the drive dynamically, shut down the Linux server to which the drive is linked.
3. Type in WRKNWSSTG. Press **Enter**. The Work with Network Server Storage Spaces display appears. Type 11 in the Opt column next to the storage space that you want to unlink. Press **Enter**. The Remove Server Storage Link display appears.
4. Optionally, you can type RMVNWSSTGL on the command line. Press **Enter**. The Remove Server Storage Link display appears.

- a. For “Network server storage space” enter the storage space name.
  - b. For “Network server description” enter the NWSD that corresponds to your Linux server.
  - c. You might need to press **F9** to see the Renumber link parameter. It is recommended that you take the default of \*YES for this parameter unless you plan to relink the drive at a later time.
5. Press **Enter**. You see a message at the bottom of the display confirming that the storage space was unlinked successfully from the NWSD.

## Deleting Linux drives

This topic discusses deleting integrated Linux server drives.

### Deleting Linux drives using iSeries Navigator

To delete a Linux drive using iSeries Navigator, complete the following steps:

1. Shut down the Linux server to which the drive is linked.
2. Unlink the Linux drive that you want to delete.
3. Click **All Virtual Disks** under **Integrated Server Administration**.
4. Right-click the drive you want to delete.
5. Select **Delete** to delete the drive.
6. Click **Delete**.

To delete multiple Linux drives simultaneously, hold down the control key (Ctrl) and click each of the drives you want to delete. Then right-click one of the selected drives and click **Delete**.

### Deleting Linux drives using CL commands

To delete a Linux drive using a CL command, complete the following:

1. Shut down the Linux server to which the drive is linked.
2. Unlink the Linux drive that you want to delete.
3. Type in WRKNWSSSTG. Press **Enter**.

The Work with Network Server Storage Spaces display appears. Type 4 in the Opt column next to the storage space that you want to delete. Press **Enter**. The Confirm Delete Network Server Storage Space display appears.

4. Alternatively, you can type DLTNWSSTG on the command line. Press **Enter**. The Delete Network Server Storage Space display appears.
5. Press **Enter**. You see a message at the bottom of the display confirming that the storage space was deleted successfully.

## Sharing iSeries tape and optical drives with Linux

This topic discusses how to share optical drives with your integrated Linux server.

### About this task

Another advantage of using integrated Linux servers is the ability for Linux to share iSeries tape and optical devices. Your ability to use iSeries tape drives from Linux depends on the tape device model, tape controller, and media type. If you have logical partitions on your iSeries, a Linux server can only see the tape and optical drives that are allocated to the i5/OS partition where the server is hosted.

The main reason why you would share an iSeries tape drive with Linux is so that you do not need to buy additional tape drives specifically for Linux to use. All Linux servers can share an iSeries tape drive in the same host i5/OS partition, although not at the same time.

Each xSeries server comes with its own optical drive. Therefore, Integrated xSeries Adapter-attached xSeries machines are unlikely to need to share an iSeries optical drive. An exception can be if the iSeries optical drive is a DVD-RAM device. However, at the time of writing, the Linux integration software did not support writing to DVD-RAM from Linux.

The Integrated xSeries Server does not come with its own optical device. Therefore, it is likely that you would need to share your iSeries optical device with your Integrated xSeries Servers.

Before you can use iSeries tape and optical devices from a Linux server, you must make the devices available to Linux. An iSeries tape device must first be made unavailable (varied off) before a Linux server can use it, whereas an optical drive must be made available (varied on) before Linux can use it.

The latest information on backup device support is available at:

<http://www.ibm.com/systems/i/bladecenter/linux/backup.html> 

## Transferring control of an optical drive from i5/OS to Linux

### About this task

Sharing an iSeries optical drive with Linux involves transferring control of the optical drive from i5/OS to Linux, and transferring control back from Linux to i5/OS.

To transfer control of an iSeries optical drive from i5/OS to Linux with CL commands, complete the following steps:

1. On an i5/OS command line, use this CL command to check the status of the optical drives (you cannot check the status of an optical device using iSeries Navigator):

```
WRKCFGSTS *DEV *OPT
```

On the Work with Configuration Status display, find the optical device you want to transfer control of to Linux. Check that it is varied on. If not, type 1 next to the device and press **Enter**.

2. To lock the optical device to Linux, start a Linux terminal session and log in as root.
3. At the command prompt enter the `ixsdev` command and press Enter. The syntax is:

```
ixsdev [-list] | [[-lock | -unlock] device name]
```

For example, to list the tape and optical drives accessible by Linux enter the following command:

```
ixsdev -list
```

To lock OPT01 to Linux you would enter the following command:

```
ixsdev -lock OPT01
```

You could also use the Linux name:

```
ixsdev -lock /dev/scd0
```

The status of OPT01 shows as LOCKED.

**Note:** In Linux, commands are case sensitive. Make sure you use the same case as in the examples. For example, use OPT01, not opt01.

4. Because Linux sees optical drives as block devices, you need to mount the optical device (**NOTE:** If the optical device has been mounted, you will need to unmount the device before unlocking it.) :
  - a. Create a mount point using the Make directory (`mkdir`) command. For example:

```
mkdir /mnt/scd0
```
  - b. Mount the optical device. For example:

```
mount /dev/scd0 /mnt/scd0
```

## Results

After the optical drive has been logically switched to the integrated Linux server, you can use it in the same way as you would use an optical drive directly attached to a standalone Linux server.

### Transferring control of an optical drive from Linux to i5/OS

#### About this task

To transfer control of an iSeries optical drive back to i5/OS all you need to do is unlock it in Linux. Note that if you either shut down the integrated Linux server before you unlock the optical drive, or the server fails before you unlock the optical drive, it unlocks automatically.

To transfer control of an iSeries optical drive back to i5/OS from Linux, complete the following steps:

1. To unlock the optical drive from Linux, start a Linux terminal session and log in as root.
2. If the optical device is mounted, it should be unmounted. For example:

```
umount/dev/scd0
```

3. At the command prompt enter the `ixsdev` command and press **Enter**. The syntax is:

```
ixsdev [-list] | [[-lock | -unlock] device name]
```

To unlock OPT01 from Linux you would enter the following command:

```
ixsdev -unlock OPT01
```

**Note:** In Linux, commands are case sensitive. Make sure you use the same case as in the examples. For example, use `OPT01`, not `opt01`.

4. You could also use the Linux name:

```
ixsdev -unlock /dev/scd0
```

The status of OPT01 shows as UNLOCKED.

5. On an i5/OS command line, follow these steps:

- a. Use the following CL command to check that the optical drive is still varied on:

```
WRKCFGSTS *DEV *OPT
```

- b. On the Work with Configuration Status display, find the optical device you previously unlocked. Its status should be VARIED ON. If not, type 1 next to the device and press **Enter**.

## Results

You can now use the optical drive from i5/OS.

### Restricting iSeries tape and optical drives

This topic discusses how to restrict tape and optical drive usage for your integrated Linux server.

#### About this task

You can prevent a Linux server from using a particular tape or optical device by specifying the tape or optical device as a restricted resource in the server's NWSD.

### Restricting iSeries tape and optical drives from iSeries Navigator

#### About this task

Go to the server properties panel, then the **System** tab. Click the **Advanced** button. Go to the **Restricted Devices** tab.



## Restricting iSeries tape and optical drives using CL commands

### About this task

To restrict devices, complete the following steps:

1. Shut down the integrated Linux server.
2. Type WRKNWSD on the i5/OS command line. Press **Enter**. The Work with Network Server Descriptions display appears. Type 2 in the Opt column next to the NWSD that you want to change. Press **Enter**. The Change Network Server Desc display appears.
3. Optional: Type CHGNWSD on the command line. Press **Enter**. The Change Network Server Desc display appears. Enter the name of the NWSD. Press **Enter**.
4. Scroll down to the Restricted device resources parameter (RSTDEVRSC) and list the devices that are not to be made available. To specify more than two resources, type a + in the + for more values field and press **Enter**.

After you have entered your resources press **Enter**. Here is an example:

```
CHGNWSD NWSD(nwsd-name) RSTDEVRSC(OPT01 TAP02)
```

5. Restart the server.

### | Deleting Linux servers

| This topic discusses how to delete and remove integrated Linux servers.

### | About this task

| A CL command, called Delete Linux Server (DLTLNXSVR), was created for OS/400® V5R4 that removes the iSeries objects associated with an integrated Linux server. This command also deletes integrated Linux servers. The DLTLNXSVR command removes the:

- | • Network server description (NWSD) and associated line descriptions
- | • Point-to-Point Virtual Ethernet LAN line descriptions
- | • TCP/IP interfaces bound to Virtual Ethernet LAN line descriptions
- | • All system created storage spaces linked to the NWSD

| You cannot delete an integrated server using iSeries Navigator. This command is available through CL commands only.

| To run the delete command, enter:

```
| DLTLNXSVR NWSD(nwsd-name)
```

| If you no longer want to run integrated Linux or Windows servers on the iSeries, you can also delete the i5/OS Integrated Server Support (5722SS1 option 29 and 5722LSV).

| To remove the license program, enter two commands:

```
| DLTLICPGM LICPGM(5722SS1) OPTION(29)  
| DLTLICPGM LICPGM(5722LSV)
```

## Viewing messages for a Linux server

This topic discusses how to view messages from the integrated Linux server.

### About this task

The iSeries Integrated Server Support software issues messages for integrated Linux and Windows servers in the following ways:

- Through the server's monitor job.



- Through the QSYSOPR message queue.

Both of these sources can contain valuable information relating to Linux servers that can help you in problem determination.

You can view the job log of the server's monitor job using either iSeries Navigator or the i5/OS command line. Note that the server needs to be started before you can view the job log.

## Viewing the job log using iSeries Navigator

### About this task

To view the job log of the server's monitor job through iSeries Navigator, complete the following steps:

1. Launch iSeries Navigator from your Windows desktop.
2. Click **iSeries connection** → **Work Management**
3. Click **Active Jobs**.
4. Scroll down the list of jobs in the right-hand window until you see the jobs running in the QSYSWRK subsystem. Look for a job with the same name as the NWSD of the Linux server.
5. Right-click the name of the server whose job log you want to view. Select **Job Log**.
6. You can view detailed information for any of the messages shown by double-clicking a message ID.

## Viewing the job log using CL commands

### About this task

To view the job log of the server's monitor job through the i5/OS command line, complete the following steps:

1. Type WRKACTJOB (Work with Active Jobs command) on an i5/OS command line. Press **Enter**.
2. Scroll down until you see the jobs running in the QSYSWRK subsystem. Look for a job with the same name as the NWSD of the Linux server. Type 5 next to the server's monitor job.
3. Type 10 to display the job log. Press **Enter**.
4. Press **F10** to display detailed messages.
5. Scroll up and down through the messages by pressing the Page up and Page down keys.

## Viewing QSYSOPR message queue using iSeries Navigator

### About this task

To view the QSYSOPR message queue through iSeries Navigator, complete the following steps:

1. Launch iSeries Navigator from your Windows desktop.
2. Click **iSeries connection** → **Basic Operations**
3. Before you can view the QSYSOPR message queue through iSeries Navigator, you need to make QSYSOPR the default queue.
  - a. Right-click **Messages** and select **Customize this view** → **Include**
  - b. Click the button with the three dots (...).
  - c. Click the **System operator** radio button. Click **OK**.
  - d. The system operator's message queue (QSYSOPR) is now the default message queue.
4. Click **Messages** in the left pane to view the QSYSOPR message queue.
5. Double-click an entry in the From user column to view detailed information for a message.

## Viewing QSYSOPR message queue using CL commands

### About this task

To view the QSYSOPR message queue through the i5/OS command line, complete the following steps:

**Note:** The server needs to be started before you can view the QSYSOPR message queue.

1. Type DSPMSG QSYSOPR (Display Messages for QSYSOPR) on an i5/OS command line. Press **Enter**.
2. You can scroll through the messages by pressing the Page up and Page down keys. Note that the most recent messages are at the bottom.

To view detailed information for a message, use the up and down arrow keys to move the cursor to the message you want to see and press **F1**.

---

## Linux-centric backup and recovery

This topic introduces the techniques you can use to perform Linux-centric backup and recovery on an Integrated xSeries Server or Integrated xSeries Adapter running Linux.

### Linux backup and recovery overview

This topic discusses Linux backup and recovery.

Linux-centric backup is inherently file oriented in nature because Linux backup utilities operate at a file level. It is very difficult to back up an entire Linux drive as a single entity using Linux backup utilities. The only way to save a complete copy of a disk volume using Linux is to buy a third-party imaging product. These products require special skills to use, and are not supported for use with integrated Linux servers.

It should be noted that when we discuss Linux file-level backup and the applications that can be used to perform this task, we are really talking about Linux flat files, that is, non-database files. In i5/OS, we usually do not draw such a distinction because we use the same backup tools to save both database and non-database files. In Linux, however, you would usually back up database files using a backup application that is specific to the installed database. Therefore, when we discuss file-level backup in this section it is in the context of flat, or non-database, files.

There are other differences between an i5/OS and Linux backup. When performing an i5/OS backup, we usually put the partition to restricted state and save objects direct to tape. While the data can be compressed, the objects are written sequentially on the tape and not saved to a single container (archive in Linux terms). When backing up to tape, this is the most logical and efficient way to perform a save. In the Linux world, however, a backup usually means saving files to a single archive file on disk. The archive file can then either be copied to another Linux server in the network for safekeeping, or dumped to tape.

When we discuss backup we usually assume that we are backing up to tape. However, we can also back up to disk. Backing up to disk is usually an intermediate step before saving to tape in a staged backup. Ultimately we must back up to tape, or to disk on another system to guard against a complete loss of the primary system. While backup to tape is fairly straightforward, there are applications and techniques that can also be used to save data to disk on remote systems so that files can be quickly recovered in the event of a data loss.

Although Linux documentation might talk about “disaster recovery” backup, we need to draw a distinction between the disaster recovery backup capability provided by i5/OS for an integrated Linux server, and disaster recovery in the context of a Linux backup application. From the Linux point of view, a disaster recovery backup is a backup of all files on a disk volume or volumes versus backing up selected files. However, Linux backup applications still save at a file level, and you can still restore individual files from a Linux “disaster recovery” backup. In contrast, an i5/OS storage space backup

saves a complete image of a Linux drive as a single entity, and you cannot directly restore individual files from it. You can restore the storage space to a new name and link it to the same server, then copy the needed files from it. However, i5/OS storage space backup is very fast, and can provide additional functions not readily available with a Linux "disaster recovery" backup, such as the ability to quickly restore a complete copy of the system drive. In the context of this information, we regard all Linux backup operations as file-level backups.

**Note:** When using a Linux backup application you can use either a native iSeries tape drive or a standalone Linux server-attached tape drive to save and restore Linux files. Both i5/OS and Linux cannot use an iSeries tape drive concurrently; it must be allocated or "locked" to one operating system or the other.

Note that the tape drive must be varied off in the hosting i5/OS partition. Tape devices show up under the /dev directory, for example, /dev/st0 or /dev/st1.

Using a Linux application to perform file-level backups can be difficult to incorporate into an unattended backup from the i5/OS side. This is because Linux cannot share a tape cartridge formatted for use by i5/OS.

**Note:** Tape libraries and automatic cartridge loaders (ACLs) on the supported devices list are not supported in random access mode. However, they are supported in manual or sequential access mode.

There are some good reasons why you might want to use an iSeries tape drive in preference to an Integrated xSeries Server or xSeries-attached tape drive. iSeries tape drives tend to be very fast, reliable, robust and high capacity; and it can be possible to consolidate a number of tape drives in your organization down to one or two iSeries devices. If you have multiple integrated Linux servers in the same i5/OS partition, they can all access the same tape drive (although not at the same time). Therefore, you might only need one iSeries tape drive to back up all your Linux servers.

You can find more information about Linux backup strategies and automating the backup process, along with general information about backup and restore, at the following Web site:

<http://www.backupcentral.com/> 

## Linux recovery options

This topic discusses different Linux recovery options.

In the case where a Linux server fails to start, there are options that allow you to boot from external media and potentially recover the server. These Linux recovery functions include:

- Rescue diskette
- Recovery mode CD-ROM

The Rescue diskette and Recovery mode CD-ROM functions are intended to provide you with a means of recovering a Linux server that fails to start. While it can be possible to use either of these techniques to recover an integrated Linux server, they are designed for standalone servers. Note that you could only use these techniques with an Integrated xSeries Adapter-attached xSeries server because the xSeries can boot directly from the diskette and CD-ROM drives. An Integrated xSeries Server card does not have its own diskette or CD-ROM drives and therefore cannot boot from the Rescue diskette or Recovery mode CD-ROM.

If you are unable to recover a failed standalone Linux server then your only other alternative is to rebuild it. In the case of an integrated Linux server, you should rarely, if ever, need to either use the Linux recovery options or rebuild the server because you can save a complete image of the server using i5/OS

storage space backup. In this case, all you need to do to recover a failed server is to restore a previously saved copy of the server's drives (storage spaces) and then restart it. Note that you might need to also restore volatile files from your file-level backup to make sure that the data is up-to-date.

You can effectively use these techniques to eliminate the need for the Rescue diskette and Recovery mode CD-ROM options that are available to recover a standalone Linux server.

**Note:** The ability to quickly and easily recover a failed Linux server is one of the major benefits of the Linux integration support.

## Choosing a tape drive for use by your Linux backup application

This topic discusses the different tape drives to be used for a Linux backup.

Linux backup applications can either save to an iSeries tape drive or a tape drive directly controlled by a Linux server somewhere in the network.

- Backing up to a native iSeries tape drive

iSeries tape drives can be accessed by backup applications running on the integrated Linux server.

Because an iSeries tape drive appears to an integrated Linux server as a directly connected tape drive, you can also save remote integrated or non-integrated Linux servers to the iSeries tape drive across the network using a utility such as `rsync`. In other words, the iSeries tape drive behaves exactly as if it were a native Linux tape drive.

- Backing up to a native Linux tape drive

If you have a Linux-based backup infrastructure already in place, it is likely that you have tape drives attached to standalone Linux servers in the network. Therefore, you can save files to a tape drive attached to an integrated Linux server in the same way as you would save files on any other standalone Linux server in the network.

## Restricting iSeries tape drives that can be used by Linux

This topic discusses how to restrict Linux to specific tape drives.

### About this task

You can optionally restrict which tape and optical drives can be allocated to integrated Linux servers in the hosting i5/OS partition. You might want to do this in order to reserve certain drives for use by i5/OS only. Tape and optical drives that are not to be made available to integrated Linux servers can only be specified after installation has been completed.

- | For IXS or IXA solutions, the installation of RHEL3 requires at least one optical drive that is not restricted.
- | For iSCSI attached solutions, you must restrict any tape or optical device that is not available, such as:
  - | • External devices that are not powered on
  - | • Any device that has been physically removed but still has a device description
- | If you do not restrict these devices, other devices (including disk, optical, and tape devices) might fail to report on your Linux server.

To restrict devices, complete the following steps:

1. Shut down the integrated Linux server.
2. On an i5/OS command line, enter the Change Network Server Description (`CHGNWSD`) command and press **F4**.
3. Scroll down to the Restrict device resources (`RSTDEVRSC`) parameter and list the devices not to be made available. Press **Enter**.

4. Restart the server.

## Setting up an iSeries tape drive for use by Linux

This topic describes the tasks you need to perform to set up an iSeries tape drive for use by an integrated Linux server.

### About this task

The iSeries supports a wide range of tape drives. If you have multiple tape drives on your iSeries, each one can be allocated separately to either i5/OS or Linux.

A number of iSeries tape drives have been tested for use with integrated Linux servers, but some models cannot be used. The latest information about which tape devices have been tested for use can be found at the following Web sites:

For IXA and IXS solutions, see [http://www.ibm.com/systems/i/bladecenter/linux/tape\\_support.html](http://www.ibm.com/systems/i/bladecenter/linux/tape_support.html)



For iSCSI attached solutions, see <http://www.ibm.com/systems/i/bladecenter/linux/>

[iscsi\\_tape\\_support.html](#)

To allocate a native iSeries tape drive for use by an integrated Linux server, the iSeries tape drive must be logically detached (varied off) from i5/OS. Then it must be logically attached (locked) to the integrated Linux server so that Linux thinks it has a physical tape drive directly attached. When this has been done, the integrated Linux server can use the iSeries tape drive as if it were a directly attached tape drive.

**Tip:** For IXS and IXA-attached solutions, the Linux integration support software supports the use of Automatic Cartridge Loaders (ACLs) in manual and automatic modes. This means that if the ACL is in automatic mode when a Linux backup application ejects a full tape, the media is unloaded and the ACL automatically loads the next tape. However, the first tape must always be loaded manually by the user.

To enable an iSeries tape drive for use by Linux backup utilities and applications, complete the following steps:

1. Format tape media for use by Linux.
2. Transfer control of a tape drive from i5/OS to Linux.
3. Transfer control of a tape drive from Linux to i5/OS.

## Formatting tape media for use by Linux

This topic discusses how to format tape media.

### About this task

The tape media formats used by i5/OS and Linux are mutually exclusive; i5/OS uses labels, Linux requires a non-labelled tape. Therefore i5/OS and Linux servers cannot share the same tape media.

**Note:** All new tapes must be initially formatted using the Initialize Tape (INZTAP) CL command. After the tape has been formatted using INZTAP, additional formatting can be done by Linux, if required.

From an i5/OS command line, enter the Initialize Tape (INZTAP) CL command:

```
INZTAP DEV(TAP01) NEWVOL(*NONE) NEWOWNID(*BLANK) VOL(*MOUNTED) CHECK(*NO) DENSITY(*DEVTYPE) CODE(*EBCDIC)
```

TAP01 is the name of the tape device. Yours might be different.

DENSITY(\*DEVTYPE) gives the best performance; however, if you receive an error with this setting when trying to initialize the tape cartridge, try using DENSITY(\*CTGTYPE). If you still receive an error, the tape cartridge is incompatible with your tape drive.

This command produces a non-labelled tape that can be used by Linux backup applications. Note that if tapes are used that are not of the default density for the drive (consult your drive's documentation), you need to reset the tape density after the iSeries partition has been restarted.

To reset the tape density after an iSeries partition has been restarted, complete the following steps:

1. Put a spare tape in the drive. Note that the next step erases all the data on the tape.
2. Issue the following command:  
`INZTAP DEV(tape-device-name) CHECK(*NO) Density(*CTGTYPE).`
3. You can now switch the blank tape with the tape that you want to use for Linux backups. Backup applications should now work normally. Failure to initialize a blank tape of the correct density can have unanticipated results. If you regularly switch tape densities you might need to repeat the above steps more often than just after an IPL.

## Transferring control of a tape drive from i5/OS to Linux

This topic discusses how to transfer control of a tape drive from i5/OS to Linux.

### About this task

Before you use an iSeries tape drive for a Linux backup application, you must make it unavailable from the i5/OS side using iSeries Navigator or a CL command, and then lock it on the Linux side through a Linux terminal session.

Note that some tape devices report in under more than one device description. Tape libraries (3570, 358x, 3590, and so on) report in as tape libraries (TAPMLBxx) as well as tape devices (TAPxx), where xx is a sequence number. The Linux integration support software does not support the tape library function. Therefore, if your device has a tape library description, both the tape and tape library devices must be made unavailable (varied off) before locking the device on the Linux server. Note that, although tape libraries are not supported as libraries in Linux, you can use them in sequential mode if the drive supports it.

If you have multiple integrated Linux servers being hosted by the same i5/OS partition, only one server at a time can use a particular iSeries tape drive. If you have multiple logical partitions on your iSeries, a tape drive that is owned by one partition cannot be shared by integrated Linux servers that are being hosted by other partitions. Note, however, that it might be possible to logically switch tape drives between i5/OS partitions, depending on the hardware configuration of the iSeries.

To transfer control of a tape drive from an i5/OS partition to an integrated Linux server, you must have i5/OS Administrator or Backup Operator authority.

To transfer control of an iSeries tape drive from i5/OS to Linux, choose one of the two following methods:

## Transferring control of a tape drive with iSeries Navigator

### About this task

To transfer control of an iSeries tape drive with the iSeries Navigator, complete the following steps:

1. From an iSeries Navigator window, expand the i5/OS partition you are working with.
2. Click **Configuration and Service** → **Expand** → **Hardware** → **Tape Devices**.



3. Click **Stand-Alone Devices** and then right-click the tape device you want to transfer control of to Linux. Select **Make Unavailable**.
4. If the tape device is also a tape library, click **Tape Libraries** and then right-click the tape library you want to transfer control of to Linux. Select **Make Unavailable**.
5. To lock the tape device to Linux, start a Linux terminal session and log in as root.
6. At the command prompt enter the `ixsdev` command and press **Enter**. The syntax is:  
`ixsdev [-list] | [[-lock | -unlock] device name]`  
 For example, to list the tape and optical drives accessible by Linux enter the following command:  
`ixsdev -list`
7. To lock TAP02 to Linux you would enter the following command:  
`ixsdev -lock TAP02`

**Note:** In Linux, commands are case sensitive. Make sure you use the same case as in our examples, for example, use TAP02 not tap02.

You could also use the Linux name. `ixsdev -lock /dev/st1` for a rewindable tape device, or `ixsdev -lock /dev/nst1` for a non-rewindable tape device. The `ixsdev -list` command shows the status of TAP02 as LOCKED.

Note that you do not need to mount the tape device because Linux does not see it as a block device. You only need to mount block devices.

8. Insert a tape cartridge that has been formatted for Linux.

## Transferring control of a tape drive with CL commands

### About this task

To transfer control of an iSeries tape drive with CL commands, complete the following steps:

1. On the i5/OS command line, use this command to vary off the tape device:  
`WRKCFGSTS *DEV *TAP`
2. On the Work with Configuration Status display, find the tape device you want to transfer control of to Linux. Type 2 next to the device and press **Enter**.
3. If the tape device is also configured as a tape library, enter the following command:  
`WRKCFGSTS *DEV *TAPMLB`
4. On the Work with Configuration Status display, find the tape library corresponding to the tape device you want to transfer control of to Linux. Type 2 next to the tape library and press **Enter**.
5. To lock the tape device to Linux, start a Linux terminal session and log in as root.
6. At the command prompt enter the `ixsdev` command and press **Enter**. The syntax is:  
`ixsdev [-list] | [[-lock | -unlock] device name]`  
 For example, to list the tape and optical drives accessible by Linux enter the following command:  
`ixsdev -list`
7. To lock TAP02 to Linux you would enter the following command:  
`ixsdev -lock TAP02`

**Note:** In Linux, commands are case sensitive. Make sure you use the same case as in our examples, for example, use TAP02 not tap02.

You could also use the Linux name. `ixsdev -lock /dev/st1` for a rewindable tape device, or `ixsdev -lock /dev/nst1` for a non-rewindable tape device. The `ixsdev -list` command shows the status of TAP02 as LOCKED.

Note that you do not need to mount the tape device because Linux does not see it as a block device. You only need to mount block devices.

8. Insert a tape cartridge that has been formatted for Linux.

## What to do next

After the tape drive has been logically switched to the integrated Linux server, you can use it in the same way as you would use a tape drive directly attached to a standalone Linux server. Using a Linux backup application, you can now direct your Linux backups to the iSeries tape drive.

## Transferring control of a tape drive from Linux to i5/OS

This topic discusses how to transfer control of a tape drive from Linux to i5/OS.

### About this task

To hand control of the tape drive back to i5/OS, unlock it on the Linux side and then make it available on the i5/OS side. This procedure is simply the reverse of the process you used to pass control of the tape drive to Linux.

Note that if you shut down the integrated Linux server, or the Linux server fails before you unlock the tape drive, it unlocks automatically. However, it is still in an unavailable state in i5/OS.

To transfer control of a tape drive from Linux to i5/OS, use one of the following methods:

### Transfer control of a tape drive with Linux commands

#### About this task

To transfer control of an iSeries tape drive back to i5/OS from Linux with Linux commands, complete the following steps:

1. To unlock the tape drive from Linux, start a Linux terminal session and log in as root.
2. At the command prompt enter the `ixsdev` command and press **Enter**. The syntax is:

```
ixsdev [-list] | [[-lock | -unlock] device name]
```

To unlock TAP02 from Linux you would enter the following command:

```
ixsdev -unlock TAP02
```

You could also use the Linux name:

```
ixsdev -unlock /dev/st1
```

**Note:** In Linux, commands are case sensitive. Make sure you use the same case as in our examples, for example, use TAP02 not tap02.

### Transfer control of a tape drive with iSeries Navigator

#### About this task

To transfer control of an iSeries tape drive back to i5/OS with the iSeries Navigator, complete the following steps:

1. From an iSeries Navigator window, expand the i5/OS partition you are working with.
2. Click **Configuration and Service** → **Hardware** → **Tape Devices**.
3. Click **Stand-Alone Devices** and then right-click the tape device you want to transfer control of to i5/OS. Select **Make Available**.
4. If the tape device is also a tape library, click **Tape Libraries** and then right-click the tape library you want to transfer control of to i5/OS. Select **Make Available**.

### Transfer control of a tape drive with CL commands

#### About this task

To transfer control of an iSeries tape drive back to i5/OS with CL commands, complete the following steps:



1. On an i5/OS command line, use the following command to vary on the tape device:  
`WRKCFGSTS *DEV *TAP`
2. On the Work with Configuration Status display, find the tape device you want to transfer control of to i5/OS. Type 1 next to the device and press **Enter**.
3. If the tape device is also configured as a tape library, enter the following command:  
`WRKCFGSTS *DEV *TAPMLB`
4. On the Work with Configuration Status display, find the tape library you want to transfer control of to i5/OS. Type 1 next to the library and press **Enter**.

## Results

You can now use the tape drive from i5/OS.

## Backing up files using Linux utilities and applications

This topic discusses backup and recovery using Linux utilities.

At the time of writing, not all the applications listed here had been tested for use with iSeries tape drives. However, even though a particular Linux backup application has not been tested for use with iSeries tape drives, this does not mean that the application will not work. You need to test your backup application to determine if it is compatible with the tape drive installed on your iSeries server.

When performing a backup you should close any operation that is using files on the Linux file system or the directory you intend to back up. Though native Linux backup utilities do back up files that are in use, we recommend that you close all open files and the operations using them so as to preserve the latest file updates and avoid any data corruption.

If the file system is network file system (NFS) then file locking occurs when the file residing on NFS is in use by an operation. Creating a tar or cpio backup archive can fail if you try to include such files. Commercial applications like IBM Tivoli® Storage Manager (TSM), ARCserve, Veritas, or Legato provide better facilities for backing up data as they provide open file agents. Even so, it would be safe practice to stop the operation that is locking the file over NFS before backing it up.

## Determining the tape device block size

Before you begin your backup using a Linux utility or application, you need to determine the block size of the tape device you are using. To do this you run the `mt` command, for example:

```
mt -f /dev/st1 status
```

Replace `/dev/st1` with the device name of your tape drive.

## Setting the tape device block size

The `ixsdev` command will attempt to set the block size to 32768 bytes when the device is locked. If the `ixsdev` command is unable to set the block size, you can use the `mt` command to set it after the media is inserted.

For example, you can run the following command at the Linux console:

```
mt -f /dev/st0 setblk 32768
```

Replace `/dev/st0` with the device name of your tape drive. Do not use a block size larger than 65,536 bytes.

---

## Troubleshooting integrated Linux servers

Resolve problems with integrated Linux servers.

- | See the System i integration with BladeCenter and System x category in the System i Support Knowledge Base for information about resolving problems with integrated Linux servers.




---

## Related information

Listed here are the product manuals and IBM Redbooks (in PDF format), Web sites, and information center topics that relate to the Linux on an integrated xSeries solution topic. You can view or print any of the PDFs.


**Note:** See “Printable PDF” on page 2 for a printable PDF version of this documentation.

### IBM Redbooks

- | • Implementing Linux on Integrated xSeries Solutions for iSeries 
- | • Linux Handbook: A Guide to IBM Linux Solutions and Resources 
- | • Linux and Windows Integration for IBM eServer xSeries Servers 

For additional IBM Redbooks, see the IBM Redbook Web site: <http://www.redbooks.ibm.com/> 

### Web sites

- System i integration with BladeCenter and System x  (<http://www.ibm.com/systems/i/bladecenter/linux/>)  
Web site featuring the latest product technical and marketing information.
- | • iSCSI install read me first (<http://www.ibm.com/systems/i/bladecenter/iscsi/readme/>)  
Instructions for planning for and installing an iSCSI attached integrated server.
- | • System i integration with BladeCenter and System x category in the System i Support Knowledge Base  
Information about resolving problems with integrated Linux servers.
- IBM Linux Portal - Linux at IBM  (<http://www.ibm.com/linux/>)  
Additional information on integrating Linux into your business and on demand world, including solutions, hardware, and software information.
- The Linux Documentation Project  (<http://www.tldp.org/>)  
An open source documentation support project for Linux on multiple platforms. This site is not affiliated with IBM.

### Other information


- Windows environment on iSeries
- Red Hat Enterprise Linux installs the documentation available with rpm packages into the `/usr/share/doc/<package_name>` directory.
- If you are not able to find a particular command then you can use the `whereis` command to find that particular command, its source, and its man page location. This will work if you have the package installation path specified in the system environment variables.
- If you know the first few letters of the command you are trying to remember then type those letters on the command prompt and press the Tab key twice, and you will get a complete list of commands starting with that sequence of letters.
- Man pages are the best source of documentation for your Linux system. Type `man` and the command name as its argument to find more information for it.

- Almost every command provides the --help option, which you can use to view other options provided by these commands.

## | **Saving PDF files**

- | To save a PDF on your workstation for viewing or printing:
  - | 1. Right-click the PDF in your browser (right-click the link above).
  - | 2. Click the option that saves the PDF locally.
  - | 3. Navigate to the directory in which you want to save the PDF.
  - | 4. Click **Save**.

## | **Downloading Adobe Reader**

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

### **Related reference**

“Printable PDF” on page 2  
Use this to view and print a PDF of this information.



---

## Appendix. Notices

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

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

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

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

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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

IBM Corporation

Software Interoperability Coordinator, Department 49XA  
3605 Highway 52 N  
Rochester, MN 55901  
U.S.A.

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

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

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

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

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

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

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

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

#### COPYRIGHT LICENSE:

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

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

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

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

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

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

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

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

---

## Trademarks

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

AIX  
AIX 5L  
BladeCenter  
e(logo)server  
eServer  
i5/OS  
IBM  
iSeries  
OS/400  
pSeries  
System i  
System i5  
System x  
xSeries  
zSeries

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

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

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

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

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

---

## Terms and conditions

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

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

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

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

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

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

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







Printed in USA